Financial Impact of Washtrading in the NFT Market: An Algorithmic Approach

Efstathios Polyzos [a,b] and Mieszko Mazur [c,d]

Abstract

This paper investigates potential market manipulation in the market for Non-Fungible Tokens (NFTs). Specifically, we focus on washtrading strategies among NFT investors. First, we propose a relatively straightforward methodology to detect washtrading transactions. Second, we test our methodological approach using data from roughly 46 million NFT sales (largest sample ever studied). We find that between 2018 and 2022 washtrading artificially inflates NFT volume as well as NFT valuations. Moreover, we find that NFTs from the washtrading subsample yield abnormally high returns. Our results suggest that NFT washtrading is less rampant than commonly believed, although we find evidence that in some recent cases washtrading constitutes more than 25% of the aggregate USD volume.

JEL: G12; G23; G18

Keywords: Market manipulation; Non-fungible tokens; NFT; Washtrading; Cryptocurrencies

[a]Zayed University, United Arab Emirates
[b]Centre for Applied Macroeconomic Analysis, Australia
[c]ESSCA School of Management, France
[d]Corresponding Author

1 Introduction

A Non-Fungible Token (NFT) can be defined as a digital certificate of ownership written permanently into the blockchain. Technically, an NFT represents a non-fungible (and thus unique) asset - be it digital or real - including an artwork, collectible, real-estate, or virtually any type of financial and legal instrument. NFTs gained sudden popularity in 2021 and have been actively traded ever since, reaching a global trading volume of roughly $60 billion.

In this paper, we investigate abnormal behavior in the NFT market. More specifically, we examine potential washtrading activity - buying and selling NFTs between at least two NFT blockchain accounts over a very short time window with the objective to manipulate the market. To this end, we: 1) develop a novel and relatively simple methodological approach that helps us identify suspicious trading activity; and 2) use a large and comprehensive sample of tens of millions of NFTs that have not been systematically examined to date. This very large dataset allows us to test the robustness of our findings and facilitates more accurate inference about the NFT market quality. We draw on existing evidence from the traditional financial markets that provides a useful and important framework for our basic understanding of the washtrading behavior among market participants.

We establish several novel stylized facts about the NFT trading behavior in the context of market manipulation. First, we document that washtrading is far from being standardized and exhibits differing patterns that emerge over time, across NFT collections, as well as across other dimensions (e.g., NFT price and volume). Intriguing, we find that at the end of 2022 washtrading accounts for abnormally high percentage of all NFT transactions reaching on several occasions up to 25% of the aggregate dollar volume. Second, we find that if the frequency of incidence of the washtrading cluster (group of suspicious NFT accounts) is above a certain reasonable threshold, then most of the doubtful NFT transactions by dollar volume take place within the three-hour trading period. Moreover, if we keep extending the estimation window by a fixed length, we find the the dollar volume of the NFT washtrading diminishes steadily. Even if we increase the window length up to 30 hours, the process does not add any significant incremental informativeness to our baseline results. To the contrary, larger windows tend to reduce the amount of the valuable information about suspicious NFT trading. It should be stressed that the median NFT is sold within 98 hours (mean being several orders of magnitude larger) implying that washtraders clear the market for the specific NFTs thirty times quicker compared to a typical NFT trade based on the full sample. The above suggests that the time in which NFT becomes liquidated by the identified washtrading cluster in a sequence of transactions (time-to-sale) can be considered an important indicator of the potential washtrading activity in the NFT market. Furthermore and consistent with the patterns observed for the NFT trading volume, we find that the NFT price levels are highest for NFTs traded within the three-hour trading period and decrease steadily together with the length of the estimation window. Thus, the increased trading volume due the perceived NFT market manipulation tends to be coupled with higher NFT prices, the evidence often found for washtrading activity in the traditional financial markets.

When we look at the realized returns from the subsample of the suspicious NFT transactions, we find that they are several orders of magnitude higher relative to the mean return on a typical NFT that hovers around 3%. Further, we notice that washtraders focus on high-end NFTs valued in excess of $200,000 at the mean and $1,500 at the median. Arguably, high-priced NFTs are a common target of washtrading activity because elevated dollar volume serves the purpose of attracting new high net-worth NFT buyers lured to the NFT market by the expectation of abnormally high gains.

It should be emphasized that the methodology used throughout the paper is not immune to the general criticism of the omitted information bias and most likely contains both Type 1 and Type 2 errors. In other words, not all the transactions captured by our model are washtrades, and vice versa, some transactions coded as washtrades might be legitimate NFT sales. In one way or the other, the simplicity of our empirical approach inevitably leaves some potentially interesting questions unanswered.

This paper contributes to the literature in several ways. First, we develop a relatively simple methodology for identifying suspicious transactions in the NFT market. More specifically, we focus on washtrading, a market manipulation trading strategy designed to deceive investors about asset trading volume and its valuation. We add to the scant literature on this issue (e.g., von Wachter, Jensen, Regner, and Ross, 2022). Second, we gather a large and comprehensive sample of roughly 46 million NFT transactions. To the best of our knowledge, this is the largest NFT sample ever evaluated in the academic literature. In contrast, existing studies use a very limited amount of NFT data that typically encompass one or, at most, several NFT collections (e.g., Kong and Lin, 2023; Borri, Liu, and Tsyvinski, 2023). Third, we extend the literature on washtrading in general by presenting evidence on the washtrading behavior in the new and unregulated NFT market that resembles early US stock exchanges and their lax and easily circumvented regulation (Mahoney, 1999; Jiang, Mahoney, and Mei, 2005; Cumming, Johan, Li, 2011). Overall, extant empirical evidence is surprisingly sparse, lacks a comprehensive analysis, and our work contributes to a better understanding of NFT market in general and washtrading practices in the NFT market in particular.

The remainder of the paper is organized as follows. Section 2 discusses washtrading in traditional financial markets viewed as a backdrop to our analysis. Section 3 presents data sources and descriptive statistics. Section 4 describes the methodological approach, while Section 5 discusses the results. Section 6 concludes the paper.

2 What is washtrading?

According to the Securities and Exchange Commission (SEC) washtrading involves purchasing and selling securities that "match each other in price, volume and time of execution, and involve no change in beneficial ownership[1]". Washtrading is closely related to at least two other market manipulation activities: matched orders and wash selling. Matched orders can be seen as washtrading performed by a number of different individuals or entities trading securities directly or indirectly within the group. Washselling, on the other hand, occurs when the investor trades securities to realize a loss and subsequently intends to offset this loss against future gains on substantially identical securities with the aim of minimizing tax liabilities. In this latter case, there is a strict time window within

[1] https://www.sec.gov/files/litigation/aljdec/id82grl.txt

which a trade can be classified as washselling[2].

In our study, we extend the definition of washtrading to include also matched orders. The reason we do so is that it is empirically unfeasible to identify whether the accounts from which NFTs are traded belong to a single or multiple individuals or entities. By construction, the NFT accounts may (and usually do) remain anonymous. Setting up an account does not require the verification of the true identity. To the extent that the NFT trades are manipulated, it is reasonable to assume that significant part of this type of activity occurs in groups of traders akin to the so-called ”pools” in the New York Stock Exchange (NYSE) in the 1920s and analyzed in Mahoney (1999). The aspect of the pool we are particularly interested is the presumable collusion among the traders that constitute a pool and its impact on the NFT market characteristics (i.e., price and volume).

Generally, we classify washtrading as buying and selling of an NFT at least twice by at least two different NFT accounts in a very short period of time. Furthermore, we posit that the NFT trade between two accounts is equivalent to an NFT trade between three different accounts (or more), if each type of an event take place at least twice over a short time horizon. Specifically, if account A sells an NFT to account B, and then B sells the NFT back to A within the three hours' time then this chain of transactions in our framework might be considered washtrading. Similarly, if A sells and NFT to B, and then B quickly sells an NFT to C then we classify this sequence as washtrading on the condition that this specific sequence repeats twice within the three-hour window. As mentioned already, we assume that trading is executed between different parties that cluster in time with a relatively high frequency, keeping in mind that a traded asset is of non-fungible nature and displays infrequent trading patterns. The fact that trades happen reasonably fast is a key indicator for a suspicious trading activity that could raise a red flag.

We caution the reader that, legally, washtrading pertains to securities, however, NFTs have not been categorized as such. Despite a lack of clarity of what NFTs are from

[2] https://www.sec.gov/litigation/admin/2018/ia-5086.pdf

the asset classification standpoint, the lifetime NFT trading volume as of June 2023 has reached roughly $60 billion[3]. Large and prominent global corporations have issued and sold NFTs to their loyal consumer base[4]. It is also worth emphasizing that NFTs have been proliferating faster and into the broader set of sectors (including arts and entertainment) compared to classic cryptocurrencies.

2.1 Washtrading in traditional financial markets

To understand better the washtrading behavior, we first turn to traditional financial markets that have a long-standing tradition of market manipulation. Unfortunately, to date there has been no systematic study investigating washtrading practices in securities trading including equities, bonds, and derivatives. Instead, we rely mostly on source documents (complaints) issued by the Securities and Exchange Commission (SEC) and made available on the SEC website. These documents are the result of the SEC investigations and report the details of the market manipulation techniques. We analyze these documents closely and gather data that allow us to understand better the underlying mechanism, motivation, and consequences of washtrading on the traditional financial markets. We then use this knowledge to explore the washtrading behavior in the NFT market.

2.2 Motives for washtrading

It appears that the principal motive behind washtrading is to augment security's trading volume. Artificially inflated volume accompanying washtrading is highlighted in almost every SEC complaint over the last three decades[5]. Increased volume creates interest in a manipulated security and captures the attention of other investors, especially if the security is thinly traded. As a result, investors might be more inclined to trade the

[3]Based on the ten largest NFT exchanges by volume tracked in https://dappradar.com/nft/marketplaces?period=all. The difference is size between the first and second (last) NFT exchange on the list is roughly 7(53)-fold.

[4]https://www.voguebusiness.com/technology/louis-vuitton-to-sell-euro39000-nfts

[5]See e.g., https://www.sec.gov/files/litigation/opinions/3440726.txt; https://www.sec.gov/litigation/complaints/2021/comp-pr2021-195.pdf

security and eventually affect its price level, even though price manipulation might not be the original intent of the washtrader.

Manipulating trading volume may lead to all kinds of direct and indirect benefits for the parties involved. Even if the direct benefits happen to be non-monetary, ultimately the benefits linked to washtrading make perpetrators better off. For example, washtrading, despite being directly unprofitable, may help meet margin requirements[6] and preserve capital value of the investor who otherwise would have incurred substantial losses. On the other hand, direct benefits of washtrading may pertain to e.g., rebates for market making[7] or sales credit[8]. In any case, it is perhaps unreasonable to speculate that artificially created trading volume has no ultimate effect on prices even if unintentional.

Obviously, there are washtrading cases where price manipulation is of the first-order importance[9]. For example, according to Mahoney (1999) the first step in the washtrading process used by the "pools" at the NYSE was to place a large enough order to move prices. However, a non-trivial portion of the artificial price levels tend to be induced (among others) by creating a false impression of trading activity through volume manipulation.

