

The Anatomy of a Cryptocurrency Pump-and-Dump Scheme

Jiahua Xu

École Polytechnique Fédérale de Lausanne (EPFL)
Imperial College London
Harvard University

Benjamin Livshits

Imperial College London
UCL Centre for Blockchain Technologies
Brave Software

Abstract

While pump-and-dump schemes have attracted the attention of cryptocurrency observers and regulators alike, this paper represents the first detailed empirical query of pump-and-dump activities in cryptocurrency markets. We present a case study of a recent pump-and-dump event, investigate 412 pump-and-dump activities organized in Telegram channels from June 17, 2018 to February 26, 2019, and discover patterns in crypto-markets associated with pump-and-dump schemes. We then build a model that predicts the pump likelihood of all coins listed in a crypto-exchange prior to a pump. The model exhibits high precision as well as robustness, and can be used to create a simple, yet very effective trading strategy, which we empirically demonstrate can generate a return as high as 60% on small retail investments within a span of two and half months. The study provides a proof of concept for strategic crypto-trading and sheds light on the application of machine learning for crime detection.

1 Introduction

While pump-and-dump schemes are a well-trodden ruse in conventional financial markets, the old-fashioned ploy has found a new playground to thrive — cryptocurrency exchanges.

The relative anonymity of the crypto space has led to it becoming a fertile ground for unlawful activities, such as currency theft (e.g. the DAO hack [1]), Ponzi schemes [26], and pump-and-dump schemes that have each risen in popularity in cryptocurrency markets over the last few years. Due to their end-to-end encryption, programmability, and relative anonymity, new social media tools such as Telegram¹ and Discord have become cryptocurrency enthusiasts' preferred communication vehicles. While pump-and-dump schemes have been discussed in the press [29], we are not aware of a comprehensive study of this phenomenon to date.

¹Note that not all Telegram traffic is end-to-end encrypted.

Regulation: In February 2018, the CFTC (Commodity Futures Trading Commission) issued warnings to consumers [8] about the possibility of cryptocurrency pump-and-dump schemes. It also offered a substantial reward to whistleblowers around the same time [12].

In October 2018, the SEC (Securities and Exchange Commission) filed a subpoena enforcement against an investment company trust and trustee for an alleged pump-and-dump ICO scheme [27].

Clearly, regulators are aiming to find perpetrators of pump-and-dump schemes and to actively prosecute them.

This paper: In this paper, we trace the message history of over 300 Telegram channels from June 17, 2018 to February 26, 2019, and identify 412 *pump events* orchestrated through those channels. We analyze features of pumped coins and market movements of coins before, during, and after pump-and-dump. We develop a predictive random forest model that provides the likelihood of each possible coin being pumped *prior* to the actual pump event. With an AUC (area under curve) of the ROC (receiver operating characteristic) curve of over 0.9, the model exhibits high accuracy in predicting pump-and-dump target coins.

Contributions: This paper makes the following contributions:

- **Longitudinal study:** This paper is the *first* research study that examines routinely organized pump-and-dump events in the cryptocurrency space. We use a unique dataset of pump-and-dump records from June 17, 2018 to February 26, 2019 across multiple crypto-exchanges and analyze crypto-market movements associated with those pump-and-dump events.
- **Analysis:** Our analysis shows that pump-and-dump activities are a lot more prevalent than previously believed. Specifically, around 100 organized Telegram pump-and-dump channels coordinate on average 2 pumps a day, which generates an aggregate artificial trading volume of 6 million USD a month. We discover that some ex-




Figure 1: A successfully organized pump event. On the right hand side of the screenshot is the message history of a Telegram channel. The first message is the final countdown; the second message is the coin announcement; the last message presents the pump result. On the left hand side is the market movement of the corresponding coin around the pump time.

changes are also active participants in pump-and-dump schemes.

- **Prediction:** We develop machine learning models that, given pre-pump market movements, can predict the likelihood of each coin being pumped with an AUC (Area Under Curve) of over 0.9 both in-sample and out-of-sample. The models confirm that market movements contain hidden information that can be utilized for monetary purposes.
- **Trading strategy:** We formulate a simple trading strategy which, when used in combination with a calibrated prediction model, demonstrates a return of 60% over a period of three weeks, even under strict assumptions.

Paper organization: The paper is structured as follows. In [Section 2](#) we provide background information on pump-and-dump activities organized by Telegram channels. In [Section 3](#) we present a pump-and-dump case study. In [Section 4](#) we investigate a range of coin features. In [Section 5](#) we build a prediction model that estimates the pump likelihood of each coin for each pump, and propose a trading strategy along with the model. In [Section 6](#) we summarize the related literature. In [Section 7](#) we outline our conclusions. Finally, the Appendix specifies parameters of the models we have used in this paper.

2 Background

A pump is a coordinated, intentional, short-term increase in the demand of a market instrument — in our study, a cryptocurrency — which leads to a price hike. With today’s chat applications such as Telegram and Discord offering features of encryption and anonymity, various forms of misconduct in cryptocurrency trading are thriving on those platforms.

2.1 Pump-and-Dump Actors

Pump organizer: Pump organizers can be individuals, or, more likely, organized groups, typically who use encrypted chat applications to coordinate pump-and-dump events. They have the advantage of having insider information and are the ultimate beneficiaries of the pump-and-dump scheme.

Pump participants: Pump participants are cryptocurrency traders who collectively buy a certain coin immediately after receiving the instruction from the pump organizer on which coin to buy, causing the price of the coin to be “pumped”. Many of them end up buying coins at an already inflated price and are the ultimate victim of the pump-and-dump scheme.

Pump target exchange: A pump target exchange is the exchange selected by the pump organizer where a pump-and-dump event takes place. Some exchanges are themselves directly associated with pump-and-dump. Yobit, for example, has openly organized pumps multiple times (see [Figure 2](#)). The benefits for an exchange to be a pump organizer are threefold:

1. With coins acquired before a pump, it can profit by dumping those coins at a higher, pumped price;
2. It earns high transaction fees due to increased trading volume driven by a pump-and-dump;
3. Exchanges are able to utilize their first access to users’ order information for front-running during a frenzied pump-and-dump.

2.2 A Typical Pump-and-Dump Process

Set-up: The organizer creates a publicly accessible group or channel, and recruits as many group members or channel subscribers as possible by advertising and posting invitation links on major forums such as Bitcointalk, Steemit, and Reddit.

Telegram *channels* only allow subscribers to receive messages from the channel admin, but not post discussions in the channel. In a Telegram *group*, members can by default post messages, but this function is usually disabled by the group admin to prohibit members’ interference. We use the terms *channel* and *group* interchangeably in this paper.

Pre-pump announcement: The group is ready to pump once it obtains enough members (typically above 1,000). The pump organizer, who is now the group or channel admin, announces details of the next pump a few days ahead. The admins broadcast the exact time and date of the announcement of a coin which would then precipitate a pump of that coin. Other information disclosed in advance includes the exchange where the pump will take place and the pairing coin². The admins advise members to transfer sufficient funds (in the form of the pairing coin) into the named exchange beforehand.

While the named pump time is approaching, the admin sends out countdowns, and repeats the pump “rules” such as: 1) buy fast, 2) “shill”³ the pumped coin on the exchange chat box and social media to attract outsiders, 3) “HODL”⁴ the coin at least for several minutes to give outsiders time to join in, 4) sell in pieces and not in a single chunk, 5) only sell at a profit and never sell below the current price. The admin also gives members a pep talk, quoting historical pump profits, to boost members’ confidence and encourage their participation.

Pump: At the pre-arranged pump time, the admin announces the coin, typically in the format of an OCR (optical character recognition)-proof image to hinder machine reading (Figure 1). Immediately afterwards, the admin urges members to buy and hold the coin in order to inflate the coin price. During the first minute of the pump, the coin price surges, sometimes increasing several fold.

Dump: A few minutes (sometimes tens of seconds) after the pump starts, the coin price will reach its peak. While the admin might shout “buy buy buy” and “hold hold hold” in the channel, the coin price keeps dropping. As soon as the first fall in price appears, pump-and-dump participants start to panic-sell. While the price might be re-boosted by the second wave of purchasers who buy the dips (as encouraged by channel admins), chances are the price will rapidly bounce back to the start price, sometimes even lower. The coin price declining to the pre-pump proximity also signifies the end of the dump, since most investors would rather hold the coin than sell at a loss.

Post-pump review: Within half an hour, after the coin price and trading volume recover to approximately the pre-pump levels, the admin posts a review on coin price change, typically including only two price points – start price (or low price) and peak price, and touts how much the coin price increased by


²A pairing coin is a coin that is used to trade against other coins. Bitcoin (BTC) is a typical pairing coin.

³Crypto jargon for “advertise”, “promote”.

⁴Crypto jargon for “hold”.



(a) Tweets from @YobitExchange.



(b) Pump timer from the Yobit website.

Figure 2: The screen-shots demonstrate that the exchange Yobit was actively involved in pump-and-dump activities.

