

NFT Bubbles

Andrea Barbon and Angelo Rinaldo*

March 2023

Abstract

By investigating nonfungible tokens (NFTs), we provide the first systematic study of retail investor behavior through asset bubbles. Given that NFTs are recorded in public blockchains, we are able to track investor behavior over time, leading to the identification of numerous price run-ups and crashes. Our study reveals that agent-level variables, such as investor sophistication, heterogeneity, and wash trading, in addition to aggregate variables, such as volatility, price acceleration, and turnover, significantly predict bubble formation and price crashes. We find that sophisticated investors consistently outperform others and exhibit characteristics consistent with superior information and skills, supporting the narrative surrounding asset pricing bubbles.

Keywords: Financial Bubbles, Nonfungible Tokens, Agent-level, Blockchain.

JEL classification: G14; G12

*Andrea Barbon (andrea.barbon@unisg.ch) and Angelo Rinaldo (angelo.rinaldo@unisg.ch) are both at the University of St. Gallen and Swiss Finance Institute. Andrea Barbon is the corresponding author. We thank Kirill Kazakov for the excellent research assistance. We acknowledge financial support from the Swiss National Science Foundation (SNSF grant 204721).

“Speculators may do no harm as bubbles on a steady stream of enterprise. But the position is serious when enterprise becomes the bubble on a whirlpool of speculation.”
([Keynes \(1936\)](#), p. 101)

1 Introduction

Ever since Keynes, the narrative on speculative bubbles has contemplated two types of investors: pure speculators, who are mainly interested in the resale price of financial securities, and more sophisticated investors, who assess the market price compared to the fundamental value. The prevalence of speculators over sophisticated investors is the breeding ground for asset bubbles ([Kindleberger, 1978](#)). Although intuitive, the literature has not yet shed full light on this idea, and it has not yet been empirically ascertained which types of investors contribute to the formation and bursting of asset bubbles. One possible reason for this is the difficulty of tracing investor behavior through speculative bubbles.

This paper provides the first systematic study of retail investor behavior through asset bubbles. To do this, we exploit the distinctive features of nonfungible tokens (NFTs), which are unique digital assets providing proof of ownership and verification of authenticity held in the blockchain. There are at least three good reasons for leveraging the NFT market for this research purpose. First, transaction attributes and traders’ identities are highly transparent since every transaction is recorded on a public blockchain. Second, even though they were introduced only a few years ago,¹ the volume traded in NFTs as of February 2023 is roughly \$73 billion.² In this market populated only by retail investors, a legitimate question that has been asked by many is whether it has been a “balloon” blown up by a technological innovation ([Van Horne, 1985](#)) or a real speculative bubble. Third, it is more difficult even for a sophisticated investor to deter-

¹The first NFT (Quantum) was created by Kevin McCoy and Anil Dash in May 2014 ([Creighton, 2022](#)).

²[Barbon and Ranaldo \(2023\)](#) explain the crucial aspects of the NFT and its market as well as quantify its size by considering the largest NFT exchanges, which are OpenSea, SuperRare, LooksRare, Foundation, and Rarible.

mine the fundamental value of these assets and arbitrage their deviations from market prices given their uniqueness (or non-fungibility), the absence of cash flow streams, and the inability to sell them short. Short selling NFTs in strictly nonfungible collections poses significant challenges due to the unique nature of each token. Covering a short position is difficult as the short seller must acquire the exact same token, which may result in a short squeeze. While short selling may be feasible for semifungible collections with large supplies, this analysis focuses on strictly non-fungible collections for consistency. All these aspects offer a unique opportunity to study how the behavior of retail investors relates to the formation and burst of asset bubbles.

The key question that motivates our research is whether the predictive power of asset bubbles is more effectively discerned by scrutinizing the conduct of individual investors. In line with the asset bubble narrative, we posit that if there exists a larger cohort of sophisticated investors with a heightened degree of skill and expertise, their augmented presence should predict fewer incidents of speculative bubble bursts. Three main results stand out from our analysis: First, aggregate variables such as high volatility, price acceleration, and low turnover significantly predict crashes. Second, some investors persistently perform better than others, and their characteristics suggest that they are more sophisticated, with superior information and skills. Third, a greater presence of sophisticated agents in the price run-up significantly decreases the risk of a subsequent crash and increases ex-post positive returns. In contrast, a smaller number of individual investors and more price manipulation through “wash trading” increase the crash risk and predict more illiquidity in the crash phase, thus exacerbating the adverse effects.