2.3 Commonalities between washtrading in the NFT and traditional financial markets

A casual inspection of the SEC complaints leads to the following conclusions. First, washtrading in its various manifestations was invented and tested in traditional financial markets, and only later on adopted to the context of NFTs. Second and somewhat surprising, there are many commonalities between washtrading in the traditional and NFT markets including: (1) use of multiple trading accounts in control of a single individual, entity or a group of investors, (2) manipulation of trading volume for direct or indirect benefit, (3) relatively short life of the washtrading activity.

[6]https://www.sec.gov/files/litigation/opinions/3440726.txt
[7]https://www.sec.gov/litigation/complaints/2021/comp-pr2021-195.pdf
[8]https://www.sec.gov/files/litigation/admin/2011/34-63964.pdf
[9]https://www.sec.gov/litigation/complaints/2023/comp25620.pdf

3 Data, sample, and descriptive statistics

Data for this paper come from the largest NFT marketplace called OpenSea[10], with the lifetime NFT trading volume of roughly $36 billion [11]. OpenSea has been in existence since 2018, but the NFT market has seen a boost in 2021, when most individuals became aware of this new asset class. This important change in the market is evident in Table 1, where the number of collections, as well as the total market capitalization increased abruptly in 2021. In addition, while the number of listed collections rose only by approximately 14% from 2021 to 2022, the total market capitalization more than doubled, driven by both a higher trading volume (as attested by the increase in soft supply) and higher asset values, as well as by increasing crypto values, albeit to a lesser extent. If we examine individual assets (Table 2), we can see that from our sample of 24 million NFT assets, the median asset has been traded only once for a trading volume of $56.74. However, we can see that the distribution of trading volumes is dominated by extreme values, particularly at the right tail since the mean is above the 90$^{\text{th}}$ percentile. In addition, using a sample of 9.6 million NFTs that have been traded more than once, investors hold on to their NFTs for a median duration of approximately 23 days (560.28 hours), with the distribution having a fat right tail.

We note here that in all tables with descriptive statistics, we report the 10$^{\text{th}}$ and 90$^{\text{th}}$ percentiles, instead of reporting the minimum and maximum values. This is an expositional choice aimed at cleaning the data from extreme prices. More specifically, there are many transactions where the price is extremely small ($<10^{-9}$, which is rounded to 0 in calculations presented in Table 2). This results in extremely high results in the returns data series and thus reporting the 10$^{\text{th}}$ and 90$^{\text{th}}$ percentiles provides more insights into the nature of the underlying data series.

We mine the OpenSea API for NFT transaction data. Each downloaded transaction holds information about the sale, including time and date, price (in crypto token and in USD), as well as details of the underlying asset and the collection it belongs to. In

[10] https://opensea.io/
[11] www.dappradar.com

Table 1: Evolution of the NFT market (Entire Sample)

Variable	2019	2020	2021	2022
Number of Collections	1,785	6,915	478,967	545,432
Soft Supply	2.4%	6.9%	4.9%	8.6%
Market Cap (USD)	$164.69 M	$366.66 M	$5,036.08 M	$10,370.13 M
Market Cap (ETH)	147.15 M	226.70 M	299.58 M	464.13 M

This table demonstrates stylized facts regarding the NFT market. The Number of Collections represents the size of the market, in terms of the existing NFT collections publicly available for trading in the OpenSea market. The Soft Supply represents the proportion of assets in an NFT collection that have been traded at least once (estimated by the number of unique NFTs that have changed hands in the secondary market divided by the number of NFTs in a collection). Finally, the market cap measures the total market capitalization, estimated using the last sale price of each asset, aggregated across all collections.

Table 2: Descriptive statistics for NFT turnover and return (Entire Sample)

Variable	N	Mean	Median	Std. Dev.	10th Perc	90th Perc
NFT Volume (USD)	24,725,313	1,438.28	56.74	488,867.79	0.00	1,061.83
NFT Volume (ETH)	24,725,313	28.13	0.05	3,248.33	0.00	1.04
NFT Turnover	24,725,313	1.85	1.00	18.80	1.00	3.00
Mean Holding Period (hrs)	9,623,188	1,554.82	560.28	2,250.77	4.62	4,865.67

This table demonstrates descriptive statistics for the NFT assets in the market. The volume computed as the sum of the lifetime sale prices for the asset, while the turnover measures the number of trades. The mean holding period is the time distance between consecutive sales of the same NFT, measured in hours and calculated as a mean for all consecutive sales of the same asset. The sample size, N, represents the number of assets that have been traded at least once (for volume and turnover) or at least twice (for the holding period).

addition, the participating parties are identified by their wallet IDs. We have identified an approach that can help us locate potential washtrading transactions, the crypto wallet addresses that are associated with them, as well as the underlying NFT assets. Once washtrading transactions are identified, we are able to analyse the price changes of the traded assets and quantify the financial impact (in terms of artificial price changes) of this practice in the NFT markets.

We present descriptive statistics of the entire sample of sales transactions in Table 3. Our sample consists of 56,562,924 sale transactions mined from OpenSea[12], with the sample period starting from the first transaction available (on 23 January 2018, based on OpenSea data) and ending on 31 December 2022. We clean this data set, removing sales with zero price or sales where one of the counterparties (buyer or seller) is blank. This leaves a total of 45,625,644 transactions, which are fed through our washtrading identification approach, explained in the next section. Apart from summarising the (USD) price

[12]Based on the OpenSea API, sales are transactions that carry the event type ”successful”.

of all transactions, we also compute summaries per each individual NFT asset (hence the sample size here represents the unique assets for which we observe at least one transaction) as well as summaries per each NFT collection (hence the sample size here represents the unique collections for which there exists at least one asset with at least one transaction). We opted to use NFT asset prices in USD rather than the prices in the corresponding crypto token (a few of them are used in OpenSea) in order to ensure that prices are comparable across different time periods, given the fluctuating prices of cryptocurrencies.

We compute returns and the holding period once repeated sales are noted for an asset (i.e. when an asset is sold at least twice in our cleaned sample). The holding period is defined as the time between two consecutive sales of the same asset. Returns and the holding period require at least two sales for an asset and this is why the number of observations is smaller than the total number of transactions in our sample.

Table 3: Descriptive Statistics of Entire Sample of NFT Transactions

Variable	N	Mean	Median	Std. Dev.	10th Perc	90th Perc
USD price	45,625,644	779.4260	50.2760	38,966.3310	0.0000	702.1200
Number of Transactions per Asset	24,725,313	1.8450	1.0000	18.8030	1.0000	3.0000
Mean Price per Asset	24,725,313	377.8030	40.7600	13,657.4700	0.0000	548.2449
Median Price per Asset	24,725,313	364.2880	39.5190	10,869.4310	0.0000	526.6920
Number of Transactions per Collection	186,645	244.4510	3.0000	2,406.6610	1.0000	138.0000
Mean Price per Collection	186,645	274.6210	21.7020	5,934.1330	0.7609	253.2414
Median Price per Collection	186,645	235.2120	17.7140	5,563.4470	0.6336	222.7322
Return	20,564,902	0.3300	0.0350	0.9980	-0.6356	1.6662
Holding Period (hrs)	20,564,902	1,146.1220	98.8340	2,126.0440	0.1961	4,502.0964

This table presents descriptive statistics for all the transactions in our sample.

4 Methodological approach

In this section, we propose the empirical approach to measuring the magnitude and persistence of the washtrading activity in the NFT market. We define several observable characteristics that help us calibrate the empirical model along its key parameters. We discuss the specifics below, while the details of our computations can be found in the Appendix.

4.1 Washtrading identification strategy

Existing literature on washtrading is scant and to our knowledge there is no systematic study on washtrading in traditional financial markets from which we could draw on. This is probably due to the inevitable scarcity of data. On the other hand, the emergence of countless NFT collections and NFT exchanges, as well as the fact that NFT data are made available to the general public free of charge, offers a unique and important laboratory for investigating NFT market manipulation.

We propose an algorithmic approach to estimating the magnitude of the washtrading behavior in the NFT market based on two parameters: (1) the time period of consecutive transactions that constitute a set of washtrades, and (2) the frequency by which washtrading groups (i.e. specific accounts involved in suspicious transactions) appear in our identification process.

First of all, because NFT trading is largely anonymous, it is impossible to establish how many traders stand behind a single transaction or a series of transactions. For simplicity, we assume that each investor uses a distinct NFT trading account, even though it must be emphasized that it is feasible for any individual to control any number of NFT accounts. We believe that this assumption does not bias our inference, as the true identity of participants does not matter for our analysis. Second, the artificial market activity reflected in e.g., trading volume inherently assumes greater trading volume per unit of time. This implies that artificially generated volume must be produced through greater frequency of transacting assuming the number of participants in the process is held constant. More specifically, if only one NFT can be traded per time, then increased trading volume through deliberate market manipulation suggests that some NFTs are traded more frequently than normal.

In traditional financial markets, identification of washtrading is typically performed by identifying transactions where the buyer and seller are the same person (Imisiker & Tas, 2018) or by locating a series of transactions between multiple participants that does not effectively alter the portfolio of the group (Pouncy, 1994; Cumming et al., 2011; Cao et al., 2015). In crypto markets, as mentioned earlier, multiple participants (crypto wallets) can effectively belong to the same person and thus this approach may yield limited results and result in underestimating of the issue (Cong et al., 2021; Le Pennec et al., 2021).

Figure 1: Example of a Suspicious Cluster of Transactions

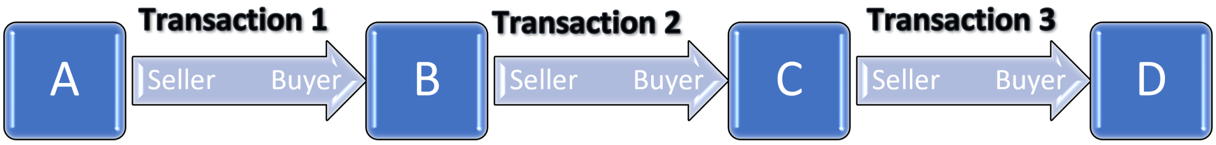

Note:This figure demonstrates the transactions in a washtrading cluster. All the transactions should occur within the specified time window.

In our approach, washtrading clusters comprise of a series of transactions where the buyer in one step is the seller in the next step, as demonstrated in Figure 1. The transactions in the cluster do need necessarily need to be for the same asset, since it is possible that washtraders may try to create artificial liquidity in the market to reduce the liquidity premium and make the assets more attractive (Gerhold et al., 2014; Darolles et al., 2015). Hence, a cluster begins when the buyer in a transaction acts as a seller within the given time frame. This will yield the next transaction(s) in the chain and the chain continues in the same manner until all transactions in the time frame are accounted for. The cluster are categorized as closed when the same user appears in two places in the chain (i.e. A sells to B, B sells to C, C sells to A) and they are categorised as open if each user appears only once in the chain.

The above strategy identifies clusters of transactions but not all such clusters are necessarily involved in washtrading. As series of transactions (even for the same asset) could, in theory, occur at random, particularly as the time window increases. However, if the cluster was identified randomly, the participating accounts (wallet IDs) should normally appear in the identified clusters only once. If a given combination appears more than once (i.e. if the same group of accounts engages in a chain of transactions similar to Figure 1 in the given time frame more than once), then there is strong evidence of colluding behavior, suggesting some form of market manipulation. We note that the combination of accounts does not need to appear in the same order (i.e. in the example provided in Figure 1, the order in another cluster could be e.g. B, C, A, D - this would be identified as the same combination of accounts). When observing the number of times that the same combination of accounts has been found in a cluster, we consider different cutoff values, which we term ”cutoff frequency”.