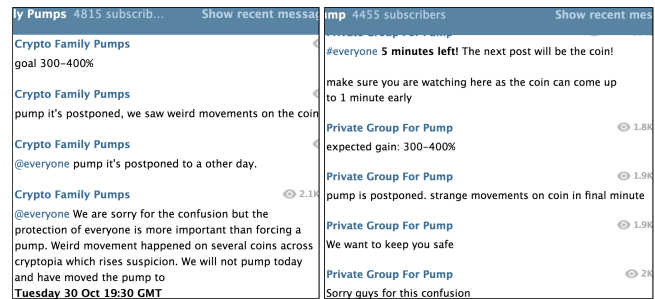


Figure 3: A pump attempt coordinated by multiple channels not executed due to unanticipated price movement of the to-be-pumped coin.

the pump (Section 2). Information such as trading volume and timescale is only selectively revealed: if the volume is high, and the pump-and-dump lasts a long time (over 10 minutes, say, would be considered “long”), then those stats will be “proudly” announced; if the volume is low or the time between coin announcement and price peak is too short (which is often the case), then the information is glossed over. Such posts give newcomers, who can access channel history, the illusion that pump-and-dumps are highly profitable.

Failed pump-and-dump attempts: Note that not every pump attempt is successful. Figure 3 shows that the admins decided not to carry through a pre-announced pump due to unanticipated price movements of the to-be-pumped coin.

While it is unknown what caused these movements, the case evidences that the admin is aware of the coin choice before the pump (as opposed to the coin being randomly selected and immediately announced at the pump time purely by algorithm), and hence has the time advantage of hoarding the coin at a low price before the coin announcement, whereas group members only purchase the coin after the coin announcement and slow buyers risk acquiring the coin at an

already (hyper)inflated price. It is generally known to pump participants that admins benefit the most from a pump. So why are there still people enthusiastic about partaking a pump, given the risk of being ripped off by the admins? Because people believe that they can sell those coins at an even higher price to other “greater fools”. The greater fool theory also forms the foundation of many other schemes, such as pyramid scams or Ponzi games [5].

One may also hypothesize that in this case, someone might have worked out the pattern of the coin selection and pre-purchased a basket of coins with high pump likelihood that happens to contain the actual to-be-pumped coin, which might explain why the admin observed peculiar movements of the coin. In the next section, we study the features of pumped coins and their price movements to understand if it is indeed possible to predict the to-be-pumped coin.

2.3 Regulatory and Ethical Considerations

Pump-and-dumps in the stock market nowadays typically involve penny stock manipulation employing deceptive campaigns on social media to amass gains and are deemed criminal [27]. However, since many cryptocurrencies cannot be neatly classified as investment or consumer products [22], the applicability of certain securities laws might be ambiguous, and to date, regulation of pump-and-dumps in the cryptocurrency market is still weak [23].

Yet, the crypto-market is likely to be considered subject to common law and general-purpose statues even though it has not been clearly regulated as either a securities market or a currency market. While offenses of market manipulation can depend on a defined market, outright fraud and deception do not. As pump-and-dump admins create information asymmetry by not showing investors the full picture of their scheme, they intentionally mislead investors for their own financial benefit. As a consequence, when it comes to US legislation, for instance, admins might be committing false advertising under the FTC (Federal Trade Commission) Act (15 USC §45) [15] or fraudulent misrepresentation. Of course, practically speaking, these admins are frequently outside of the US jurisdiction.

Pump-and-dump admins, aiming to profit from price manipulation, are certainly unethical. Nevertheless, other pump-and-dump participants are also culpable since their behaviour enables and reinforces the existence of such schemes; ironically, most participants become the victim of their own choices.

3 A Pump-and-Dump Case Study

We further study in depth the pump-and-dump event associated with Figure 1. The pump-and-dump was organized by at least four Telegram channels, the largest one being **Official McAfee Pump Signals**, with a startling 12,333 members. Prior to the coin announcement, the members were notified




Figure 4: Tick-by-tick movement of the BVB/ BTC market during the first four minutes after the coin announcement.

that the pump-and-dump would take place on one of the Cryptopia’s BTC markets (i.e., BTC is the pairing coin).

Announcement: At 19:30 GMT, on November 14, 2018, the channels announced the target coin in the form of a OCR-proof picture, but not quite simultaneously. **Official McAfee Pump Signals** was the fastest announcer, having the announcement message sent out at 19:30:04. **Bomba bitcoin “cryptopia”** was the last channel that broadcast the coin, at 19:30:23.

The target coin was BVB, a dormant coin that is not listed on CoinMarketCap. The launch of the coin was announced on Bitcointalk on August 25, 2016.⁵ The coin was claimed to be have been made by and for supporters of a popular German football club, Borussia Dortmund (a.k.a. BVB). The last commit on the associated project’s source code on GitHub was on August 10, 2017.⁶

Although it has an official Twitter account, @bvbcoin, its last Tweet dates back to 31 August, 2016. The coin’s rating on Cryptopia is a low 1 out of possible 5. This choice highlights the preference of pump-and-dump organizers for coins associated with unserious projects.

During the first 15 minutes of the pump, BVB’s trading volume exploded from virtually zero to 1.41 BTC (illustrated by the tall grey bar towards the right end of the price/volume chart), and the coin price increased from 35 Sat⁷ to its three-fold, 115 Sat (illustrated by the thin grey vertical line inside the tall grey bar).

Price fluctuations: Further dissecting the tick by tick transactions (Figure 4), we note that the first buy order was placed and completed within 1 second after the first coin announcement. With this lightning speed, we conjecture that such an order might have been executed by automation. After a mere 18

⁵<https://bitcointalk.org/index.php?topic=1596932.0>

⁶<https://github.com/bvbcoin/bvbcoin-source>

⁷One Satoshi (Sat) equals 10⁻⁸ Bitcoin (BTC).

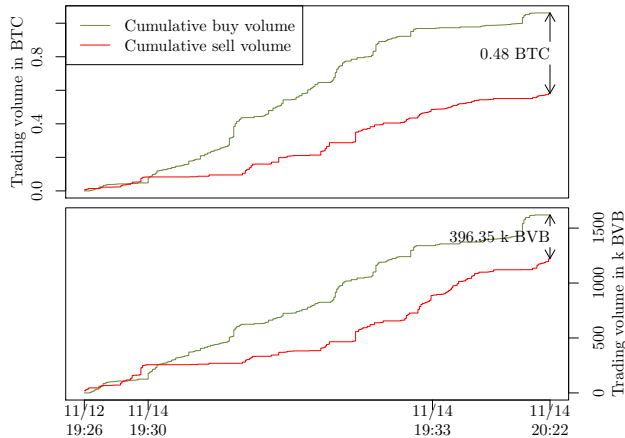


Figure 5: Gap between buy volume and sell volume caused by the BVB pump-and-dump. The figure shows a timeline from 48 hours before up to 1 hour after the pump-and-dump. For the illustration purposes, the timeline is scaled with non-linear transformation to better display the development of volume gaps during the pump-and-dump.

seconds of a manic buying wave, the coin price already skyrocketed to its peak. Note that **Bomba bitcoin “cryptopia”** only announced the coin at the time when the coin price was already at its peak, making it impossible for investors who solely relied on the announcement from the channel to make any money.

Not being able to remain at this high level for more than a few seconds, the coin price began to decrease, with some resistance in between, and then plummeted. Three and half minutes after the start of the pump-and-dump, the coin price had dropped below its open price. Afterwards, transactions only occurred sporadically.

Volume: Figure 5 shows that the pump-and-dump induces fake demand and inflates buy volume. While every pump-and-dump participant would hope for a quick windfall gain during a minute-long pump, the majority would not manage to act fast enough to sell at a high price. Those investors would either end up selling coins at a loss, or, if reluctant to sell low, would hold the virtually worthless coins. This is demonstrated by Figure 5, which shows that the buy volume exceeds the sell volume, whether measured by the target coin BVB or by BTC. The figure also shows small volume movements shortly before the pump-and-dump, also observable in Figure 4(a), which can be indicative of organizers’ pre-purchase conduct. As the BVB blockchain is not being actively maintained and the coin itself is extremely illiquid, any market movement may be deemed unusual.

Figure 5 illustrates that the total buy volume (also including the pre-purchased volume, though negligible) in BTC associated with the pump-and-dump amounts to 1.06 BTC, the sell volume only 0.58 BTC; the total buy volume measured in BVB is 1,619.81 thousand BVB, the sell amount 1,223.36 thousand BVB. This volume discrepancy between the sell and

Exchange	Volume (30d)	No. markets	Launch	Country
Binance	\$21,687,544,416	385	Jul 2017	China
Bittrex	\$1,168,276,090	281	Feb 2014	U.S.A.
Cryptopia	\$107,891,577	852	May 2014	New Zealand
YoBit	\$797,593,680	485	Aug 2014	Russia

Figure 6: Exchanges involved in pump-and-dump schemes, sorted by 30-day volume: No. markets is the number of trading pairs (eg. DASH/BTC, ETC/USDT) in the exchange. Volume and No. markets were extracted from CoinMarketCap on November 5, 2018.

the buy sides indicates a higher trading aggressiveness on the buy side.⁸ This further suggests that many investors may be “stuck” with BVB which they are unwilling to liquidate at the low market price after the pump-and-dump. Those coin holders can only expect to reverse the position in the next pump, which might never come.