We proceed in three steps. First, we build a representative data set of the entire NFT market, including the 1,000 most traded NFT collections³ on the largest NFT marketplace, [OpenSea](#). Over the sample period from January 2021 to September 2022, this data set includes more than 15 million transactions from 1.3 million unique wallets,

³An NFT collection is a set of NFT tokens based on the same smart contract. Usually, they are produced by the same artist and share a common theme. For instance, the [Bored Ape Yacht Club](#) or the [CryptoPunks](#) are examples of collections.

generating a total trading volume of more than 18 billion USD (8.5 million ETH). To identify run-up and crash phases, we adapt the methodology of [Greenwood et al. \(2019\)](#) to determine high-frequency price movements that can be considered short-term bubbles in the same vein as the dot-com bubble literature (e.g., [Ofek and Richardson, 2003](#); [Brunnermeier and Nagel, 2004](#)). We identify around 1,000 price run-up events, half of which turn into ex-post crashes. The first question we address is whether crashes can be predicted using aggregate variables. We find strong empirical evidence that collection-specific variables, rather than market-wide variables, predict price crashes. More precisely, the higher the volatility and acceleration of the price, the more likely it is to crash. On the contrary, high turnover is associated with a lower probability of crashing, suggesting that, due to an inability to sell short, a lower concentration of optimistic investors ([Scheinkman and Xiong, 2003](#)) and smaller liquidity risk predict smoother ex-post price adjustments ([Nneji, 2015](#)). The picture that shines through this first analysis is that NFT bubbles are predictable, and they are not very dissimilar to those in equities since they share similar asset-specific predictors of price crashes.

As a second step, we turn to the next question that guides our research: Are there more sophisticated retail investors? If so, what role do they play in the formation and bursting of asset bubbles? A priori, it is not obvious that there may be sophisticated investors in this market. On the one hand, this market is populated only by retail investors rather than institutional investors or hedge funds that presumably have more advanced skills ([Brunnermeier and Nagel, 2004](#)). On the other hand, the fundamental value of these new digital assets is more challenging to determine than in other financial securities given the absence of expected cash flow streams originated, for example, by dividends or coupon payments, and that their uniqueness (or nonfungibility) makes it arduous to find close substitutes for hedging or arbitrage purposes. In addition, an investor who believes the price is excessive relative to the fundamental value cannot sell short in this market. Against this background, we ascertain whether there is heterogeneity among investors, whether that heterogeneity is persistent, and whether it

can explain investor performance.

Two main findings emerge: First, the agent’s profitability is persistent across run-up events, suggesting that some agents possess more advanced abilities to perform consistently better over time. Second, outperforming agents are distinguished in terms of (a) greater market presence, (b) financial leverage, and (c) sharper timing. In addition to trading more volume and frequently, their activity spans multiple exchanges (cross-market trading), they are more likely to provide liquidity and to trade in decentralized exchanges (DEX), like Uniswap, Balancer, or Curve, and they leverage up their positions to exploit profit opportunities and access arbitrage capital through decentralized lending platforms. Furthermore, these investors are more skilled in timing the bubbles, by both buying early into the run-up phase and selling close to the peak of the bubble. All in all, the more sophisticated behavior of outperforming investors suggests that some individuals possess superior information and skills in identifying NFT investment opportunities, timing the market movements, and avoiding being caught in the bursting of NFT bubbles.

Our study brings forth a third innovation by demonstrating that, in addition to (collection-specific) aggregate variables, agent-level variables can assist in predicting asset price crashes and ex-post returns. Specifically, we include three metrics that are measured ex-ante as additional predictive variables for crashes and find that the prediction accuracy improves. First, we provide compelling evidence that a greater presence of sophisticated investors is negatively related to ex-post crashes and positively related to ex-post returns, consistent with the narrative about asset price bubbles and the idea that sophisticated investors have superior skills in choosing profitable investments less prone to overvaluation and doing so with good timing. Rather than a causal interpretation, these results may simply highlight the ability of sophisticated agents to estimate ex-ante the probability of ex-post crashes and, thus, self-select into noncrash events. Second, we calculate a measure of ownership concentration based on the number of individual investors owning an NFT of the collection relative to the available supply.