Once the final washtrading clusters have been identified and we have, thus, isolated the underlying transactions, we can proceed to analyse the data further by computing the return achieved by the washtrading practice and by examining the assets and collections that are being traded. Regarding the return, we consider the price of the asset during its first transaction in the cluster and compare it to the price recorded in the first transaction outside the cluster. Any further transactions of the same asset within the cluster are ignored as spurious. We compute the washtrading return using Equation **??**, using the price in the first ”clean” transaction (i.e. outside the cluster) as the current price and the price in the first washtraded transaction as the previous price. As mentioned before, we use asset prices in USD (as opposed to prices in crypto tokens) to compute the returns for standardization purposes in order to avoid complications arising from fluctuations in the values of cryptocurrencies.

Since the current paper introduces this approach and thus uses new terminology, the reader can refer to Table 4 for a dictionary of the terms used in our discussion.

Table 4: Detailed definitions of the variables used in the study

Variable	Definition
Suspicious Cluster	A chain of transactions where the buyer in one step is the seller in the next step and which occur with the washtrading time frame (see Figure 1).
Washtrading Cluster	A suspicious cluster which fulfils the conditions to be characterised as washtrading. More specifically, the accounts involved appear in suspicious clusters more times than the cutoff frequency.
Washtrading Group	A group of accounts that is involved in suspicious clusters more times than the cutoff frequency
Washtrading Time Frame	The time period within which a chain of transactions where the buyer in one step is the seller in the next one is characterised as suspicious.
Cutoff Frequency	The number of times a certain group of accounts (washtrading group) has to appear in suspicious clusters before these clusters can be considered washtrading clusters.

This table presents a dictionary of the terms used in this paper.

5 Results

5.1 Model selection

We start the process by selecting the appropriate model. As mentioned above, there are two critical choices that we need to make: (a) duration of the washtrading time frame, and (b) cutoff frequency. Consequently, we examine different choices for the duration of the washtrading time frame, ranging from 3 to 30 hours. Regarding the cutoff frequency, we initially considered the value of 2 for our identification approach. Given randomness, we should not be able to observe the same combination of accounts in a chain of transactions during the washtrading time frame more than once. This would suggest that observing two clusters with the same accounts should immediately result in classifying these as washtrading clusters. However, we considered other alternatives, ranging from 3 to 10, and found that the results differ significantly.

The results of this process are demonstrated in Figures 2 and 3. In Figure 2, we see the number of washtrading groups and the numbers of transactions in the washtrading clusters of these groups. First, we note that, except for the 30-hour window, changing the duration of the washtrading time frame does not significantly impact the results. In all cutoff frequencies, the results are similar across different durations. This means that lengthening the washtrading time frame does not provide any extra results in terms of washtrading identification. Furthermore, in higher cutoff frequencies the number washtrading groups and transactions is actually higher for the shorter time window. This occurs because as the time frame increases, it is easier for chains of transactions to be formed, perhaps with longer gaps between them, since the time permitted for cluster identification is longer. The identified clusters (”suspicious” clusters, based on Table 4), however, include different combinations of accounts and, thus, when the frequency cutoff is applied, they are removed from the identification process and not considered washtrading.

We also see a significant increase in the identified washtrading transactions and the groups involved when the duration of the time frame increases to 30 hours. This result is more evident in lower cutoff frequencies and it is intuitively accurate. Increasing the du-

Figure 2: Number of Groups and Clusters for Different Combinations of Model Parameters

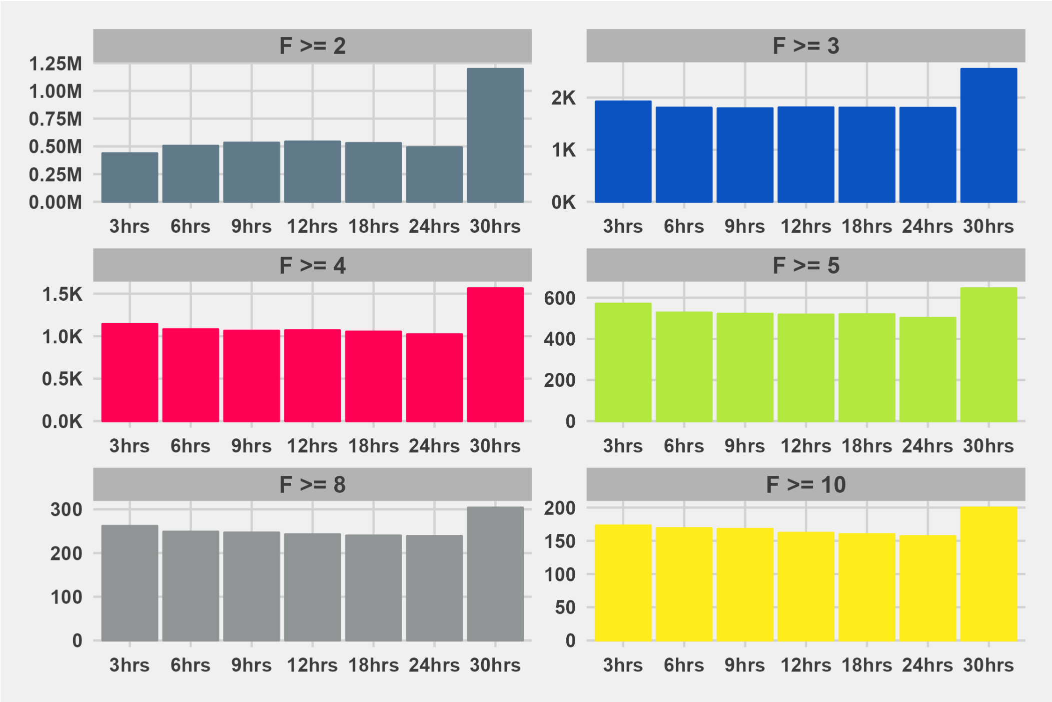

(i) Washtrading Groups

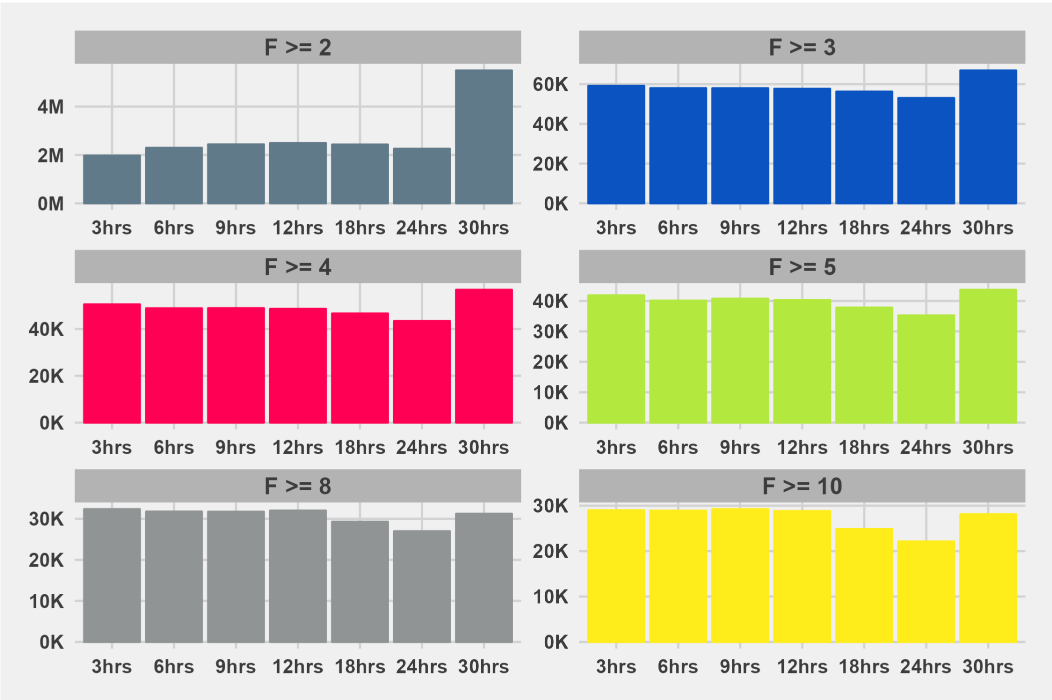

(ii) Washtrading Transactions

The figures above show the number of washtrading groups and washtrading clusters resulting from the identification process when different model parameters are being used. These parameters are the duration of the washtrading time frame (different values demonstrated on the horizontal axis) and the cutoff frequency (different figures in each panel).

ration of the time frame means that it is easier for chains of transactions to be formed but longer gaps are observed between individual transactions. However, we believe that this finding might be spurious and does not necessarily constitute washtrading, particularly given that the difference is lower as the cutoff frequency increases.

In addition, Figure 2 demonstrates the important difference of the results as we increase the cutoff frequency from 2 to 3, which reduces the number of washtrading transactions by approximately 99%. Erring on the side of caution, we choose to use a cutoff frequency greater than 2. In the two panels, we can see that the number of washtrading transactions drops somewhat when we change the frequency from 3 to 4, but there is no significant loss of information after that. The implication here is that groups of accounts that appear in at least 4 different washtrading clusters (i.e. the same combination of accounts has at least 4 different chains of transactions), then they tend to work together more consistently. We thus select the cutoff frequency to be 4, as the more conservative choice that balances out the risk of a Type I and Type II error in the identification process.

In Figure 3[13], we examine the impact of our choices regarding the duration of the time frame and the cutoff frequency on other outcomes of the identification process. We can see that for unique accounts, collections and assets involved in washtrading, the numbers are significantly smaller as the cutoff frequency increases, with the difference being more pronounced in the lower range of the frequencies. However, the number of unique accounts does not increase, as we increase the duration of the time frame. These results are similar for unique collections and unique assets involved in washtrading. This substantiates our earlier claim that examining longer time frames for washtrading identification does not improve the results as the information gain is minimal, if any, thus confirming our choice to use 3 hours as the time frame.