Low participation ratio: It is worth noting that the total count of trading transactions associated with this pump-and-dump is merely 322. That number appears very low compared to the 1,376 views of the coin announcement message, let alone the over 10,000 channel members. This indicates that the majority of group members are either observers, who want no skin in the game, or have become aware of the difficulty in securing profit from a pump-and-dump.

4 Analyzing Pump-and-Dump Schemes

In this section we explain how we obtain data from both Telegram and the various exchanges, which allows us to analyze and model pump-and-dump schemes.

4.1 Collecting Pump-and-Dump Events

In this study, we examine routinely organized pump-and-dump events that follow the pattern of “set-up → pre-pump announcement → pump → dump → post-pump review” as described in Section 2. This type of pump-and-dump involves live instructions from organizers (see Figure 1 and Figure 3), so encrypted chat applications such as Telegram and Discord are ideal for broadcasting those events.

We are confident that it suffices to focus solely on pump-and-dump events orchestrated on Telegram as every active pump-and-dump group we found on Discord was also on Telegram.⁹ Telegram is among the primary media for pump-and-dump activities and announcements, and it would be both unreasonable and unlikely for any pump-and-dump organizer

⁸Note that Cryptopia is a peer-to-peer trading platform which lets users trade directly with each other; the exchange takes no risk position and only profits from charging trading fees. Therefore, buying volume implies that the trade is initiated by the buyer, which typically drives the market price up; similarly, sale volume is initiated by the sell side and would drive the price down.

⁹This observation has also been confirmed by the PumpOlymp team, an online information provider specialized in cryptocurrency pump-and-dump.

to restrict the platform to only Discord, since the key to the success of a pump-and-dump is the number of participants.

Telegram channels: Our primary source on pump-and-dump Telegram channels and events is provided by PumpOlymp,¹⁰ a website that hosts a comprehensive directory of hundreds of pump-and-dump channels.

PumpOlymp discovers those channels by searching pump-related keywords — e.g. “pump”, “whales”, “vip” and “coin” — on Telegram aggregators such as <https://tgstat.com/> and <https://telegramcryptogroups.com/>. Another source for new pump-and-dump channels is cross-promotion on the known channels.¹¹ To validate the incoming data from PumpOlymp, we conduct an independent manual search for pump-and-dump channels. We are not able to add new channels to the existing channel list from PumpOlymp, and we are not aware of any other, more comprehensive pump-and-dump channel list. Therefore, we believe the channel list from PumpOlymp is a good starting point.

Next, we use the official Telegram API to retrieve message history from those channels, in total 358, to check their status and activity. Among those channels, 43 have been deleted from the Telegram sever, possibly due to inactivity for an extended period of time. Among the existing ones, over half (168/315) have not been active for a month, possibly because cautious admins delete pump-and-dump messages to eviscerate their traces. This might also imply that the Telegram channels have a “hit-and-run” characteristic. As described in the section above, one learns from participation in pump-and-dump activities that quick bucks are not easy to make. Therefore, curious newcomers might be fooled by pump-and-dump organizers’ advertising and lured into the activity. After losing money a few times, participants may lose faith and interest, and cease partaking. This forms a vicious circle, since with fewer participants, it would be more difficult to pump a coin. Therefore, channel admins might desert their channel when the performance declines, and start new ones to attract the inexperienced.

Pump-and-dump history: Starting June 2018, PumpOlymp has been gleaning pump-and-dump events organized on Telegram. Using their API,¹² we acquire an initial list of historical pump-and-dump activities over the period of June 17, 2018 and February 26, 2019. For each listed pump-and-dump event, the data set contains the pumped coin, the target exchange, the organizing Telegram channel, the coin announcement time, plus the price and volume data on the tick-by-tick level from coin announcement up to 15 minutes afterwards.

We run plausibility checks to validate each record’s qualification as a pump-and-dump. For example, if an alleged pump-and-dump is recorded to have started at a time that is

¹⁰<https://pumpolymp.com>

¹¹This is based on a conversation with a PumpOlymp staff member.

¹²<https://pumpolymp.com:5001/api/allPumps> and <https://pumpolymp.com:5001/api/PumpMarketHistory/raw>, only available for premium users.




Figure 7: Cumulative counts of pumps and pumped coins on four exchanges from June 2018 to February 2019.

far from a full hour (6:00, 7:00, etc.) or a half hour, then we would be suspicious, because an organizer would normally not choose a random time for a pump-and-dump. If there is no significant increase in volume or high price around the pump time, we would also be skeptical. In such a circumstance, we manually check the message history to make a final judgment. In most cases, the message either discusses the potential of a coin or the record is simply a mistake. Note that we exclusively consider message series with count-downs (e.g. “3 hours left”, “5 mins left”) and coin announcement; messages on pump signal detection are eliminated from our sample.

In the end, we trace 429 pump-and-dump coin announcements from June 17, 2018 to February 26, 2019, each of which is characterized by a series of messages similar to those presented in Figure 1. One pump-and-dump can be co-organized by multiple channels; if two coin announcements were broadcast within 3 minutes apart from each other and they target the same coin at the same exchange, then we consider them to be one pump-and-dump event. In total, we collected 412 unique pump-and-dump events.

Excluded data points: All the pumped coins in our sample were paired with BTC. We also observed and manually collected a few ETH-paired pumps, most of which took place in other exchanges.¹³ Inclusion of those cases would require data collection with other methods and resources. Due to their rarity, we do not consider ETH-paired pump-and-dumps in our study.

4.2 Obtaining Coin Data

Apart from consulting the online pump-and-dump information center PumpOlymp, we retrieve additional information on features and price movements of coins from other sources,

¹³For example, PLX on October 10, 2018 in CoinExchange, ETC on April 22, 2018 in Bibox.




Figure 8: Aggregate trading volume of pumped coins before and during a pump.

in order to establish a connection between the information and the pump-and-dump pattern.


Specifically, we use the public API from CryptoCompare¹⁴ for coins’ hourly OHLC (open, high, low, close) and volume data on 189 exchanges, including Binance, Bittrex, Cryptopia and Yobit. The API provides live data, which means users are able to obtain price information up to the time point of data retrieval. While historical minute-level data are also available on CryptoCompare, they are restricted to a 7-day time window and thus not utilized.

In the conventional stock market, pump-and-dump operators favor microcap stocks due to high manipulability of their price [3]; we expect to observe a similar phenomenon in the crypto-market. To collect coins’ market cap data, we use the public API from CoinMarketCap. Because we are interested in coins’ “true” market cap that is uninfluenced by any maneuver, we purposefully chose to retrieve the data at 08:42 GMT, November 5. We believe the market cap data retrieved are not contaminated by Telegram organized pump-and-dumps, since they typically start on the hour or the half hour and last only a few minutes.


In addition to market trading data, we also retrieve coins’ non-financial features. Specifically, we use exchanges’ public API¹⁵ to collect information on coins’ listing status, algorithm, and total supply. We also collect coins’ launch dates using CryptoCompare’s public API. For information that is not contained in the API but viewable online (such as coins’ rating data on Cryptocurrency), we use either page source scraping or screen scraping, depending on the design of the desired webpage. All our data on coin features are from publicly accessible sources.

¹⁴<https://min-api.cryptocompare.com/>

¹⁵<https://api.binance.com/api/v1/ticker/allPrices> for Binance, <https://bittrex.com/api/v1.1/public/getcurrencies> for Bittrex, <https://www.cryptopia.co.nz/api/GetCurrencies> for Cryptopia, and <https://yobit.net/api/3/info> for Yobit.



(a) Pump and dump activities from June 2018 to February 2019



(b) Enlarged section of the highlighted area in (a) that shows one of the most recent pump-and-dump

Figure 9: Pump and dump timeline. A green bar represents price increase through pump, calculated as $\frac{\text{high price} - \text{open price}}{\text{open price}}$; a red bar represents price drop after pump, calculated as $\frac{\text{close price} - \text{high price}}{\text{close price}}$. All prices are denominated in BTC, and from a 3-hour window around pump activities. Visit <http://rpubs.com/xujiahuyz/pd> for the full, interactive chart.

4.3 Role of Exchanges

Pump-and-dump schemes take place within the walled gardens of crypto-exchanges. Binance, Bittrex, Cryptopia, and Yobit are among the most popular exchanges used by pumpers (see Figure 6). While those exchanges differ vastly in terms of their volume, markets, and user base, each of them has its own appeal to pumpers. Large exchanges such as Binance and Bittrex have a large user base, and abnormal price hype caused by pump activities can quickly attract a large number of other users to the exchange. Smaller exchanges such as Cryptopia and Yobit tend to host esoteric coins with low liquidity, whose price can be more easily manipulated compared to mainstream coins such as Ether (ETH) or Litecoin (LTC).