We find that an increase in individual owners predicts lower crash risk and greater liquidity in the ex-post price correction phase. These results support our hypothesis that wider heterogeneity reduces the risk of concentration in optimistic investors and mitigates the adverse effects of the short sale impossibility. It also reduces liquidity risk in future transactions as a larger and more diverse community of traders facilitates resale opportunities, reduces search costs, and increases risk sharing.

Third, we analyze wash trading⁴ at the aggregate and transaction levels, identifying suspicious trades empirically. Although this practice of price manipulation is not widespread in OpenSea and in our data set, we find that when wash trading is more pervasive ex-ante, ex-post crashes and liquidity deterioration are more likely.

Our paper contributes to two strands of the literature. First, we add to an extensive literature on asset bubbles.⁵ Some papers suggest that bubbles could be driven by asymmetric information (e.g., [Allen et al., 1993](#)), limit to arbitrage (e.g., [Shleifer and Vishny, 1997](#)), heterogeneous beliefs and individual characteristics (e.g., [Abreu and Brunnermeier, 2003](#); [Hong and Stein, 1999](#); [Scheinkman and Xiong, 2003](#)), and risk premia combined with learning about new technologies (e.g., [Pastor and Veronesi, 2009](#)). We extend the previous literature by empirically analyzing the run-up and crash of asset bubbles on an *agent* basis. The NFT market provides us with an ideal laboratory for analyzing a large number of episodes subject to the impossibility of short selling and accurately identifying individual retail investors. In contrast to prior research, we do not need to conjecture the investor’s sophistication by contrasting institutional and retail investors ([Ofek and Richardson, 2003](#)) or imposing a priori that sophistication pertains only to hedge fund managers ([Brunnermeier and Nagel, 2004](#)). In doing so, we highlight new characteristics of individual investors who outperform in bubble episodes such as advanced skills in timing, leveraging, and cross-market activities. Our results suggest that (i) rather than pure speculators, outperforming investors are more

⁴We broadly define wash trading following the Commodity Futures Trading Commission (CFTC) concept: “transactions to give the appearance that purchases and sales have been made, without incurring market risk or changing the trader’s market position” (CFTC, 2022).

⁵See [Brunnermeier \(2016\)](#) for a recent survey of this literature.

sophisticated in terms of information and skills, and (ii) a larger share of sophisticated investors is negatively (positively) related to bubble bursts (ex-post returns), a result so far only documented with experimental approaches ([Bosch-Rosa et al., 2018](#)).

Second, we add to the nascent literature on blockchain-based assets. So far, only a few studies on NFT have been carried out. There are three notable exceptions related to our work: First, [Borri et al. \(2022\)](#), who construct an NFT market index and conduct an asset pricing analysis. Second, [Ito et al. \(2022\)](#) conduct a time-series analysis of NFT prices that can turn into bubbles using logarithmic periodic power law (LPPL) models. Third, [Oh et al. \(2022\)](#) analyze whether investors' experience predicts higher NFT returns. Other papers focus on specific NFTs such as (digital) real estate in Decentraland ([Goldberg et al., 2021](#); [Dowling, 2022a](#)), CryptoKitties ([Kireyev and Lin, 2021](#)), and CryptoPunks ([Kong and Lin, 2021](#)). [Nadini et al. \(2022\)](#) focus on the investor's propensity to buy similar NFTs. Other papers deal with the empirical features of NFT prices. For example, [Dowling \(2022b\)](#) analyzes the relation between NFT and cryptocurrency prices, [Yousaf and Yarovaya \(2022\)](#) examine return-volume and return-volatility connectedness in NFTs, and [Ante \(2022\)](#) examines spillover effects among NFT sales and between NFT and cryptocurrency prices using error correction models. We add to the NFT literature by investigating the individual characteristics of NFT investors and how these predict aggregate NFT price movements. Since those who invest in NFTs are typically retail investors rather than traditional financial institutions, our work can also be considered a contribution to the FinTech and household finance literature ([Campbell, 2006](#)).

The paper is organized as follows. Section 2 presents the data. Section 3 identify price run-ups and crashes. Section 4 analyzes individual investors. Section 5 concludes.