This choice is also confirmed by the results regarding the mean and median prices of the washtrading transactions. The corresponding panels in Figure 3 provide two insights. First, increasing the duration of the time frame does not seem to change the mean and median prices significantly. There is a slight decrease, as the time frame duration increases (the path is actually somewhat U-shaped for the mean price), but the impact is not very

[13]The figure does not include the results for the cutoff frequency of 2 and the time frame of 30 hours in order to make it easier to read. We supply detailed results for each variable in Figure 3 in Appendix A

Figure 3: Identification Results for Different Combinations of Model Parameters

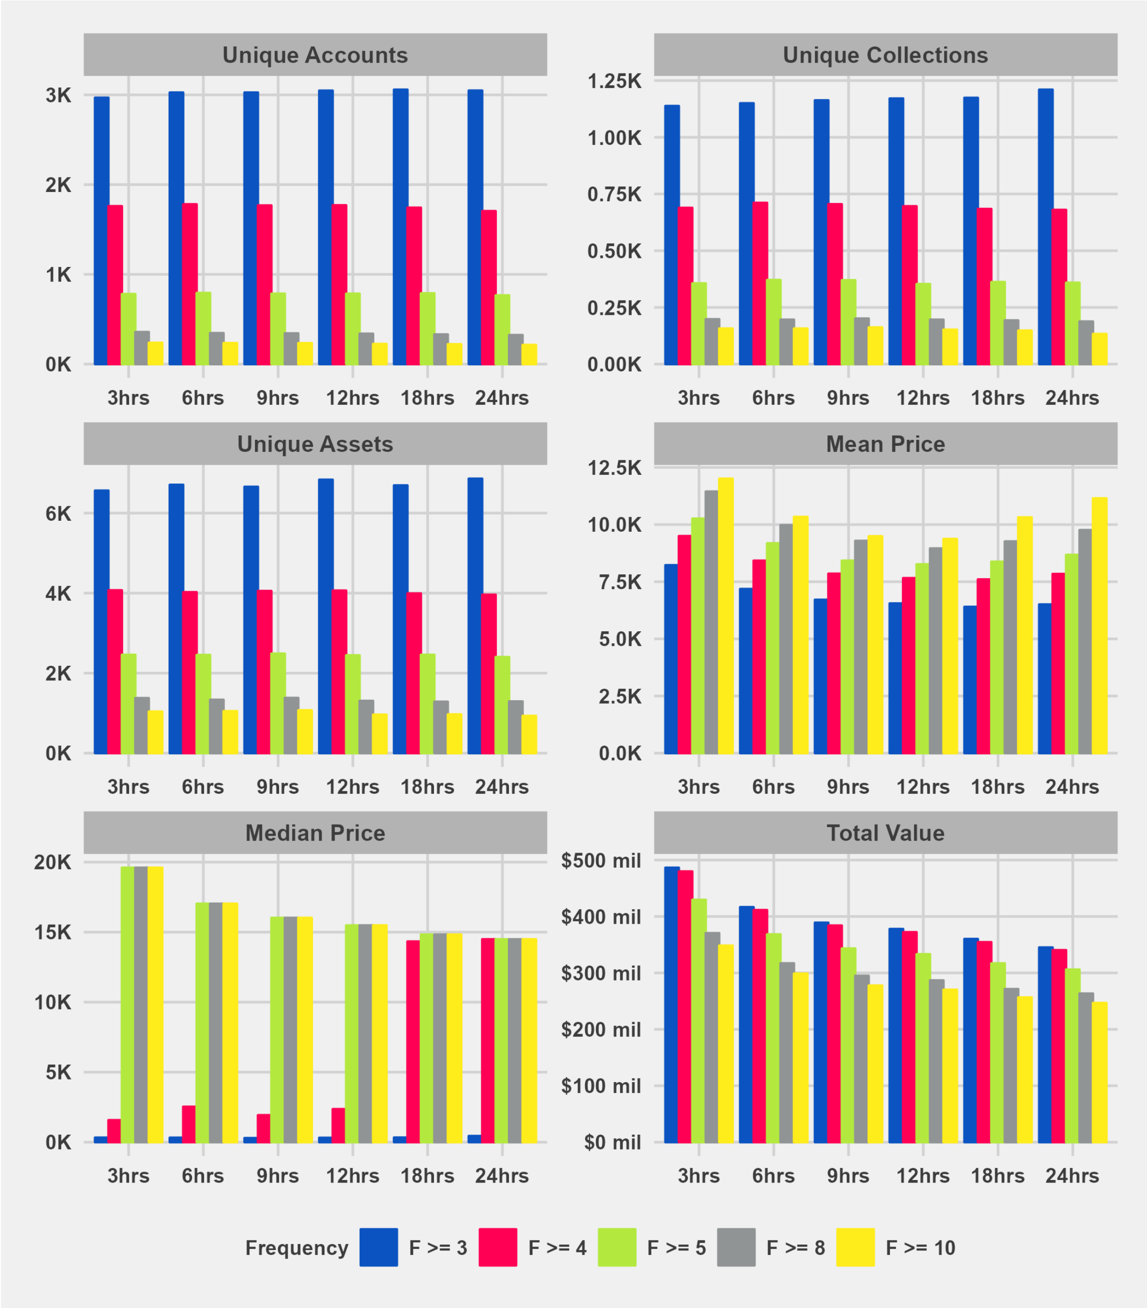

The figures above show the results of the identification process when different model parameters are being used. These parameters are the duration of the washtrading time frame (different values demonstrated on the horizontal axis) and the cutoff frequency (different figures in each panel). The figure does not include the results for the cutoff frequency of 2.

large. On the other hand, increasing the cutoff frequency results in increasing mean and median washtrading prices with the change from 4 to 5 resulting in a huge increase in the median price. We will explain what this means regarding the behaviour of washtraders later on, but from the technical perspective, this confirms that the cutoff frequency should be higher than 3, as this improves the results of the identification process. Given the significant impact of the change from 4 to 5, there is perhaps an argument to be made here for setting the cutoff frequency to 5. However, given the results shown in Figure 2, we believe that this would result in a loss of information since the number of washtrading clusters is significantly reduced. At the same time, as mentioned before, we would lose information on unique accounts, assets and collections. Hence, we maintain our choice of 4 for the cutoff frequency. Finally, we note that the total value of washtrading transactions is higher for the 3-hour window.

5.2 Characteristics of NFT washtrading

Table 5: Descriptive Statistics in Washtrading Transactions

Variable	N	Mean	Median	Std. Dev.	10th Perc	90th Perc
USD price	50,516	227,710.9650	1,573.8100	481,477.3620	0.8813	674,724.6022
Number of Transactions per Asset	4,071	12.4090	2.0000	89.6520	1.0000	8.0000
Mean Price per Asset	4,071	35,529.2350	25.3530	500,517.2570	0.0000	6,360.9705
Median Price per Asset	4,071	34,957.0860	25.2500	497,872.6350	0.0000	6,393.0040
Number of Transactions per Collection	689	73.3180	6.0000	543.9350	1.0000	58.4000
Mean Price per Collection	689	12,189.8450	15.8020	78,640.9260	0.0000	1,656.9797
Median Price per Collection	689	11,929.9660	13.3020	79,603.9210	0.0000	1,587.9440
Return	5,614	15.3170	0.0000	909.7770	-0.0961	0.5350

This table presents descriptive statistics for the washtrading transactions identified as a result of our approach.

Table 5 demonstrates the descriptive statistics for the transactions in involved in washtrading, based on our identification strategy. For consistency and ease of comparison, the table follows a similar format with Table 3. A first observation relates to the price of the traded assets, which is much higher in washtrading transactions. Both the mean and the median are higher, confirming our previous suggestion, that washtraders in the NFT market trade with higher prices, since they cannot benefit from the traded volume of the underlying asset. To further demonstrate the differences between the distribution of variables, we present violin plots of the two samples (i.e. the entire sample of transactions and the washtrading sample) in Figure 4. The violin plots is similar to a box plot, but shows the entire distribution of the data and thus is very helpful particularly when the distribution is multimodal (i.e. has multiple peaks). We note that the data is winsorized at the 10^{th} and the 90^{th} percentiles for consistency with the previous tables and to improve the readability of the graphs. The figures show that while the shape of the distributions is similar across the two samples, they have different magnitude, reflecting higher prices and more transactions per asset in the washtrading sample.

Figure 4: Data Distribution of Summary Statistics

(i) USD Price

(ii) Median Price per Asset

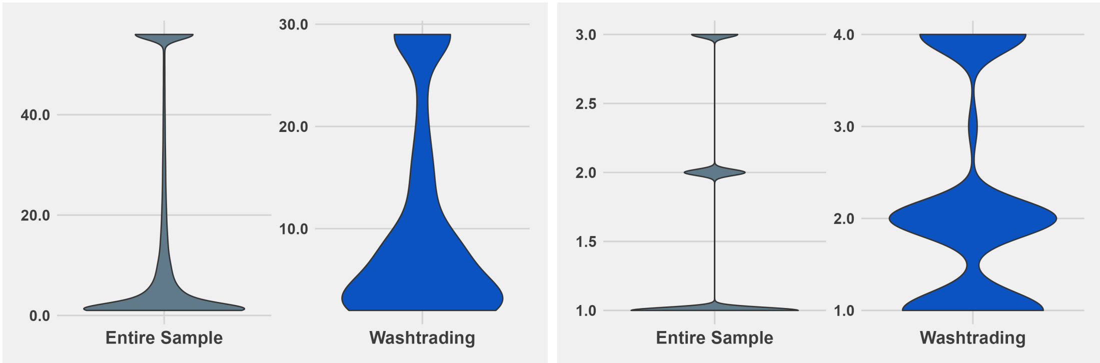

(iii) Number of Transactions per Collection

(iv) Number of Transactions per Asset

This figure demonstrates violin plots for the prices (for each transaction), the median price per asset and number of transactions per collection and per asset, as computed over the entire sample of transactions (grey plot) and the washtrading transactions (blue plot). The vertical axes show the values of the corresponding variable, while the width of the figure depicts the frequency density.

We also note that only a small subset of assets and collections is involved in washtrading (4,071 unique assets in the washtrading transactions vs 24,725,313 in the entire sample, 689 unique collections vs 186,645), confirming that this practice, while popular, probably does not involve the entire market. We have more transactions per asset and per collection, with higher standard deviations, suggesting fat tails at the upper end of each distribution (since the lower end is bound by 1). Also, the distribution of mean and median prices over the sample of assets and collections is higher. Finally, mean returns are higher under washtrading. We also note the presence of some negative returns in washtrading, suggesting that the practice is not always successful in generating profits for washtraders.

Figure 5: Washtrading Transactions in the NFT Market

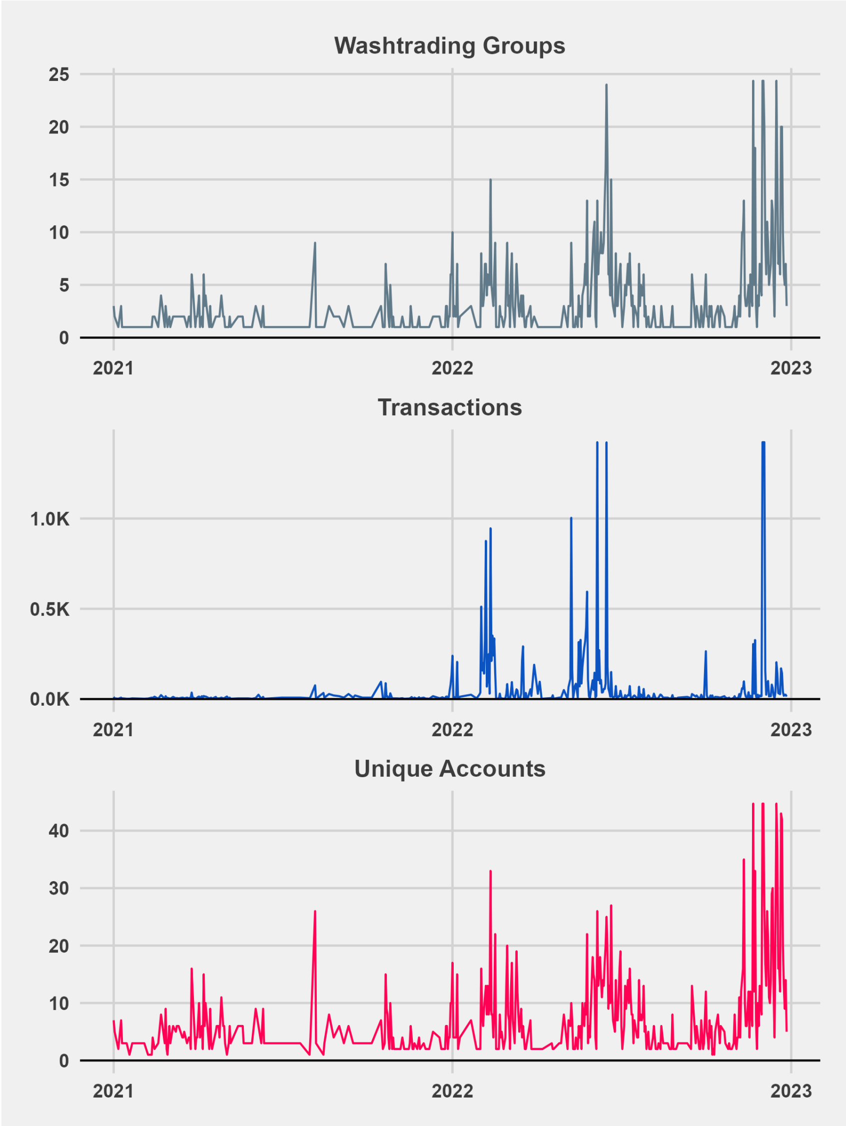

This figure shows the number of groups, the number of transactions and the number of unique accounts involved in washtrading, based on our identification process. The numbers are absolute numbers calculated with daily frequency.