In general, larger exchanges are more reliable than smaller ones. While both Binance and Cryptopia were hacked recently,¹⁶ the former managed to remain operative, while the

¹⁶<https://www.bloomberg.com/news/articles/2019-05-08/crypto-exchange-giant-binance-reports-a-hack-of-7-000-bitcoin>

Exchange	Number of PD's	Admins' profit (BTC), aggregated	Admins' return, aggregated
Binance	51	148.97	15%
Bittrex	15	0.92	7%
Cryptopia	180	44.09	57%
Yobit	102	5.54	52%
Total	348	199.52	18%

Table 1: Number of pump-and-dumps (348) considered in this analysis deviates from the total number of pump-and-dumps (412) due to lack of price data for some events.

latter halted trading and fell into liquidation.

Activity distribution by exchange: Among the 412 pump-and-dump activities, 68 (17%) took place in Binance, 21 (5%) in Bittrex, 211 (51%) in Cryptopia and 112 (27%) in Yobit. In aggregate, 35% (146/412) of the time, the selected coin had previously been pumped in the same exchange (see Figure 7).

Figure 8 compares the aggregate three-hour trading volume in BTC of pumped coins before and during a pump-and-dump, and the artificial trading volume generated by those pump-and-dump activities is astonishing: 8,793 BTC (93% from Binance), roughly equivalent to 50 million USD,¹⁷ of trading volume during the pump hours, 9 times as much as the pre-pump volume (943 BTC), and that only over a period of eight months.

Figure 9 illustrates the occurrence and the effectiveness of individual pump-and-dump activities. In terms of frequency, Bittrex is most rarely chosen; Binance started to gain traction only since September, but still witnesses far less pump-and-dump occurrence than Yobit and Cryptopia. Turning to Yobit with Cryptopia, we find that the two exchanges have complemented each other: when Yobit was inactive (most notably October 2018 to January 2019), Cryptopia experienced more traffic; when Cryptopia went silent (since the hack in mid-January 2019), Yobit regained popularity. In terms of percentage of coin price increase, pumps in both Yobit and Cryptopia appear to be more powerful than those in Bittrex and Binance. What goes hand-in-hand with price surge is price dip: coin prices also drop more dramatically during the dump in Yobit and Cryptopia compared to their peer exchanges.

Profit for admins: Even with tick-by-tick data for each pumped coin during their respective pump-and-dump period, due to lack of trader ID we cannot precisely match individuals' buy and sell transactions. Therefore, to estimate profit for admins, we need to make a few assumptions:

1. Admins purchase coins and enter sell orders only prior to the pump.

and https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12231209.

¹⁷This is calculated based on the unit BTC price of 5,715 USD, which is the mean of the high price of 8,250 USD and the low price 3,180 USD during the data period.

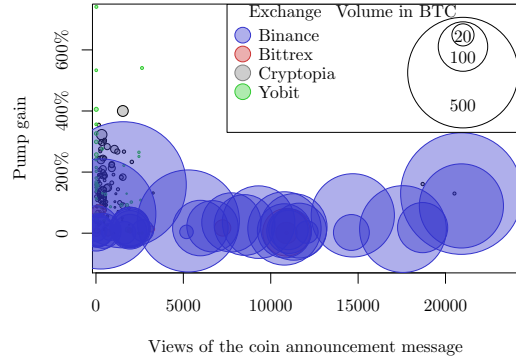


Figure 10: Views of coin announcement message versus coin price increase during the pump. The figure illustrates the relationships between coin price increase through pump, views of coin announcement message, pump volume, and pump exchange.

2. Admins purchase coins at the price immediately before the pump begins.
3. During the pump period — before the price reaches the peak, investors lift the admin's offers and push the price higher; during the dump period — when the price drops, investors transact with each other.

With those assumptions, we arrive at the estimation as presented in Table 1. We estimate that admins made a net profit of 199.52 BTC, equivalent to 1.1 million USD, through 348 pump and dump events during our sample period. The estimated return of insiders averages 18%, which aligns perfectly with Li et al. [23].

So, what is the investors' payout? Some investors win; others lose. Since trading is a zero-sum game, the aggregate investor loss would be on the equivalent scale as the aggregate admin win.

Coin announcement views: While investigating the degree of exposure in coin announcement messages distributed by Telegram channels, we find a negative correlation (-0.162) between number of views of coin announcement and pump gain, which is rather counter-intuitive, because one would think that more views would indicate more participation, which would result in higher pump gain. Two extreme examples: the coin announcement of the pump on MST had 325 views and the pump gain was 12.6%; another coin announcement of the pump on PARTY had only 18 views, and the pump gain was a whopping 533.3%.

This finding suggests that the number of views cannot accurately proxy number of participants, possibly because: (1) only a fraction of message viewers would actually participate in a pump-and-dump; (2) if a user reads the message history after the pump, his/her view would still be counted; (3) if a user re-views a message 24 hours after his/her first view, the user's view would be counted twice;¹⁸ (4) some participants

¹⁸<https://stackoverflow.com/questions/42585314/telegram-channels-post-view-count>

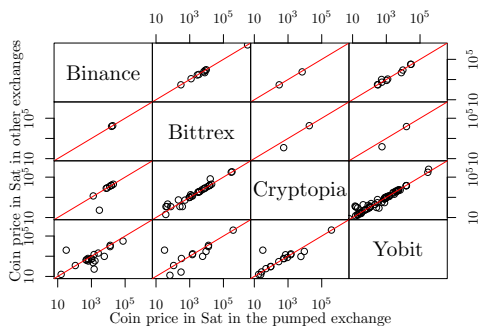


Figure 11: Arbitrage opportunities: coin price (highest during the pump hour) in pumped exchange versus price in other exchanges

might have retrieved messages via bots, which would not be counted in number of views.¹⁹

Price increase: We further notice that although pump-and-dumps in Binance generate more trading volume during the pump hour (Figure 8),²⁰ thanks to its large user base, coin price increase through pumps is generally at a much smaller scale than that in Cryptopia and Yobit (Figure 9 and Figure 10). This is possibly caused by high bid and sell walls on the order book that are typical for large crypto exchanges like Binance, which prevent the price from fluctuating significantly even at coordinated pump-and-dump events.

Arbitrage: Pump-and-dump activities not only engender abnormal returns within the pumped exchange, but also arbitrage opportunities across different exchanges. Figure 11 shows the presence of a price discrepancy of the same coin during the pump hour across different exchanges. Interestingly, coin price can sometimes be higher in exchanges other than the pumped one. It is also worth noting that most coins pumped in Cryptopia are also listed in Yobit but not in Bittrex or Binance, and vice versa. This is because the former two have more conservative coin listing strategies, which results in a different, more mainstream portfolio of listed coins compared to the latter two. While there may be trading strategies resulting from these arbitrage opportunities, they are outside the scope of this work.

4.4 Capturing Features

Market cap: Figure 12 presents the market cap distribution of coins pumped in different exchanges. Pumped coins' market cap ranges from 1 BTC (Royal Kingdom Coin (RKC), pumped in Cryptopia) to 27,600 BTC (TrueUSD (TUSD), pumped in Yobit). Half of those coins have a market cap below 100 BTC, most of which were pumped in Cryptopia.

¹⁹<https://stackoverflow.com/questions/49704911/is-it-possible-for-a-telegram-bot-increase-post-view-count>

²⁰A pump hour refers to the clock hour during which a pump occurs.




Figure 12: Distribution of coin market caps. Market cap information was extracted from CoinMarketCap on November 5, 2018.

Pump-and-dump organizers' preference for small-cap coins resembles equity market manipulators' taste for microcap stocks [3, 24], and can be explained by the empirical finding of Hamrick et al. [18] and Li et al. [23]: the smaller the market cap of the pumped coin, the more successful the pump would be.

Price movement: Figure 13 depicts time series of hourly log returns of pumped coins between 48 hours before and 3 hours after a pump. We detect anomalous return signals before pump-and-dump admins' announcement of the pumped coin. The signals appear most jammed one hour prior to the pump, and less so before that. This is to a certain degree in accord with Kamps et al. [20] who find that a shorter, 12-hour rolling estimation window is more suitable for anomaly detection in the crypto-market than a longer, 24-hour one.

The return signal before the pump is the strongest with Cryptopia, where in numerous pumps, coin prices were elevated to such an extent that the hourly return before the pump even exceeds the hourly return during the pump. This can be explained by the assumption that pump organizers might utilize their insider information to purchase the to-be-pumped coin before the coin announcement, causing the coin price elevation and usual return volatility before the pump. The analysis above provides grounds for predicting the pumped coin before coin announcement using coin features and market movement.