2 Data and Variables

We construct a representative data set for the total NFT market by downloading transaction-level data for the 1,000 most traded NFT collections on the largest NFT marketplace, [OpenSea](#).⁶ The choice of collections is based on the total volume traded (denominated in ETH) over the entire lifespan of each collection. We limit the sample to collections based on smart contracts complying with the ERC-721 standard, thus excluding from the sample semifungible collections.⁷ This provides us with a homogeneous universe of collections in which single NFTs are practically impossible to sell short. The sample period ranges from January 2021 to September 2022, covering 621 trading days. The data include roughly 15 million transactions from 1.3 million unique wallets, generating a total trading volume of about 20 billion USD (9 million ETH). The average daily volume is 11.5 million USD. According to data from Dune Analytics, our sample covers around 5% of the total dollar volume of NFTs across all exchanges and all blockchains over the same historical period. We use average transaction prices to construct a panel at the hourly frequency consisting of 3,535,662 observations. The panel includes the following list of variables measured at the collection hour level:

Price: Average of sale prices, expressed in ETH.

Floor: Lowest price at which an NFT of the collection is listed for sale.

Return: Percentage change in price.

Volume: Trading volume, expressed in ETH.

Sales: Count of the number of NFTs sold.

Minted: Number of NFTs minted from the contract, purchased or airdropped.

Supply: Number of NFTs available in the collection (cumulative sum of mints).

Turnover: Ratio of sales divided by supply, expressed in percentage terms.

⁶OpenSea is a decentralized marketplace for buying, selling, and creating NFTs on the Ethereum blockchain. It provides a platform for creators and users to trade and showcase their virtual assets, such as digital art, gaming items, and more, in a secure and transparent manner.

⁷The [Ethereum Request for Comments \(ERC\) 721](#) is a standard for smart contracts operating nonfungible tokens. There exist more general standards, for instance, the [ERC-1155](#), operating both non-fungible and semifungible tokens.

Table 1. Summary Statistics. The table provides summary statistics for the variables contained in our panel at the collection-hour level. An NFT collection is a set of NFT tokens based on the same smart contract, usually produced by the same artist, and sharing a common theme. The variables include: *Price*, the average of sale prices, expressed in ETH; *Floor*, the lowest price at which an NFT of the collection is listed for sale; *Return*, the percentage change in price; *Volume*, the trading volume expressed in ETH; *Sales*, the count of the number of NFTs sold; *Minted*, the number of NFTs minted from the contract, purchased, or airdropped; *Supply*, the number of NFTs available in the collection (cumulative sum of mints); *Turnover*, the ratio of sales divided by supply, expressed as a percentage; *Market Capitalization*, the product of supply and floor price; *Age*, the hours elapsed from the creation of the collection. All variables are winsorized at the 1% level.

	Count	Mean	Std	10%	50%	90%
<i>Price</i> (Ξ)	3,511,676	0.92	2.19	0.04	0.22	2.01
<i>Floor</i> (Ξ)	3,534,040	0.80	2.05	0.03	0.18	1.73
<i>Return</i> (%)	3,508,406	5.63	44.04	-16.46	0.00	19.62
<i>Volume</i> (Ξ)	3,535,662	1.24	4.42	0.00	0.00	2.38
<i>Sales</i>	3,535,662	1.66	4.42	0.00	0.00	4.00
<i>Minted</i>	3,535,662	8,102.61	7,998.05	623.00	7,795.00	11,312.00
<i>Supply</i>	3,440,668	6,467.23	9,027.44	341.00	3,974.00	11,111.00
<i>Turnover</i> (%)	3,439,122	0.33	1.34	0.00	0.02	0.46
<i>Market Cap</i> (Ξ)	3,534,040	5,254.65	15,264.94	70.83	1,100.31	9,745.52
<i>Age</i> (hours)	3,535,662	2,739.42	2,095.13	443.00	2,307.00	5,648.00

Market Capitalization: Product of supply and floor price.

Age: Hours elapsed from the creation of the collection.

The findings of our study regarding the average transaction price and total trading volume of NFT collections are presented in Figure 1. To provide a detailed overview of the sample data, we report the summary statistics of the panel data, which have been winsorized at the 1% level, in Table 1. As illustrated in the table, the average price of an NFT in our sample is approximately 0.92 ETH, equivalent to roughly 2,000 USD. Additionally, the average market capitalization is 10 million USD. These relatively high valuations are consistent with our focus on collections that have generated significant trading volume. Additionally, the hourly percentage price changes, on average, are higher than 5%, but with a wide dispersion, reaching 44% volatility. The mean turnover ratio is 0.33% at the hourly frequency, implying that, on average, 8% of the outstanding NFTs of a collection are traded on any given day. Finally, the average time elapsed since the creation of a collection is 114 days, which indicates that most collections are present during a large portion of our sample period.