The results of the washtrading identification process are also demonstrated in Figures 5 and 6, where we display the results across time, based on daily aggregations. We note that washtrading activity is not constant but has time-varying behavior. In Figure 5, we note peaks and troughs in washtrading numbers with all three metrics (the numbers of groups, the number of transactions and the number of unique accounts) moving roughly symmetrically. In addition, we note peaks in washtrading activities in the summer of 2022 (which was a period of high activity in the NFT market in general) and in the end of the same year. These results are cross-validated with those in Figure 6, where we demonstrate the number and value of washtrading transactions as percentages of total

Figure 6: Washtrading Transactions (Percentages) in the NFT Market

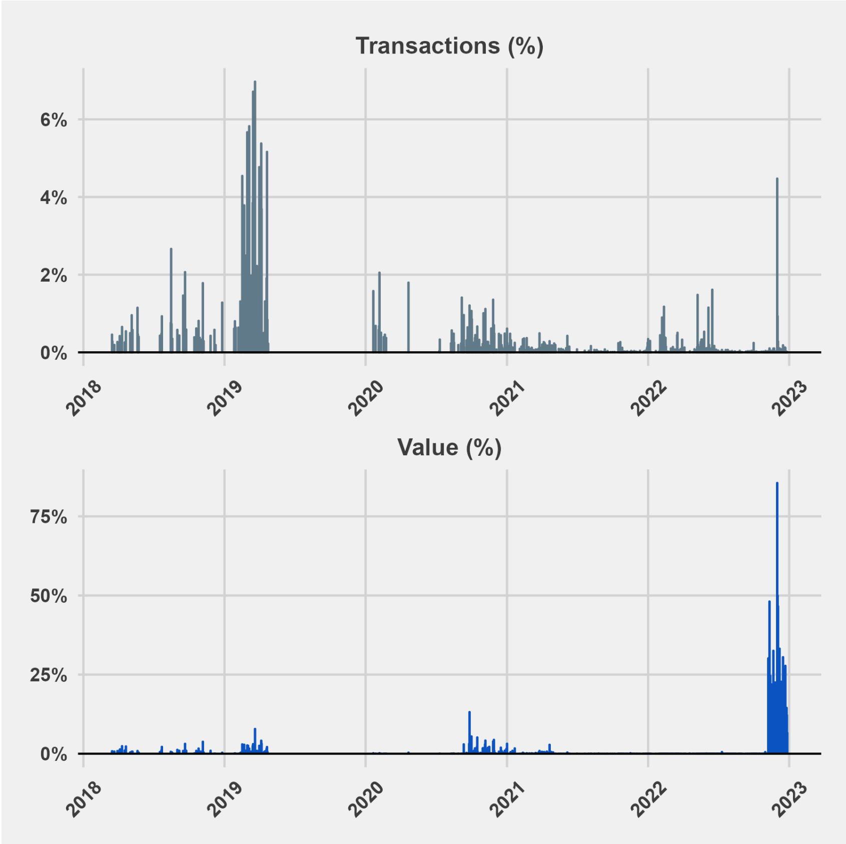

This figure shows the number and value of transactions involved in washtrading, based on our identification process. The numbers are percentages over the total number and volume of transactions in the NFT market, calculated with daily frequency.

activities. This figure yields two important findings. First, it appears that there is a strong surge in washtrading in the end of 2022, where such transactions account for as high as 75% of the total value of daily transactions, despite being only 4% of the total number of transactions. This surge does not occur on one day only, but seems to be spread out across roughly a week. In addition, we note that peaks in the percentage of washtrading transactions in early 2019, early 2020 and late 2020. Finally, the figure demonstrates the time-varying nature of washtrading activities, which are not ubiquitous in all trading days.

5.3 Characteristics of washtrading groups

Our next step is to examine the characteristics of washtrading groups and we present some summary information in Table 6. We note that washtrading groups appear as frequently as 13 times (90^{th} percentile), with the median value being at 5 but the mean being higher

Table 6: Descriptive Statistics for Washtrading Groups

Variable	N	Mean	Median	Std. Dev.	10th Perc	90th Perc
Frequency	1,030	8.4720	5.0000	14.5090	4.0000	13.1000
Number of Accounts	1,030	2.4950	2.0000	0.5740	2.0000	3.0000
Number of Transactions	1,030	47.6200	14.0000	140.6560	8.0000	84.1000
Mean Duration (hrs)	1,030	0.5580	0.1450	0.7540	0.0006	1.7953
Number of Collections	1,030	1.5550	1.0000	1.1220	1.0000	3.0000
Number of Assets	1,030	4.3870	3.0000	5.1810	1.0000	8.0000
Total Value	1,030	11,166,622.3690	2,614.0510	68,900,401.1630	1.0642	9,486,637.8958
Mean Return	1,030	40.4440	0.0000	1,179.9380	-0.0834	1.1182

This table presents descriptive statistics for the washtrading groups identified as a result of our approach.

at 8.47 times. Most clusters are made up of 2 to 3 accounts (suggesting cooperation between a tight group of NFT traders) and they can generate as many as 84 transactions. However, most groups have a small number of transactions (the median is at 14), suggesting that the washtrading process in the NFT market is not particularly sophisticated and the manipulation efforts typically involve a small number of transactions, over a limited number of assets. We note that the median (mean) total value of transactions generated by each washtrading group is USD2,614.05 (USD11,166,622.37), suggesting a fat tail towards the higher end of the distribution. Finally, in terms of the returns generated, the median mean return of each group is 0%, suggesting that the typical washtrading group may not be successful in generating returns. This result will be explored further in the next section.

Figure 7 demonstrates scatter-plots of various characteristics of washtrading groups and the (mean) return that they generate. First, we note that while most groups are able to generate positive returns for the underlying NFT assets, this is not the case for every group, suggesting that the washtrading practice is not always successful in the NFT market. We believe that this is mainly due to the non-fungible nature of the traded assets may render washtrading meaningless, if the specific asset is not sold after being manipulated[14]

The figure demonstrates that ”successful” washtrading groups (i.e. groups that were able to consistently generate high positive returns) typically appear a few number of times and do not generate a particularly high number of transactions, nor do they generate a

[14]The washtrade can still be a success as long as it raises the value of other assets in the collection; however, this cannot be observed in our approach.