5 Predicting Pump-and-Dump Target Coins

5.1 Feature Selection

Based on the preliminary analysis in the last section, we believe pump-and-dump organizers have specific criteria for coin selection and they generally purchase the to-be-pumped coin before naming it to the investors. Thus, it should be possible to use coin features and market movements prior to a

Feature	Description	Notation
Market cap	Market cap information extracted from CoinMarketCap at 08:42 GMT, November 5, 2018 when no pump-and-dump activity in Telegram channels was observed *	<i>caps</i>
Returns before pump	x -hour log return of the coin within the time window from $x + 1$ hours to 1 hour before the pump	$return[x]h^\dagger$
Volumes in coin before pump	Total amount of the coin traded within the time window from $x + 1$ hours to 1 hour before the pump	$volumefrom[x]h^\dagger$
Volumes in BTC before pump	Total trading volume of the coin measured in BTC within the time window from $x + 1$ hours to 1 hour before the pump	$volumeto[x]h^\dagger$
Return volatilities before pump	Volatility in the hourly log return of the coin within the time window from $y + 1$ hours to 1 hour before the pump	$returnvola[y]h^\ddagger$
Volume volatilities in coin before pump	The volatility in the hourly trading volume in coin within the time window from $y + 1$ hours to 1 hour before the pump	$volumefromvola[y]h^\ddagger$
Volume volatilities in BTC before pump	The volatility in the hourly trading volume in BTC within the time window from $y + 1$ hours to 1 hour before the pump	$volumetovola[y]h^\ddagger$
Last price before pump	Open price of the coin one hour before the coin announcement	<i>last price</i>
Time since existence	The time difference between the time when the first block of the is mined and the pump time	<i>age</i>
Pumped times before	Number of times the coin been pumped in Cryptopia before	<i>pumpedtimes</i>
Coin rating	Coin rating displayed on Cryptopia, 0 being the worst, 5 being the best. The rating considers the following criteria wallet on {Windows, Linux, Mac, mobile, web, paper}, premine ratio, website and block explorer	<i>rating</i>
Withdrawal fee	Amount of coin deducted when withdrawing the coin from Cryptopia	<i>WithdrawFee</i>
Minimum withdrawal	Minimum amount of coin that can be withdrawn from Cryptopia	<i>MinWithdraw</i>
Maximum withdrawal	Daily limit on the amount of coin that can be withdrawn from Cryptopia	<i>MaxWithdraw</i>
Minimum base trade	Minimum base trade size of the coin	<i>MinBaseTrade</i>

Table 2: Features included in the prediction model. *The feature is designed to represent a coin’s market cap in a normal setting, i.e. absent market manipulation. While it might be useful to also collect coins’ historical market cap before each pump-and-dump, we have not found a public source that provides this type of data. $^\dagger x \in \{1, 3, 12, 24, 36, 48, 60, 72\}$. $^\ddagger y \in \{3, 12, 24, 36, 48, 60, 72\}$.

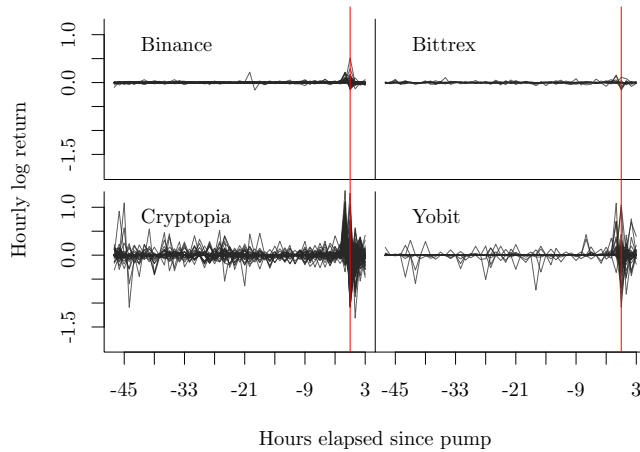


Figure 13: Time series of coin returns before and after pump. In each subplot, the hourly log return of each pumped coin before and shortly after the pump is superimposed. The vertical red line represents the pump hour during which the coin was announced.

coin announcement to predict which coin might be pumped.

In the following exercise, we focus on predicting coins pumped in one specific exchange for the ease of data harmo-

nization. We choose Cryptopia due to sufficient data collected for modelling. Although the exchange ceased to operate on May 15, 2019, our exercise demonstrates a proof of concept for strategic crypto-trading that can be adapted for any exchange.

For each coin before a pump event, we predict whether it will be pumped (TRUE) or not (FALSE). The formula for the prediction model is:

$$Pumped = M(feature_1, feature_2, \dots)$$

where the dependent variable *Pumped* is a binary variable that equals 1 (TRUE) when the coin is selected for the pump, and 0 (FALSE) otherwise. Table 2 lists the features considered in the prediction model.

Previous analyses indicate unusual market movements prior to the pump-and-dump might signal organizers’ pre-pump behavior, which could consequently give away the coin selection information. Therefore, we place great emphasis on features associated with market movements, such as price, returns and volatilities covering various lengths of time. Those features, 46 in total, account for 85% of all the features considered.

5.2 Model Application

Sample specification: We consider all the coins listed on Cryptopia at each pump-and-dump event. On average, we have 296 coin candidates at each pump, out of which one is the actual pumped coin. The number of coins considered varies for each event due to constant listing/delisting activities on the part of exchanges. The full sample contains 53,208 pump-coin observations, among which 180 are pumped cases,²¹ accounting for 0.3% of the entire sample population. Apparently, the sample is heavily skewed towards the unpumped class and needs to be handled with care at modelling.

For robustness tests, we split the whole sample into three chronologically consecutive datasets: training sample, validation sample and and test sample:

Pumped?	Training	Validation	Test	Total	
TRUE	60	60	60	180	(0.3%)
FALSE	17,078	17,995	18,135	53,028	(99.7%)
Total	17,138	18,055	18,195	53,208	(100.0%)

The training sample covers the period of June 19, 2018 to September 5, 2018 and consists of 17,138 data points (32.2% of full sample); the validation sample covers September 5, 2018 to October 29, 2018 and consists of 18,055 data points (33.9% of full sample); the test sample covers October 29, 2018 to January 11, 2019 and consists of 18,195 data points (34.2% of full sample).

Model selection: We test both classification and logit regression models for the prediction exercise. Specifically, for the classification model, we choose random forest (RF) with stratified sampling; for the logit regression model, we apply generalized linear model (GLM). Both RF and GLM are widely adopted in machine learning and each has its own quirks.

RF is advantageous in handling large quantities of variables and overcoming overfitting issues. In addition, RF is resilient to correlations, interactions or non-linearity of the features, and one can be agnostic about the features. On the flip side, RF relies upon a voting mechanism based on a large number of bootstrapped decision trees, which can be time-consuming, and thus challenging to execute. In addition, RF provides information on feature importance, which is less intuitive to interpret than coefficients in GLM.

GLM is a highly interpretable model [28] that can uncover the correlation between features and the dependent variable. It is also highly efficient in terms of processing time, which is a prominent advantage when coping with large datasets. However, the model is prone to overfitting when fed with too many features, which potentially results in poor out-of-sample performance.

²¹Due to missing data on several delisted coins, this number deviates from the total number of 211 pump events in Cryptopia, as presented in Figure 7.

Hyperparameter specification: Due to the heavily imbalanced nature of our sample, we stratify the dataset when using RF [9], such that the model always includes TRUE cases when bootstrapping the sample to build a decision tree. Specifically, we try the following three RF variations:

Model	Sample size per tree		Total	Number of trees
	TRUE	FALSE		
RF1	60	20,000	20,060	5,000
RF2	60	5,000	5,060	10,000
RF3	60	1,000	1,060	20,000

We fix the number TRUES at 60 for each RF variation, so that the model may use the majority of TRUES to learn their pattern when building each tree. Model RF1 stays loyal to our sample’s original TRUE/FALSE ratio, with 0.3% of TRUES contained in each tree-sample. RF2 and RF3 raise the TRUE/FALSE ratio to 1.2% and 6%, respectively. Note that while the sample size per tree decreases from RF1 to RF2 to RF3, we are mindful to increase the number of trees accordingly to ensure that whichever model we use, every input case is predicted a sufficient number of times. We use the R package `randomForest` to model our data with RF1, RF2 and RF3.

With conventional binomial GLM, problems can arise not only when the dependent variable has a skewed distribution, but also when features are skewed. With heavy-tailed coin price distribution and market cap distribution, conventional binomial GLM can be insufficient to handle our sample. Therefore, we apply LASSO (least absolute shrinkage and selection operator) regularization to the GLM models. After preliminary testing, we choose to focus on three representative LASSO-GLM models with various shrinkage parameter values (λ):

Model	Shrinkage parameter (λ)
GLM1	10^{-8}
GLM2	10^{-3}
GLM3	5×10^{-3}

Higher values of λ causes elimination of more variables. We use the R package `glmnet` to model our data with GLM1, GLM2, and GLM3.

Variable assessment: By applying the specified models on the training sample, we are able to assess the features’ relevance to coin prediction. Figure 17 presents features’ importance based on mean decrease in Gini coefficient with RF models. We find that:

- Coin market cap *caps* and last hour return before the pump *return1h* appear to be the two most important features in predicting pumped coin using RF models.
- Features describing market movements shortly before the pump, e.g. *return1h*, *volumeto1h* and *volumefrom1h*, appear to be more important than features describing longer-term movements.




Figure 14: Model performance on the training sample measured by Precision, F_1 (left) and ROC AUC (right) at different threshold levels.