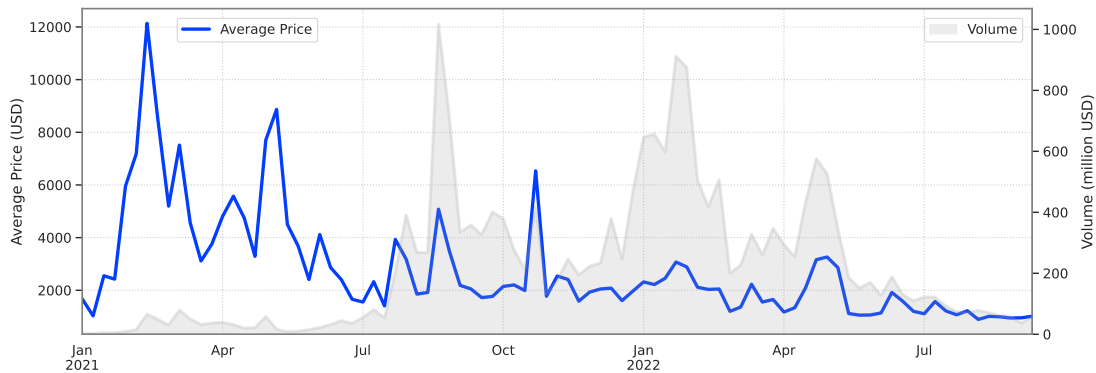


Figure 1. NFT market. The figure presents the weekly time series of average transaction prices (in USD) and total trading volume (in million USD) based on our data set, comprising the 1,000 most traded NFT collections on OpenSea from January 2021 to September 2022. This sample covers around 5% of the total dollar volume on NFTs across all exchanges and all blockchains over the same period.

3 Price Run-Ups and Bubbles

We now examine the evolution of prices for our sample of NFT collections to identify price run-ups and subsequent crashes. We then run cross-sectional regressions on the identified events to assess the ex-ante predictability of crashes.

3.1 Price Run-up Events

The approach used to identify price run-up events in our study is based on the methodology proposed by [Greenwood et al. \(2019\)](#). However, we have adapted their approach to suit our intraday data that cover the NFT market and to neatly link individual investors' behavior to high-frequency abnormal price movements that can be considered short-term bubbles, similarly to the dot-com bubble literature (e.g., [Ofek and Richardson, 2003](#); [Brunnermeier and Nagel, 2004](#)). Specifically, we have made two key modifications. First, we have substituted monthly returns with high-frequency (hourly) returns. Second, we have utilized a shorter window of 24 hours to detect run-ups. This ensures that the number of observations utilized to identify run-up events remains unaltered. More specifically, a given collection is defined to experience a *price run-up* at

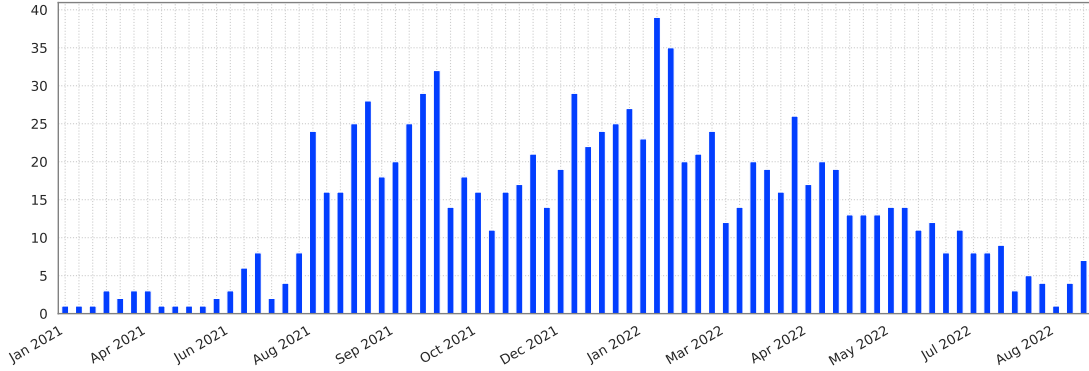


Figure 2. Events distribution. Distribution of price run-up events over time, aggregated at a weekly frequency. Run-up events are defined as situations in which the average sell price of an NFT collection increases by more than 100% within 24 hours.