Figure 7: Washtrading Group Characteristics

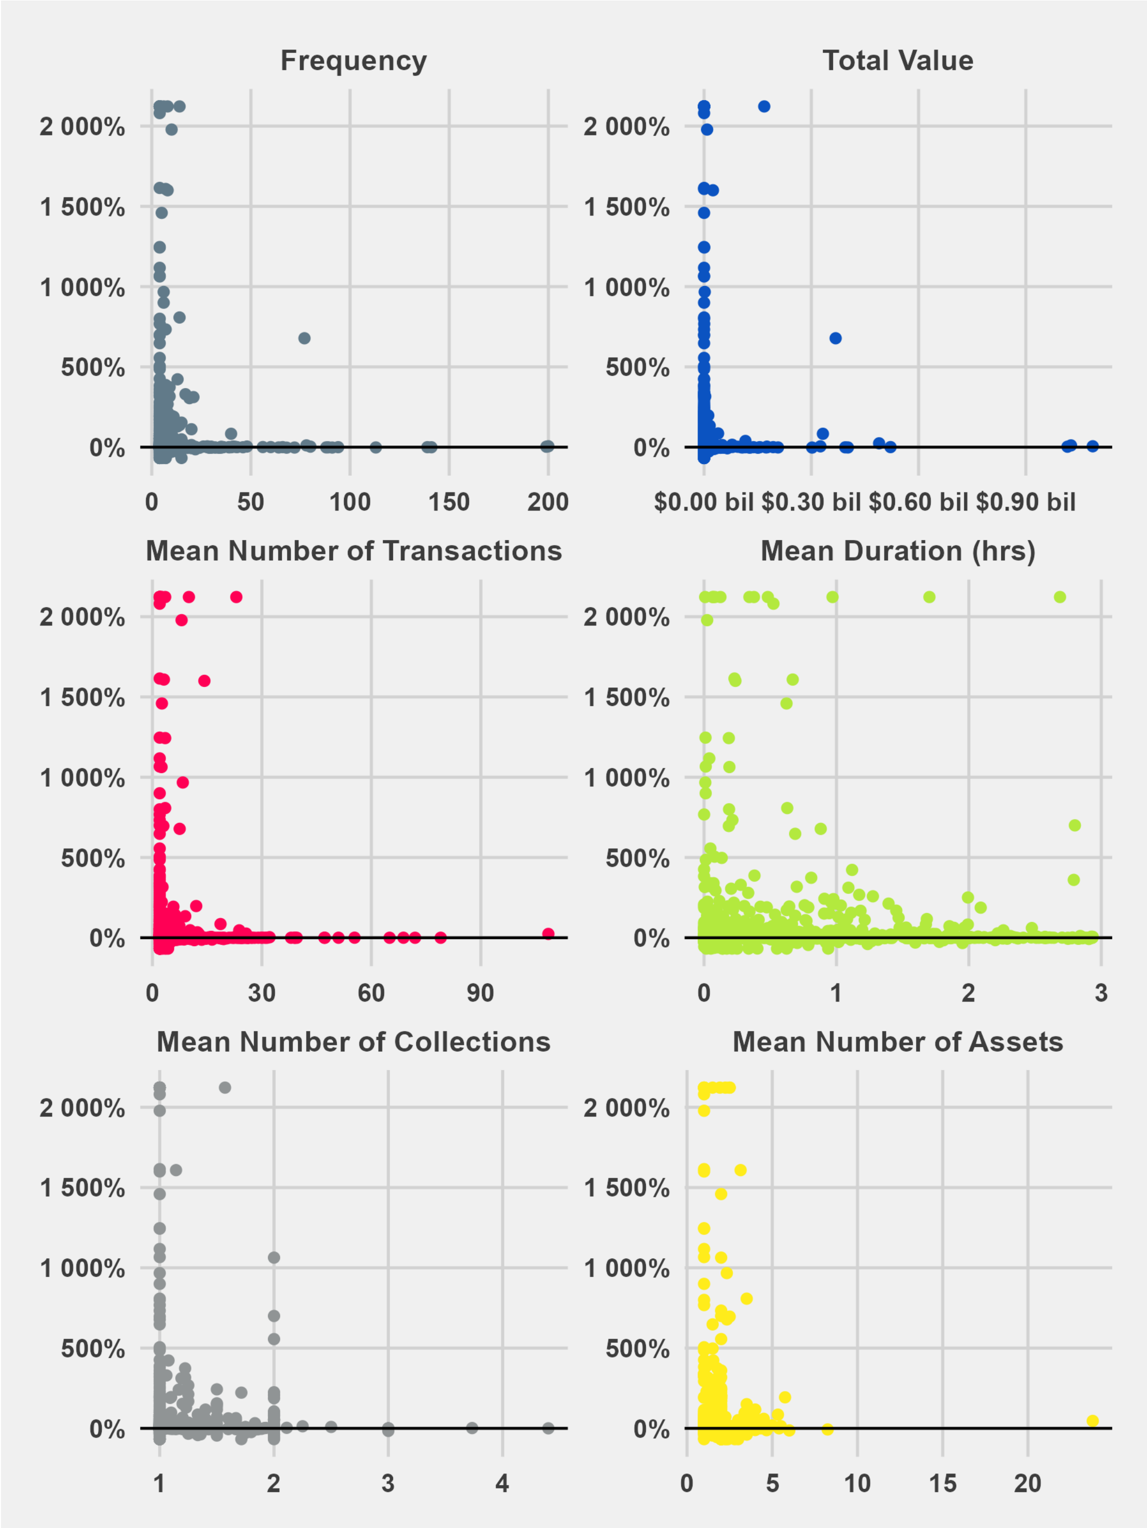

This figure includes bivariate scatterplots of different washtrading groups characteristics (in each panel) with the mean returns of the group in the vertical axis. The data in this graph is not winsorized to demonstrate the underlying distribution.

high number or volume of transactions. However, we note that groups with high returns work on few collections and few assets. Given that most collections contain numerous

NFTs but the mean number of assets in washtrading transaction is typically below five, this suggests that washtraders focus on few assets in the collection and aim and most likely aim at generating interest in the collection as a whole. Finally, in terms of duration, most groups seem to operate quickly, performing the clustered transactions in the space of around one hour.

We close this section by examining common members in different washtrading clusters. We demonstrate this discussion in Figure 8, which shows Venn diagrams of the top seven groups, according to certain metrics (e.g. returns, total value, number of transactions, etc). We examine the top seven groups for each metric due to visualisation limitations, since adding more groups would reduce readability. The numbers in the graphs represent the common elements among different groups, so that elements (accounts) that belong to a single group are in the outwards slices, while elements in the inwards slices belong to more groups. We can see that many groups share accounts (which in the crypto market correspond to wallet addresses), suggesting possible collusion between washtrading groups or, most likely, that certain accounts are part of multiple groups. We note here that it is possible for investors in crypto markets to control more than one addresses either directly or indirectly and this is further facilitated by anonymity.

5.4 Economic benefit to washtrading

Following the results of the washtrading identification strategy, we are now able to examine the returns generated by the washtrading activity or, in other words, the economic benefit of washtrading. We compute mean and total returns[15] (Figure 9) and per collection (Figure 10) and examine the distribution. We note again that the underlying data has been winsorized at the 10^{th} and 90^{th} percentile to improve readability.

Figure 9 confirms our earlier finding that washtrading groups generate low returns, since most values seem to be concentrated around zero. However, we do show that certain groups are able to increase the prices of the assets, since all distributions show long, fat tails at the higher end of the distribution. Focusing on specific collections and comput-

[15]Total returns are the sum of the returns generated in each washtrading cluster

Figure 8: Common Accounts among Washtrading Groups

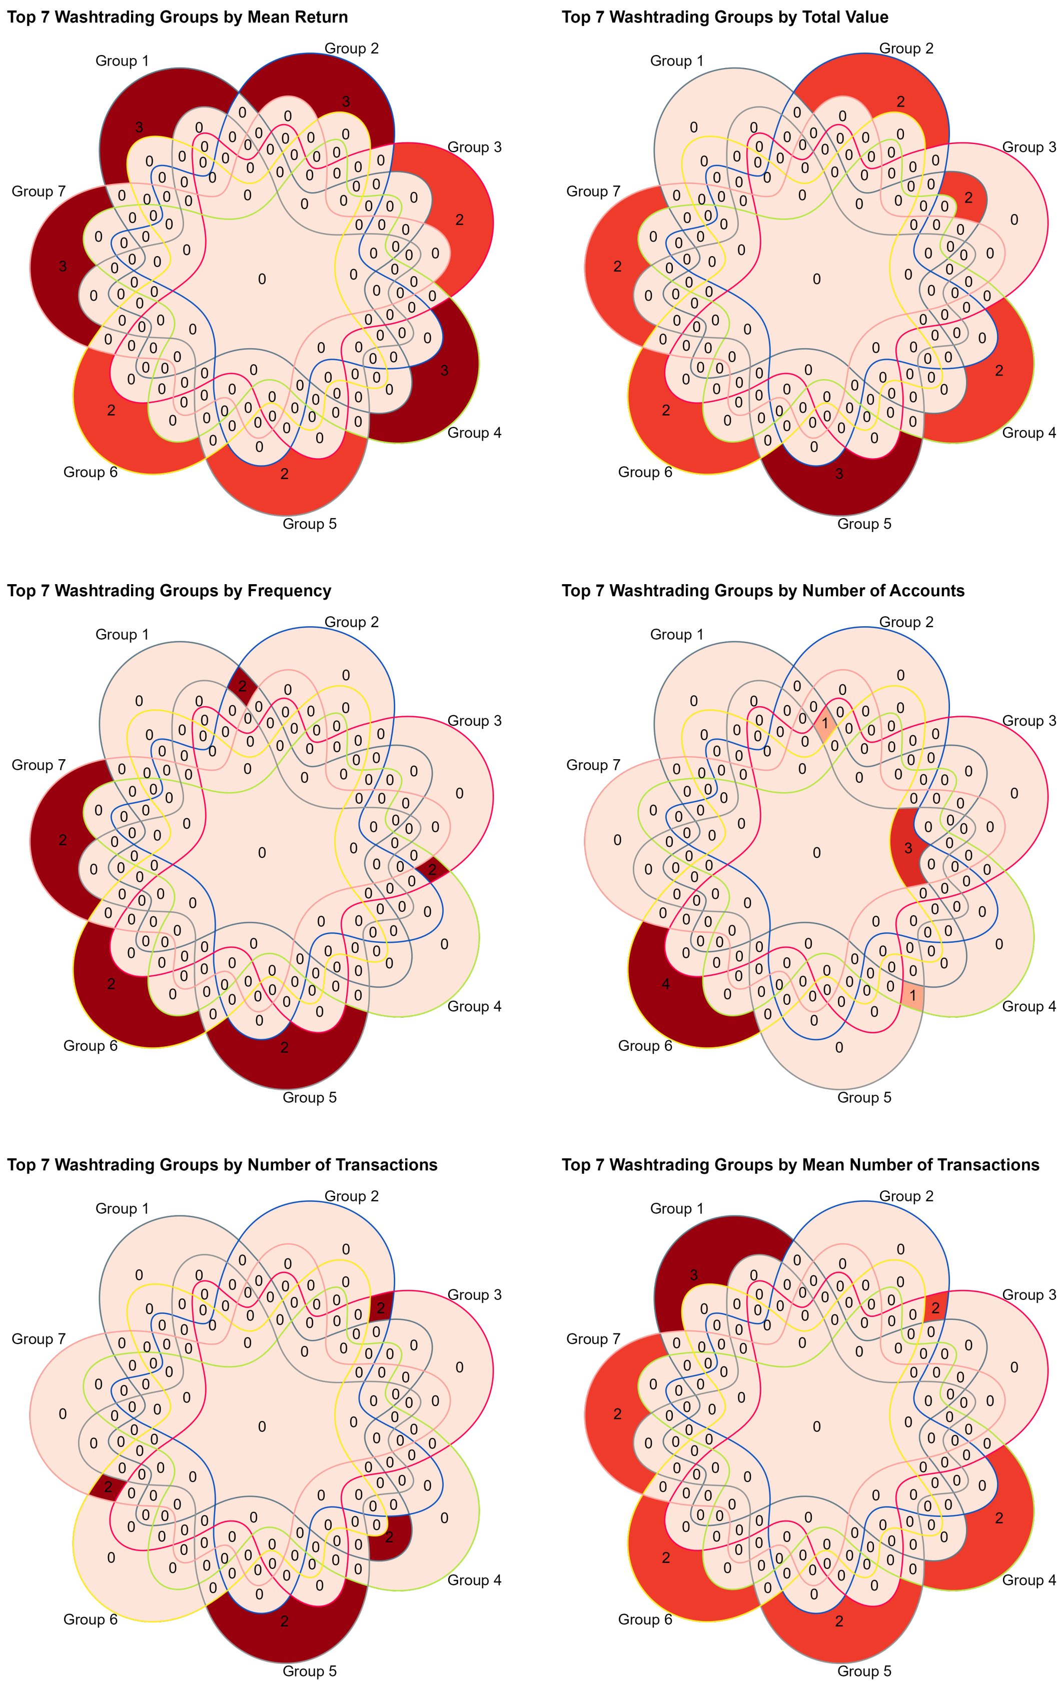

Figure 9: Distribution of Returns per Washtrading Group

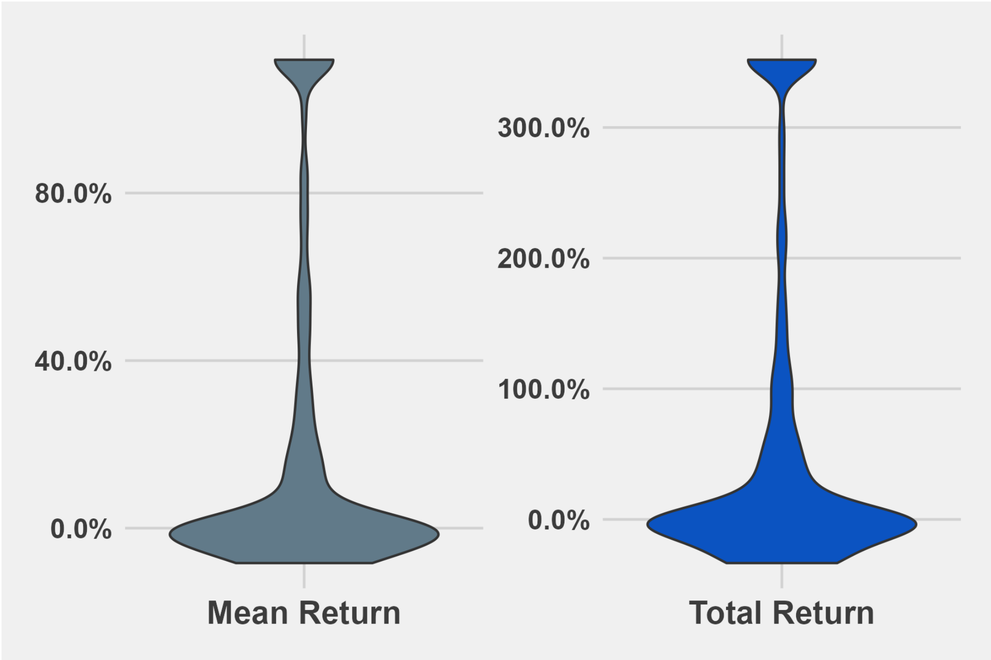

This figure demonstrates a violin plot of mean and total returns, resulting from washtrading activities, calculated for each washtrading group. The vertical axes show the values, while the width of the figure depicts the frequency density.