- Among all the features related to market movements, return features are generally more important than volume or volatility features.
- Exchange-specific features including *MinBaseTrade*, *MinWithdraw*, *MaxWithdraw*, and *WithdrawFee* are least important.

Figure 18 presents the estimated coefficients of variables with GLM models, from which we obtain several findings in line with what is indicated by RF models above. Specifically, we notice that:

- When only one variable is included, *return1h* appears to have the highest explanatory power on coins' pump likelihood;
- The positive coefficients of return features imply that the higher the return a coin shows before the pump, the more likely the coin is to be pumped;
- The positive coefficient of *pumpedtimes* implies that pumped coins are more likely to get pumped again.

The variable assessment performed by RF and GLM is coherent in that both find features representing market movement shortly before the pump to be more important than longer-term features. This echoes our exploratory analysis illustrated in Figure 13 and aligns with Kamps et al. [20]. The finding suggests the spontaneity of admins' coin selection, and the importance for strategic traders to obtain real-time market data.

5.3 Assessing Prediction Accuracy

Both the random forest model and GML predict whether a given coin will be pumped as a likelihood ranging between 0

and 1. We apply thresholding to get a binary TRUE/FALSE answer.

Figure 14 depicts the in-sample fitting of model candidates with the training sample as the threshold value changes. The fitting measurements include precision, the F_1 measure and area under ROC (Receiver operating characteristic) curve. Figure 14(a) describes the performance of RF models and Figure 14(b) GLM models.

Precision represents the number of true positive divided by number of predicted positive, and the precision line ends when the denominator equals zero, i.e. when no TRUE prediction is produced. Figure 14 shows that, among the three RF models, the threshold value at which the line ends is the lowest with RF1, and highest with RF3. This indicates that absent balanced bootstrapping, an RF model tends to systematically underestimate pump likelihood, leading to zero predicted TRUE cases even when the threshold value is small.

Compared to RF models, none of the GLM models is able to produce high precision.

In terms of F_1 measure, RF models again appear superior to GLM models. Among the three RF models, the RF1 performs best at a low threshold range (< 0.2), while RF3 performs best at a high threshold range (> 0.4). RF2 resides in between.

The RF models' superiority to GLM models is further demonstrated by the ROC (Receiver operating characteristic) curve in Figure 14. Among the three RF models, no discernible difference can be found in terms of ROC AUC: all exhibit high performance with $AUC > 0.94$. The GLM models, in contrast, render an AUC between 0.63 and 0.88.

Due to their obvious inferiority, we eliminate GLM models from further analysis. Figure 15 illustrates the out-of-sample performance of RF models. The model performance with the validation sample resembles that of the training sample, remaining strong with regard to all three indicators (precision, F_1 and AUC). This suggests that the classification model trained and calibrated on one period of data can accurately predict a later period.

Both Figure 14(a) and Figure 15 suggest that balancing the sample with various TRUE/FALSE ratios only changes the absolute value of the pump likelihood output, but not the relative one. This means the three RF models can perform similarly in terms of Precision and F_1 measure, when the appropriate threshold value is chosen in correspondence with the model (specifically, $Threshold_{RF1} < Threshold_{RF2} < Threshold_{RF3}$).

5.4 Testing an Investment Strategy

To explore the model's practical utility, we devise a simple investment strategy. At each pump, we check which coin's predicted pump likelihood surpasses a predetermined threshold, and we purchase all those coins before the actual coin announcement (if no coin's vote exceeds the threshold, we will not pre-purchase any coin). Note that if we had the ability to short or use margin trading on the exchanges we use,




Figure 15: Performance of RF models on the validation sample measured by Precision, F_1 (left) and ROC AUC (right) at different threshold levels.




Figure 16: Investment return using different models at different threshold levels.

potentially more options would open up for us.

Strategy: Specifically, for each coin that we pre-purchase, we buy the coin at the open price one hour before the coin announcement with the amount of BTC equivalent to k times the vote where k is a constant. That is to say, with all the coins we purchase, the investment, measured in BTC, on each coin is proportionate to its vote supplied by the random forest model. This is logical because a higher vote implies a higher likelihood of being pumped, and thus worth a higher investment.

We further assume that among all the coins we purchased, those coins that do not get pumped (false positive, “false alarms”) will generate a return of zero, i.e. their price will remain at the same level as the purchase price; those coins that get pumped (true positive, “hits”) will be sold at an elevated price during the pump. To be conservative, we assume that with each purchased coin that gets pumped we only obtain half of the pump gain, expressed as:

$$\text{pump gain} = \frac{\text{high price} - \text{open price}}{\text{open price}}.$$

Returns: Figure 16 presents the relationship between the aggregate return and the threshold choice.

Figure 16(a) illustrates the performance of the trading strategy with the training sample. The figure shows that, in general, the *higher* the threshold, which means we buy coins with higher pump likelihoods and disregard others, the *higher* the return.

Figure 16(b) illustrates the performance of the trading strategy with the validation sample. As the threshold increases, the return first increases and then decreases. This is because the coins with the highest predicted pump likelihood in the validation sample happen to have very low pump gain. When the threshold is high, only those coins with high likelihood but

		Predicted		Total
		TRUE	FALSE	
Actual	TRUE	9	51	60
	FALSE	0	18,135	18,135
Total		9	18,186	18,195

Table 3: Confusion matrix of RF1 with threshold value 0.3 applied to test sample.

low gain are included in the investment portfolio, resulting in a low overall return.

As already mentioned at the end of Section 5.3, every model has its own optimal threshold value. In terms of the magnitude of the profit, with the right combination of threshold and model, investors would theoretically enjoy a return of 140% with the training sample cases (RF1 with threshold of 0.7), and a return of 80% with the validation sample cases (RF1 with threshold of 0.3).

One should be mindful that if the threshold is set *too high* (e.g., greater than 0.8), then the investor might end up not buying any coins, and consequently gaining no profit. In addition, although high threshold comes with high precision, it also leads to a low number of coins being purchased, increasing the risk associated with an undiversified investment portfolio, as demonstrated in Figure 16(b).

5.5 Final Test

Based on the training and validation results of specified models, we need to select one model and an accompanying threshold value to apply to the test sample. Our ultimate goal to maximize the trading profit using the selected model in combination with the proposed trading strategy on a set of out-of-sample data. Therefore, we base our decision primarily on Figure 16(b). We apply RF1 and a threshold of 0.3 — the combination that delivers the highest return in Figure 16(b) — on our test sample.

To determine the investment amount in BTC for our trading strategy, we need to examine the market depth. This is particularly important for exchanges with low trading volume such as Cryptopia and Yobit. When trading in those exchanges, it has to be ensured that during the pump-and-dump, the market would provide sufficient depth for us to liquidate the coins purchased prior to the pump. For example, if the total trading volume in one event is 0.4 BTC, it would make no sense to spend 0.8 BTC on the coin.

To this end, we calculate the average trading volume per pump-and-dump at Cryptopia. We only consider “uptick” transactions, i.e. where the buyer is the aggressor. This yields a ballpark estimation of the market depth on the buy side. We use this number, 0.37 BTC, as the baseline investment quantity. This baseline amount, discounted by the predicted pump likelihood, would be the investment value in BTC.

Coin	Date	Pumped?	weight <i>wt</i>	BTC	Pump	Assumed	BTC
				invested $q = \bar{Q} \times wt$	gain <i>pg</i>	gain $ag = pg/2$	gained $q \times ag$
BVB	Nov 14	TRUE	0.30	0.11	283%	142%	0.16
CON	Nov 16	TRUE	0.44	0.16	33%	17%	0.03
FLAX	Nov 10	TRUE	0.58	0.21	135%	67%	0.14
MAGN	Nov 13	TRUE	0.37	0.14	70%	35%	0.05
MAGN	Dec 16	TRUE	0.39	0.14	85%	43%	0.06
OSC	Nov 13	TRUE	0.65	0.24	297%	148%	0.36
OSC	Nov 25	TRUE	0.52	0.19	100%	50%	0.10
SOON	Nov 01	TRUE	0.58	0.21	10%	5%	0.01
UMO	Nov 15	TRUE	0.55	0.20	60%	30%	0.06
				1.61			0.96

Table 4: Purchased coins based on pump likelihood predicted by RF1. Only coins with predicted pump likelihood of greater than 0.3 are purchased. Investment weight equals pump likelihood. $\bar{Q} = 0.37$, the average of total transaction volume in a pump-and-dump event in Cryptopia. Only transaction volume where the buyer is the aggressor is considered.

Table 3 displays the confusion matrix of the model prediction with the test sample. The model suggests us to purchase 9 coins in total, all of which are ultimately pumped. Table 4 lists those 9 coins, their respective investment weight and assumed profit. The return on the investment amounts to 60% (2.61/4.38) over the test sample period of two and a half months. Note that the effect of transaction fees (0.2% on Cryptopia) on the investment profitability is negligible. The result of the final test is very similar to that with both the training sample and the validation sample when the same combination of model (RF1) and threshold (0.3) is applied (Figure 16), confirming the model’s robustness.