hour h if the cumulative return over the preceding 24 hours is greater than or equal to 100%. We consider a window of 48 hours centered around h for each run-up event, so that the event time $t = 0$ corresponds to the hour in which the run-up is identified. We refer to negative event times as *ex-ante* and to positive event times as *ex-post*. We impose the condition that two run-up events for the same collection have nonoverlapping event windows, and further, we require the trading volume during the event window to sum up to at least 10 ETH. The procedure identifies 1,017 price run-up events, distributed over the entire sample period but more frequent in the interval between August 2021 and May 2022, as Figure 2 shows. Figure 3 presents the distribution of event level trading volume (in USD) and the total number of active wallets per event on a logarithmic scale. The average volume per event is 1.5 million USD, with 1,300 wallets actively participating, on average. Figure 4 exhibits the characteristic event time pattern. On average, the event time cumulative returns are slightly positive in the *ex-post* period. At the same time, the median is significantly negative, implying that the distribution of *ex-post* returns is highly skewed. Interestingly, immediately after the run-up is identified, at $t = 1$ the median is positive at around 25%, while it declines quickly and reaches -50% in the following four hours.

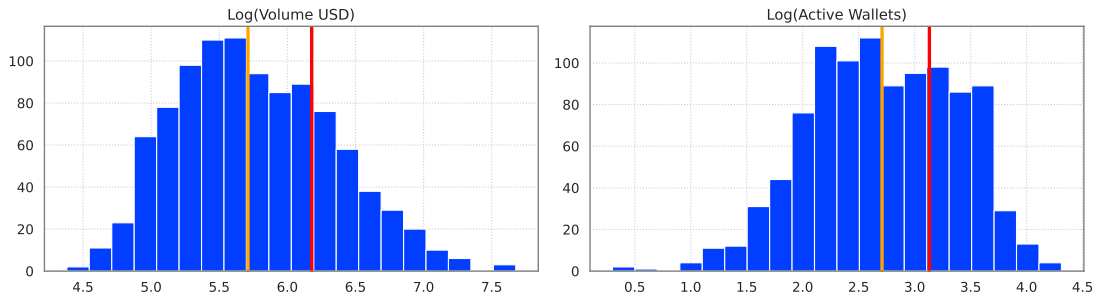


Figure 3. Events participation. The figure shows the trading volume (in USD) and the number of active wallets per event on a logarithmic scale. The red line indicates the average, while the orange line indicates the median.

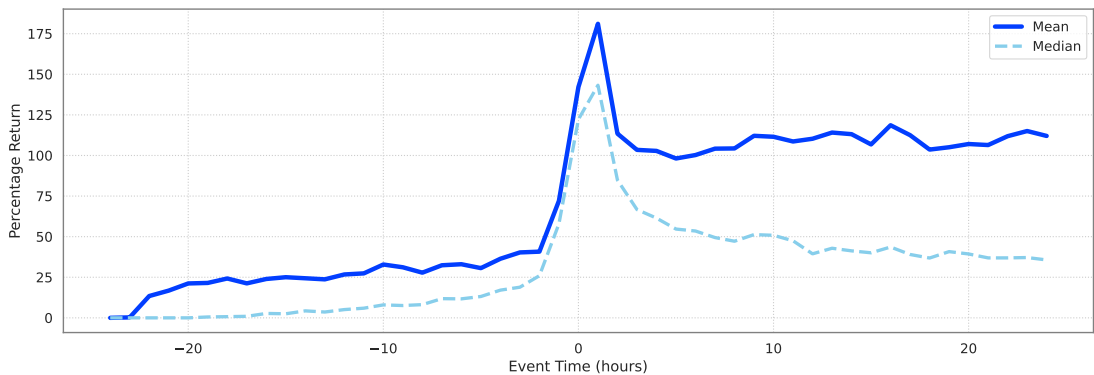


Figure 4. Events returns. Event time average (solid, in blue) and median (dashed, in light blue) of cumulative returns for the identified events, based on changes in the average sell price of the relevant NFT collections. The event time is in hours. The central point $t = 0$ corresponds to the first hour in which the cumulative return is larger than 100%, and the window extends to the 24 hours before and after that event. The mean and median are quite different, implying that event returns are highly skewed.