Figure 10: Distribution of Returns per Collection

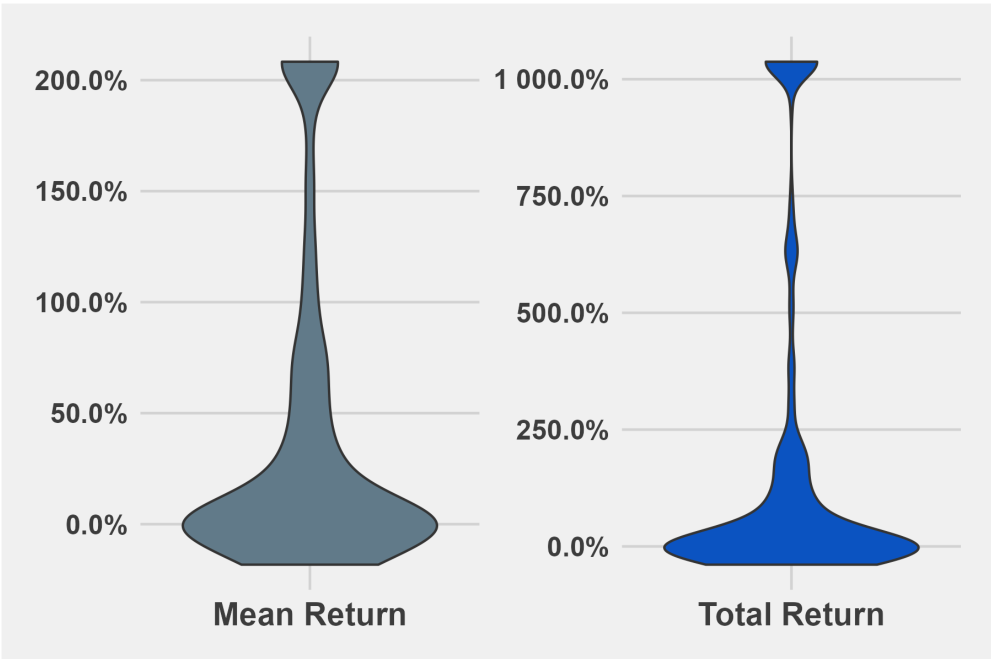

This figure demonstrates a violin plot of mean and total returns, resulting from washtrading activities, calculated for each collection. The vertical axes show the values, while the width of the figure depicts the frequency density.

ing mean and total returns per collection, we can see that the mean return for certain collections is as high as 200% with total returns reaching 1,000%. This implies washtrading has been successful in certain collections while it has failed for other collections. We will discuss possible reasons for this in the last part of this section, but we believe this asymmetry in the results reflects the non-fungible nature of the underlying assets.

5.5 Robustness Checks

Figure 11: Identification Results for Baseline Model and Robustness Checks

The figures above show the results of the identification process for the baseline model and when varying levels of small transactions are excluded. The duration of the time frame is 3 hours, while different values for the cutoff frequency are shown with different colors. The figure does not include the results for the cutoff frequency of 2 in order to facilitate interpretation.

We mentioned earlier that washtraders use high asset values, since they cannot profit from transaction volumes and have to, thus, focus on the price of the underlying assets. In order to confirm this finding, we run a series of robustness checks that exclude different levels of small transactions, varying from USD 0.1 to USD 10. If our identification process is accurate, our results should not change significantly since most transactions involved should have higher dollar values. This is confirmed, as demonstrated in Figure 11 where we see similar outcomes of the identification process for the baseline model and the robustness checks.

5.6 Unique characteristics of NFT Washtrading

Based on the above results, we can make some inferences regarding washtraders in the NFT market. We can see that washtrading in the NFT markets takes place in a relatively small time frame and that examining longer time frames can lead to spurious identification and loss of information regarding the underlying transactions. In addition, contrary to some results found in traditional markets (Cong et al., 2021), washtraders exchange NFTs at relatively high prices. This demonstrates the important differences between the markets for fungible (stock, currencies, etc.) and non-fungible assets. In the former category, washtraders can profit through volume even if the price increase of the asset is small, since inflated prices are multiplied by the quantity traded. In non-fungible asset markets, however, this is not possible as there is only one of each asset, meaning that the price increase is the only way to profit. Hence, the higher mean and median prices observed in the higher cutoff frequencies are, we believe, an indication of better identification results. In addition, washtrading groups work together consistently, engaging in washtrading clusters with the same set of accounts typically 3 to 5 times, with higher frequencies also observed, but more rarely and with few transactions.

We also demonstrated that expertise is not ubiquitous among washtrading groups. Our findings suggest that certain groups (working with specific collections) are more successful in generating high returns for the NFTs. In addition, there is an unobservable component in this process whereby washtraders increase the value of one NFT within the cluster in an effort to increase the popularity of the entire collection, thus manipulating the prices of the other assets. This outcome cannot be observed through our process and its presence is suggested by the fact that washtrading is limited to few assets in each collection (typically no more than five), despite the fact that most collections have numerous NFTs.

6 Conclusion

This study investigates potential NFT market manipulation. Specifically, we examine the incidence and extent of washtrading - a market manipulation strategy whereby investors trade assets among themselves with the aim of misleading other market participants and realize a monetary gain. This paper is the first to systematically study NFT washtrading based on large and comprehensive NFT dataset.

First of all, we find that washtrading activity follows different patterns over time, across NFT types, as well as with respect to other dimensions, e.g., NFT market valuation. Second, we propose a new methodology to determine the incidence and magnitude of washtrading practices. We document that whashtrading volume is maximized within a very short trading window, abnormally short for typical NFT transactions that are relatively infrequent. We also find that, on average, washtraders prefer NFTs with higher price levels, several orders of magnitude higher than the ones observed for normal NFT trading. Further, we show that the manipulated NFTs deliver significantly higher returns as compared to the full sample average. Overall, our findings suggest that washtrading elevates NFT trading volume and NFT market valuations, a result in line with washtrading outcomes documented for the traditional financial markets. However, NFT washtrading appears to be much less rampant than what has been portrayed in business media outlets.

A natural extension of our analysis would be to consider other types of market manipulation strategies that might be utilized in the NFT market besides washtrading. This presents a potentially fruitful avenue for future research.

References

Cao, Y., Li, Y., Coleman, S., Belatreche, A., & McGinnity, T. M. (2015). Detecting wash trade in financial market using digraphs and dynamic programming. *IEEE transactions on neural networks and learning systems*, *27*(11), 2351–2363.

Cong, L. W., Li, X., Tang, K., & Yang, Y. (2021). Crypto wash trading. *arXiv preprint arXiv:2108.10984*.

Cumming, D., Johan, S., & Li, D. (2011). Exchange trading rules and stock market liquidity. *Journal of Financial Economics*, *99*(3), 651–671.

Darolles, S., Le Fol, G., & Mero, G. (2015). Measuring the liquidity part of volume. *Journal of Banking & Finance*, *50*, 92–105.

Gerhold, S., Guasoni, P., Muhle-Karbe, J., & Schachermayer, W. (2014). Transaction costs, trading volume, and the liquidity premium. *Finance and Stochastics*, *18*, 1–37.

Imisiker, S., & Tas, B. K. O. (2018). Wash trades as a stock market manipulation tool. *Journal of behavioral and experimental finance*, *20*, 92–98.

Le Pennec, G., Fiedler, I., & Ante, L. (2021). Wash trading at cryptocurrency exchanges. *Finance Research Letters*, *43*, 101982.

Mahoney, P. G. (1999). The stock pools and the Securities Exchange Act. *Journal of Financial Economics*, *51*(3), 343–369.

Pouncy, C. R. (1994). The scienter requirement and wash trading in commodity futures: The knowledge lost in knowing. *Cardozo L. Rev.*, *16*, 1625.

Appendix A Identification analysis

The figures below present the results of the identification process based on the two key parameters: (1) the time elapsed between the first and the last NFT sale in a sequence of NFT buys and sells truncated arbitrarily at 3, 6, 9, 12, 24, or 30 hour time window (illustrated by the histograms) and (2) the frequency of incidence of the same accounts in washtrading clusters (denoted by F), where cluster is defined as a sequence of NFTs buys and sells. NFT data are extracted through the OpenSea API.

Panel A1: Number of Unique Accounts for Different Combinations of Model Parameters

Panel A2: Number of Unique Collections for Different Combinations of Model Parameters

Panel A3: Number of Unique Assets for Different Combinations of Model Parameters

Panel A4: Mean Transaction Price for Different Combinations of Model Parameters

Panel A5: Median Transaction Price for Different Combinations of Model Parameters

Panel A6: Total Value of Washtrading Transactions for Different Combinations of Model Parameters

Figure A1: Number of Unique Accounts for Different Combinations of Model Parameters

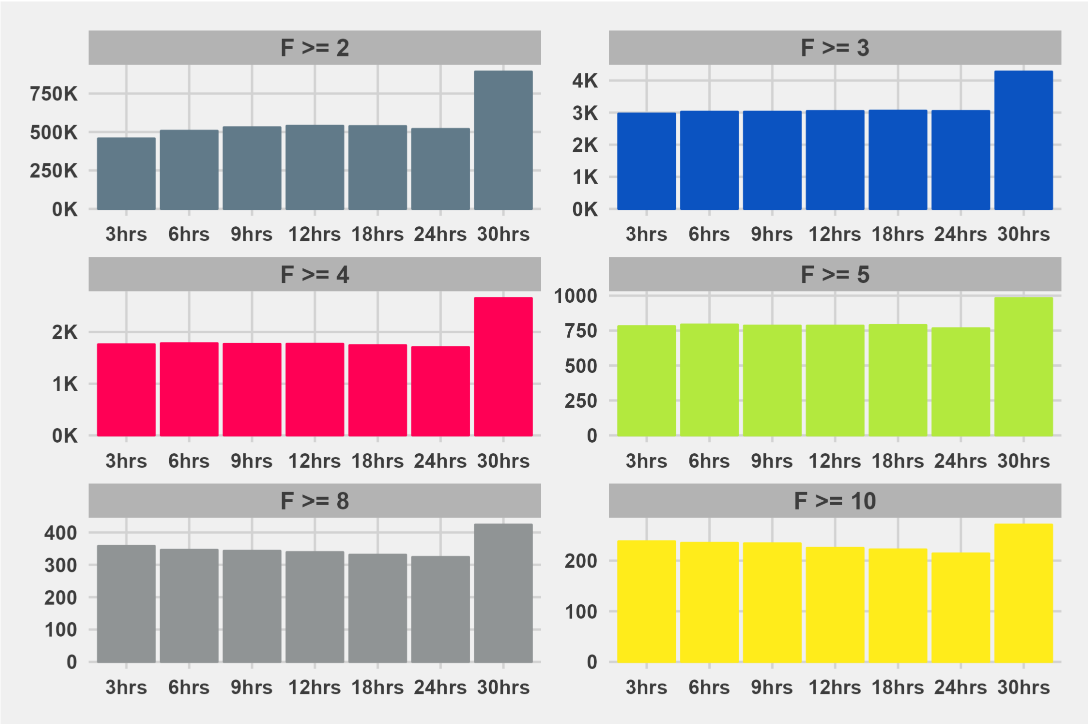

Figure A2: Number of Unique Collections for Different Combinations of Model Parameters