5.6 Caveats and Improvement Potential

Data: Upon availability, order book data, tick-by-tick data before a pump and traders’ account information can also be included as features.

Modelling method: Random forest with unsupervised anomaly detection has the potential to improve the model performance. In addition, other classification (e.g. k-NN) and regression (e.g. ridge) models are worth considering.

Additional considerations: Regarding investment weights, one may consider coin price increase potential (based on e.g. historical returns) in combination with coin pump likelihood. One must beware that in liquid exchanges, the trading strategy only applies to *small* retail investment, since big purchase orders prior to a pump can move the market, such that pump organizers may cancel the pump or switch the coin last-minute. Also worth factoring in is the market risk (e.g. security risk, legal risk) associated with the nascent crypto-market.

6 Related Work

Over the past year, a handful of studies researching cryptocurrency pump-and-dump activities have been conducted, notably Kamps et al. [20], Li et al. [23] and Hamrick et al. [18].

Our work differs from the aforementioned studies in terms of motivation, methodology, data, and contribution. We aim for *prospective* prediction as opposed to *retrospective* investigation of pump-and-dump activities. We use a homogeneous set of data that only includes clearly announced pump-and-dump events on Telegram.²² Regarding the sample period, our data cover a recent time span of June 17, 2018 to February 26, 2019. (Table 5).

Our paper is also closely linked to literature on market manipulation in non-cryptocurrency contexts. Lin [24] explains potential damage of various manipulation methods including pump-and-dump, front running, cornering and mass misinformation, and argues for swift regulatory action against those threats. Austin [3] calls for authorities’ demonstration of their ability to effectively deter market manipulation such as pump-and-dump in exchanges for small-capped companies, in order to recover investors’ confidence in trading in those markets, which would consequently foster economic growth.

Our paper is further related to research on crypto trading. Gandal et al. [17] demonstrate that the unprecedented spike in the USD-BTC exchange rate in late 2013 was possibly caused by price manipulation. Makarov et al. [25] probe arbitrage opportunities in crypto markets. Aune et al. [2] highlight potential manipulation in the blockchain market resulting from the exposure of the footprint of a transaction after its broadcast and before its validation in a blockchain, and proposes a cryptographic approach for solving the information leakage problems in distributed ledgers.

Our paper is also akin to existing literature on cryptocurrencies’ market movements. The majority of related literature still orients its focus on Bitcoin. Many scholars use GARCH models to fit the time series of Bitcoin price. Among them, Dyhrberg et al. [13] explore the financial asset capabilities of Bitcoin and suggests categorizing Bitcoin as something between gold and US Dollar on a spectrum from pure medium of exchange to pure store of value; Bououiour et al. [7] argue that Bitcoin is still immature and remains reactive to negative rather than positive news at the time of their writing; 2 years later, Conrad et al. [10] present the opposite finding that negative press does not explain the volatility of Bitcoin; Dyhrberg [14] demonstrates that bitcoin can be used to hedge against stocks; Katsiampa [21] emphasizes modelling accuracy and recommends the AR-CGARCH model for price retro-fitting. Bariviera et al. [4] compute the Hurst exponent by means of the Detrended Fluctuation Analysis method and conclude that the market liquidity does not affect the level of long-range dependence. Corbet et al. [11] demonstrate that Bitcoin shows characteristics of a speculative asset rather than a currency also with the presence of futures trading in Bitcoin.

Among the few research studies that also look into the financial characteristics of other cryptocurrencies, Fry et al. [16]

²²As suggested earlier, all the coin announcements we found on Discord overlap with our Telegram data

	Kamps et al. [20]	Hamrick et al. [18]	Li et al. [23]	This paper
Motivation	Locating suspicious transactions patterns through automated anomaly detection	Identifying success factors for historical pumps	Examining how pump-and-dumps are correlated with cryptocurrency price	Predicting the coin to be pumped with input of Telegram signals
Methodology	Breakout indicators & reinforcers	Ordinary least squares (OLS)	OLS, difference in difference	RF, GLM
Data	Market data of cryptocurrencies from April 2018 to May 2018 on Binance, Bittrex, Kraken, Kucoin and Lbank	Explicit (with coin announcement) and suspected (no coin announcement) pump-and-dumps from January 2018 to July 2018	Pump-and-dump events from May 2017 to August 2018 on Binance, Bittrex, and Yobit, with a focus on Bittrex	Pump-and-dump events from June 2018 to February 2019 on Binance, Bittrex, Cryptopia and Yobit, with a focus on Cryptopia
Main finding / contribution	The authors develop a defining criteria set for detecting suspicious activity like pump-and-dumps.	Pumping obscure, small-market-cap coins is more likely to be successful.	Pump-and-dumps are detrimental to the liquidity and price of cryptocurrencies.	Pump-and-dumps schemes can be found and foiled by machine learning.

Table 5: Comparison of studies on cryptocurrency pump-and-dump.

examine bubbles in the Ripple and Bitcoin markets; Baur et al. [6] investigate asymmetric volatility effects of large cryptocurrencies and discover that in the crypto market positive shocks increase the volatility more than negative ones. Jahani et al. [19] assess whether and when the discussions of cryptocurrencies are truth-seeking or hype-based, and discover a negative correlation between the quality of discussion and price volatility of the coin.

7 Conclusions

This paper presents a detailed study of pump-and-dump schemes in the cryptocurrency space. We start by presenting the anatomy of a typical attack and then investigate a variety of aspects of real attacks on crypto-coins over the last eight months on four crypto-exchanges. The study demonstrates the persisting nature of pump-and-dump activities in the crypto-market that are the driving force behind tens of millions of dollars of phony trading volumes each month. The study reveals that pump-and-dump organizers can easily use their insider information to profit from a pump-and-dump event at the sacrifice of fellow pumpers.

Through market investigation, we further discover that market movements prior to a pump-and-dump event frequently contain information on which coin will be pumped. Using LASSO regularized GML and balanced random forests, we build various models that are predicated on the time and venue (exchange) of a pump-and-dump broadcast in a Telegram group. Multiple models display high performance across all subsamples, implying that pumped coins can be predicted based on market information. We further propose a simple but effective trading strategy that can be used in combination with the prediction models. Out-of-sample tests show that a return of as high as 60% over two and half months can be consistently exploited even under conservative assumptions.

In sum, we wish to raise the awareness of pump-and-dump schemes permeating the crypto-market through our study. We

show that with fairly rudimentary machine learning models, one can accurately predict pump-and-dump target coins in the crypto-market. As such, we hope our research could, on one hand, lead to fewer people falling victim to market manipulation and more people trading strategically, and on the other hand, urge the adoption of new technology for regulators to detect market abuse and criminal behavior. If such advice would be heeded, admins' schemes would crumble, which would in turn lead to a healthier trading environment, accelerating the market towards a more fair and efficient equilibrium.

References

- [1] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. A Survey of Attacks on Ethereum Smart Contracts. *Proceedings of the 6th International Conference on Principles of Security and Trust*, 10204:164–186, 2017.
- [2] Rune Tevasvold Aune, Adam Krellenstein, Maureen O’hara, and Ouziel Slama. Leakage in Distributed Ledgers Trading and Information Footprints on a Blockchain. *Journal of Trading*, 12(3):5–13, 2017.
- [3] Janet Austin. How Do I Sell My Crowdfunded Shares? Developing Exchanges and Markets to Trade Securities Issued by Start-Ups and Small Companies. *Harvard Business Law Review*, 8:21–35, 2018.
- [4] Aurelio F. Bariviera, María José Basgall, Waldo Hasperué, and Marcelo Naiouf. Some stylized facts of the Bitcoin market. *Physica A: Statistical Mechanics and its Applications*, 484:82–90, 2017.
- [5] Massimo Bartoletti, Salvatore Carta, Tiziana Cimoli, and Roberto Saia. Dissecting Ponzi schemes on Ethereum: identification, analysis, and impact. 2017.
- [6] Dirk G. Baur and Thomas Dimpfl. Asymmetric volatility in cryptocurrencies. *Economics Letters*, 173:148–151, 2018.
- [7] Jamal Bouoiyour and Refk Selmi. Bitcoin: A beginning of a new phase? *Economics Bulletin*, 36(3):1430–1440, 2016.
- [8] CFTC. Customer Advisory: Beware Virtual Currency Pump-and-Dump Schemes. 2018.
- [9] Chao Chen and Andy Liaw. Using Random Forest to Learn Imbalanced Data. Technical report, 2004.
- [10] Christian Conrad, Anessa Custovic, Eric Ghysels, Christian Conrad, Anessa Custovic, and Eric Ghysels. Long- and Short-Term Cryptocurrency Volatility Components: A GARCH-MIDAS Analysis. *Journal of Risk and Financial Management*, 11(2):23, 2018.
- [11] Shaen Corbet, Brian Lucey, Maurice Peat, and Samuel Vigne. Bitcoin Futures – What use are they? *Economics Letters*, 172:23–27, 2018.
- [12] Crypto Insider. CFTC offers \$100,000+ bounty for crypto pump and dump whistleblowers.
- [13] Anne Haubo Dyrberg. Bitcoin, gold and the dollar - A GARCH volatility analysis. *Finance Research Letters*, 16:85–92, 2015.
- [14] Anne Haubo Dyrberg. Hedging capabilities of bitcoin. Is it the virtual gold? *Finance Research Letters*, 16:139–144, 2016.
- [15] Federal Reserve. Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices, 2016.
- [16] John Fry and Eng-Tuck Cheah. Negative bubbles and shocks in cryptocurrency markets. *International Review of Financial Analysis*, 47:343–352, 2016.
- [17] Neil Gandal, JT Hamrick, Tyler Moore, and Tali Oberman. Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics*, 95(4):86–96, 2018.
- [18] JT Hamrick, Farhang Rouhi, Arghya Mukherjee, Amir Feder, Neil Gandal, Tyler Moore, and Marie Vasek. The Economics of Cryptocurrency Pump and Dump Schemes. In *Workshop on the Economics of Information Security*, 2019.
- [19] Eaman Jahani, Peter M. Krafft, Yoshihiko Suhara, Esteban Moro, and Alex Pentland. ScamCoins, S*** Posters, and the Search for the Next Bitcoin™: Collective Sense-making in Cryptocurrency Discussions. *Proceedings of the ACM on Human-Computer Interaction*, 2:79, 2018.
- [20] Josh Kamps and Bennett Kleinberg. To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1):18, 2018.
- [21] Paraskevi Katsiampa. Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, 158:3–6, 2017.
- [22] Jiasun Li and William Mann. Initial Coin Offering and Platform Building. 2018.
- [23] Tao Li, Donghwa Shin, and Baolian Wang. Cryptocurrency Pump-and-Dump Schemes. 2019.
- [24] Tom C.W. Lin. The New Market Manipulation. *Emory Law Journal*, 66:1253–1314, 2016.
- [25] Igor Makarov and Antoinette Schoar. Trading and Arbitrage in Cryptocurrency Markets. 2018.
- [26] SEC. Ponzi schemes Using virtual Currencies Ponzi Schemes Generally. *Investor Alert*, (153), 2017.
- [27] SEC. SEC Files Subpoena Enforcement Against Investment Company Trust and Trustee for Failure to Produce Documents. 2018.
- [28] Lin Song, Peter Langfelder, and Steve Horvath. Random generalized linear model: a highly accurate and interpretable ensemble predictor. *BMC bioinformatics*, 14:5, 2013.
- [29] Oscar Williams-Grut. ‘Market manipulation 101’: ‘Wolf of Wall Street’-style ‘pump and dump’ scams plague cryptocurrency markets. *Business Insider*, 2017.

Appendix

	RF1	RF2	RF3
<i>caps</i>	8.52	8.53	7.67
<i>return1h</i>	4.60	6.65	7.30
<i>return3h</i>	2.88	3.83	4.62
<i>return12h</i>	2.89	3.22	3.88
<i>return24h</i>	2.53	2.59	2.63
<i>return36h</i>	2.45	2.84	3.68
<i>return48h</i>	3.84	3.95	4.17
<i>return60h</i>	2.65	2.71	3.10
<i>return72h</i>	2.55	2.95	3.60
<i>volumefrom1h</i>	2.22	2.45	2.84
<i>volumefrom3h</i>	1.53	1.34	1.21
<i>volumefrom12h</i>	1.58	1.41	1.14
<i>volumefrom24h</i>	1.70	1.60	1.38
<i>volumefrom36h</i>	1.85	1.68	1.37
<i>volumefrom48h</i>	1.84	1.69	1.41
<i>volumefrom60h</i>	1.93	1.81	1.53
<i>volumefrom72h</i>	1.95	1.87	1.66
<i>volumeto1h</i>	2.64	3.27	3.19
<i>volumeto3h</i>	1.85	1.87	1.60
<i>volumeto12h</i>	1.86	1.70	1.45
<i>volumeto24h</i>	2.23	2.07	1.79
<i>volumeto36h</i>	2.42	2.18	1.87
<i>volumeto48h</i>	2.31	2.19	1.86
<i>volumeto60h</i>	2.40	2.30	2.01
<i>volumeto72h</i>	2.81	2.51	2.16
<i>returnvola3h</i>	2.02	2.30	3.07
<i>returnvola12h</i>	2.08	1.94	1.87
<i>returnvola24h</i>	2.17	1.96	1.73
<i>returnvola36h</i>	2.22	1.99	1.78
<i>returnvola48h</i>	2.44	2.10	1.69
<i>returnvola60h</i>	2.39	2.17	1.80
<i>returnvola72h</i>	2.30	2.09	1.67
<i>volumefromvola3h</i>	1.39	1.34	1.31
<i>volumefromvola12h</i>	1.57	1.42	1.16
<i>volumefromvola24h</i>	1.65	1.51	1.25
<i>volumefromvola36h</i>	1.75	1.55	1.21
<i>volumefromvola48h</i>	1.81	1.56	1.22
<i>volumefromvola60h</i>	1.79	1.56	1.25
<i>volumefromvola72h</i>	1.81	1.66	1.33
<i>volumetovola3h</i>	1.86	2.06	1.96
<i>volumetovola12h</i>	1.77	1.74	1.52
<i>volumetovola24h</i>	2.10	1.94	1.70
<i>volumetovola36h</i>	2.16	1.94	1.65
<i>volumetovola48h</i>	2.12	1.96	1.64
<i>volumetovola60h</i>	2.15	1.99	1.67
<i>volumetovola72h</i>	2.26	2.04	1.70
<i>lastprice</i>	2.14	2.02	1.66
<i>age</i>	2.20	1.88	1.69
<i>pumpedtimes</i>	1.31	1.65	2.52
<i>rating</i>	1.77	1.64	1.37
<i>WithdrawFee</i>	0.73	0.71	0.63
<i>MinWithdraw</i>	1.02	1.03	0.98
<i>MaxWithdraw</i>	0.43	0.36	0.28
<i>MinBaseTrade</i>	0.00	0.00	0.00

Figure 17: Features' importance indicated by mean decrease in Gini coefficient. Higher importance is marked by darker cell color.

	GLM1	GLM2	GLM3
<i>caps</i>	0.00	-	-
<i>return1h</i>	2.76	4.75	5.02
<i>return3h</i>	-0.04	-	-
<i>return12h</i>	1.08	-	-
<i>return24h</i>	-4.81	-	-
<i>return36h</i>	1.41	0.11	-
<i>return48h</i>	3.64	2.33	-
<i>return60h</i>	0.07	-	-
<i>return72h</i>	1.21	-	-
<i>volumefrom1h</i>	0.00	-	-
<i>volumefrom3h</i>	-0.00	-	-
<i>volumefrom12h</i>	-	-	-
<i>volumefrom24h</i>	-	-	-
<i>volumefrom36h</i>	-	-	-
<i>volumefrom48h</i>	0.00	-	-
<i>volumefrom60h</i>	-	-	-
<i>volumefrom72h</i>	-	-	-
<i>volumeto1h</i>	1.61	-	-
<i>volumeto3h</i>	5.99	-	-
<i>volumeto12h</i>	-	-	-
<i>volumeto24h</i>	-	-	-
<i>volumeto36h</i>	-	-	-
<i>volumeto48h</i>	-	-	-
<i>volumeto60h</i>	-2.88	-	-
<i>volumeto72h</i>	-0.49	-	-
<i>returnvola3h</i>	3.94	-	-
<i>returnvola12h</i>	4.41	-	-
<i>returnvola24h</i>	-9.39	-	-
<i>returnvola36h</i>	10.40	-	-
<i>returnvola48h</i>	9.10	-	-
<i>returnvola60h</i>	-12.57	-	-
<i>returnvola72h</i>	-3.93	-	-
<i>volumefromvola3h</i>	-0.00	-	-
<i>volumefromvola12h</i>	0.00	-	-
<i>volumefromvola24h</i>	-	-	-
<i>volumefromvola36h</i>	-	-	-
<i>volumefromvola48h</i>	-	-	-
<i>volumefromvola60h</i>	-	-	-
<i>volumefromvola72h</i>	-0.00	-	-
<i>volumetovola3h</i>	-7.46	-	-
<i>volumetovola12h</i>	1.32	-	-
<i>volumetovola24h</i>	-9.96	-	-
<i>volumetovola36h</i>	-2.13	-	-
<i>volumetovola48h</i>	18.83	-	-
<i>volumetovola60h</i>	-	-	-
<i>volumetovola72h</i>	8.65	-	-
<i>lastprice</i>	-91.74	-	-
<i>age</i>	0.00	-	-
<i>pumpedtimes</i>	0.69	0.66	-
<i>rating</i>	-0.16	-	-
<i>WithdrawFee</i>	-0.00	-	-
<i>MinWithdraw</i>	-0.00	-	-
<i>MaxWithdraw</i>	0.00	-	-
<i>MinBaseTrade</i>	-	-	-
(Intercept)	-5.43	-6.15	-5.95

Figure 18: Variable coefficients (unstandardized) using GLM. Coefficients of variables not selected by the model are shown as "-".