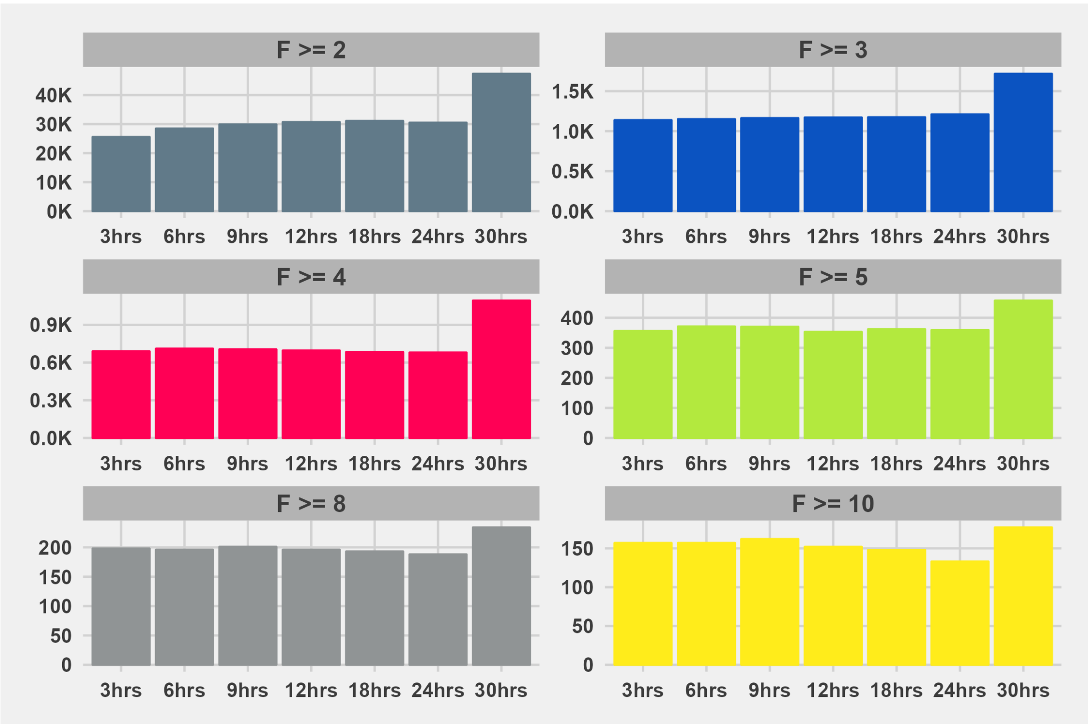

Figure A3: Number of Unique Assets for Different Combinations of Model Parameters

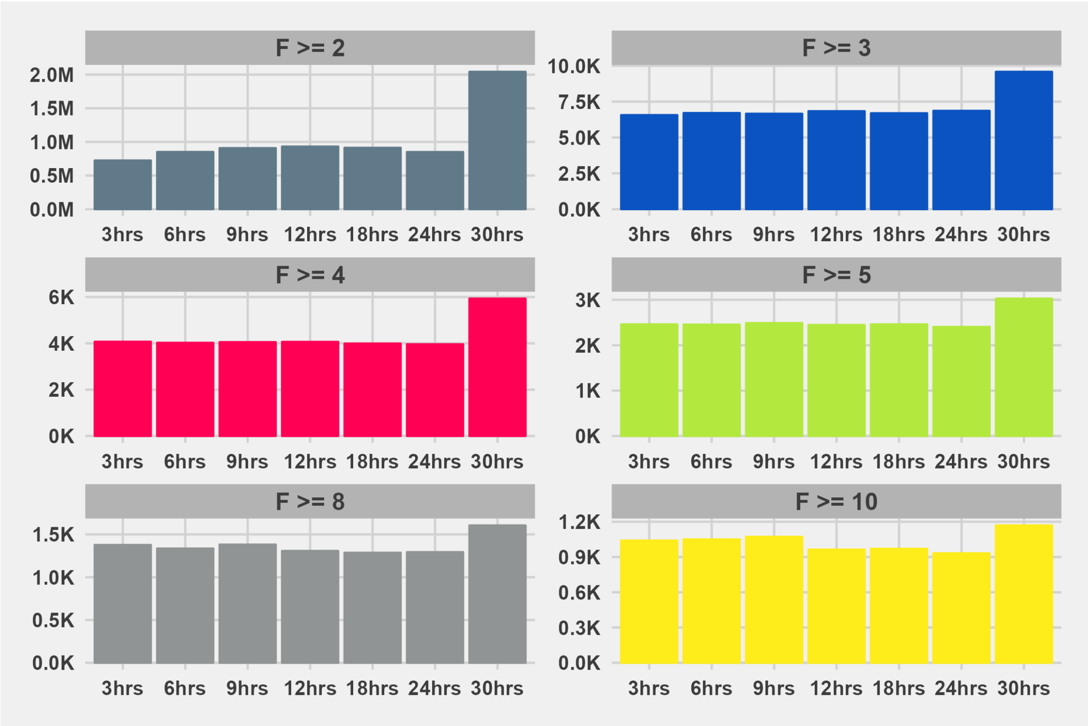

Figure A4: Mean Transaction Price for Different Combinations of Model Parameters

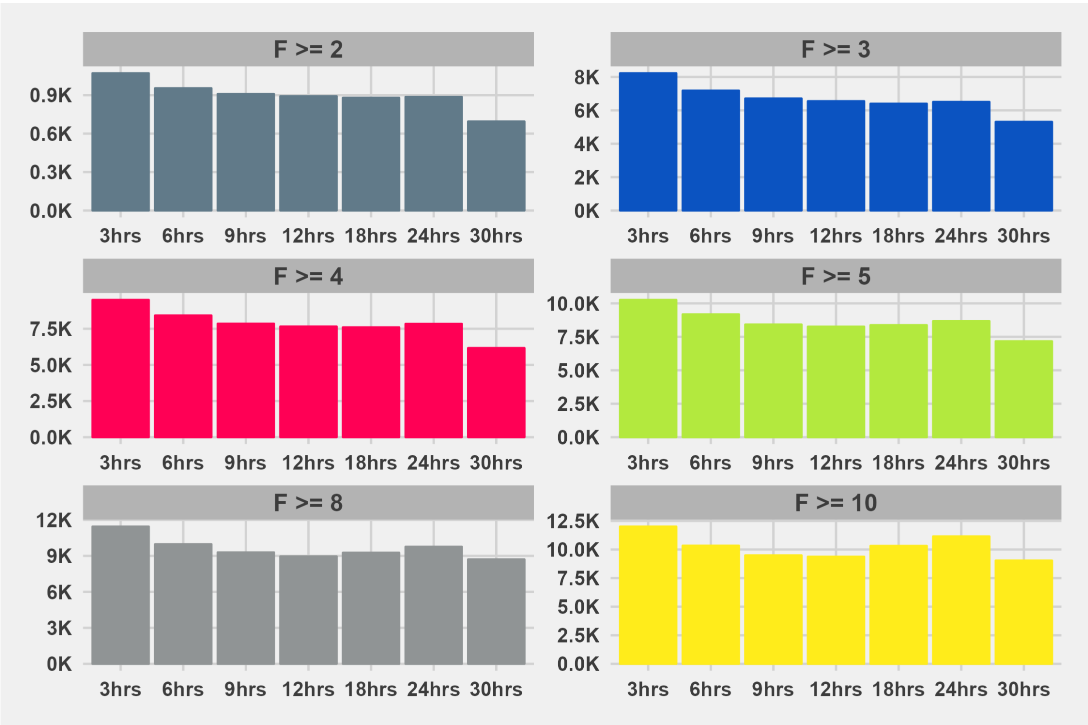

Figure A5: Median Transaction Price for Different Combinations of Model Parameters

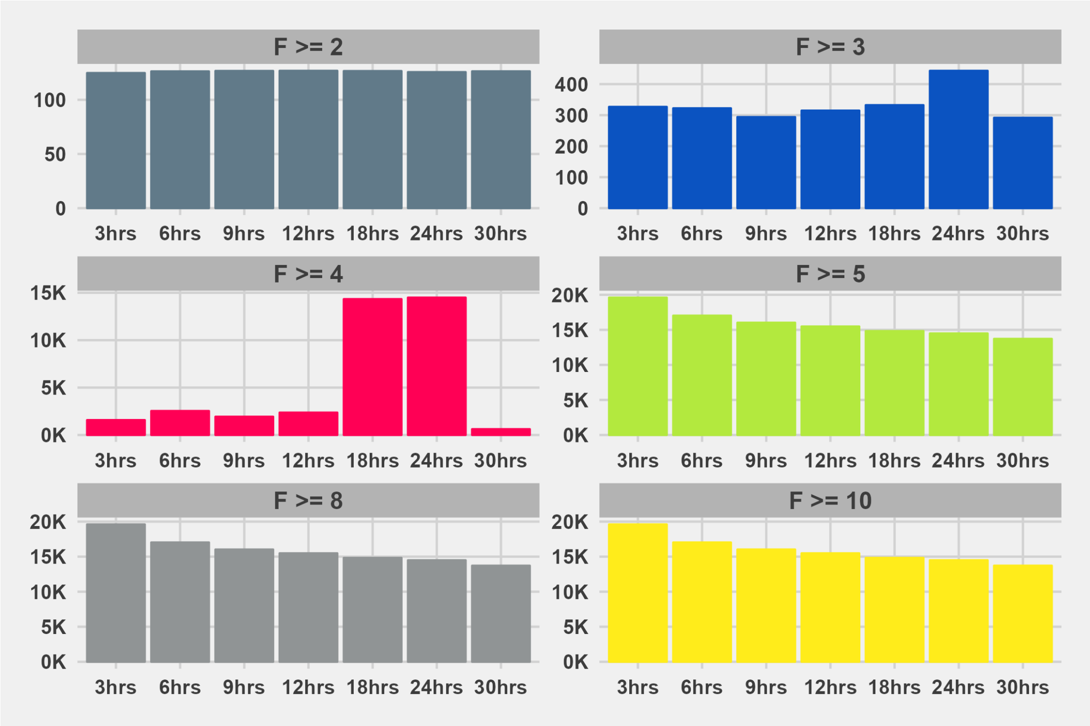

Figure A6: Total Value of Washtrading Transactions for Different Combinations of Model Parameters