3.2 Crashes and Noncrashes

Next, we define an event as a *crash* if the ex-post return realized during the 24 hours following the run-up identification is lower than -40% .⁸ According to this criterion, we find that about 52% of the events result in a crash, ex-post. This is not surprising: a high proportion of incidences was expected given the high volatility and fast-paced growth of the NFT market.

The temporal distribution of price run-up events, classified as crashes and noncrashes, is

⁸The threshold does not constitute an arbitrary choice but, rather, it is taken from [Greenwood et al. \(2019\)](#). Our results are robust to varying the threshold from -20% to -80% .

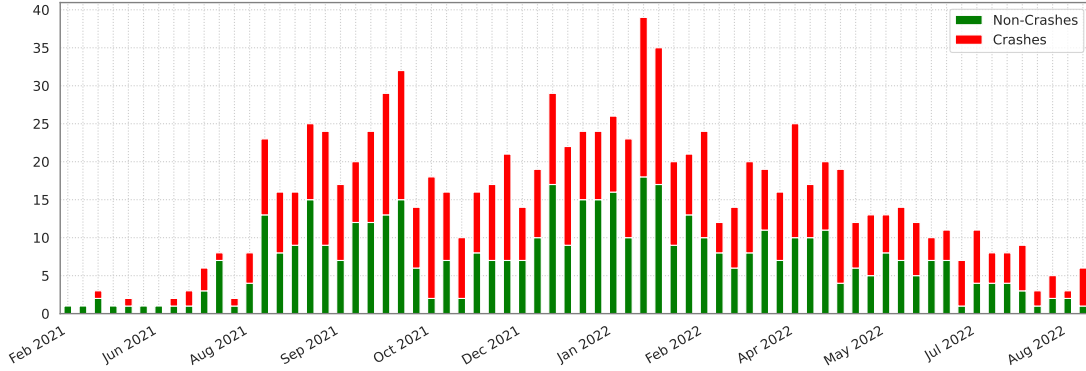


Figure 5. Events distribution. Distribution of price run-up events over time, aggregated at the weekly frequency and split between ex-post crashes (in red) and noncrashes (in green). Crash events are defined as those in which the ex-post return is lower than -40% after 24 hours from the run-up identification date.

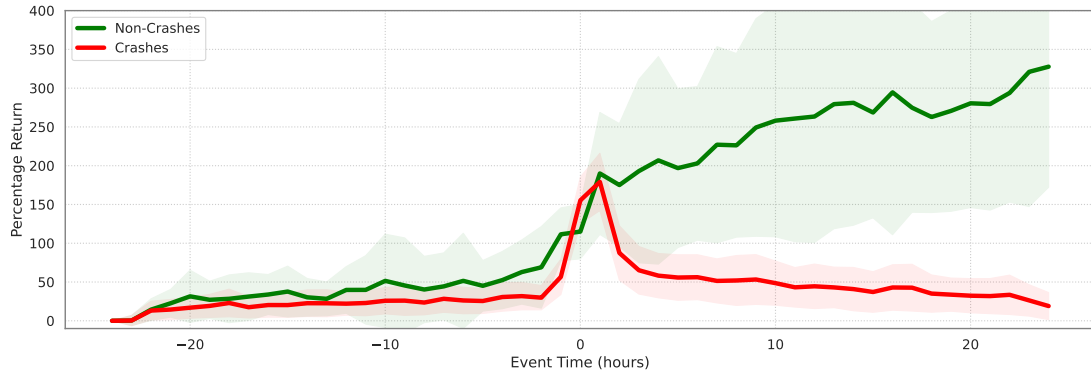


Figure 6. Crashes. Event time (in hours) average cumulative returns for the identified events split between ex-post crashes (in red) and noncrash (in green). The central point $t = 0$ corresponds to the first hour in which the cumulative return is larger than 100%, and the window extends to the 24 hours before and after that event.

illustrated in Figure 5. The depicted clusters notwithstanding, both categories appear to be distributed uniformly over time. To further illustrate the contrasting characteristics of these two categories, Figure 6 presents the cumulative returns of crashes and noncrashes separately. Notably, a conspicuous wedge separates the two categories. Specifically, while the appreciation in value from the ex-ante period is substantially wiped out for crashes, the noncrash events experience a significant surge in value, on average.

Figure 7 plots the distribution of trading volume and the number of active wallets

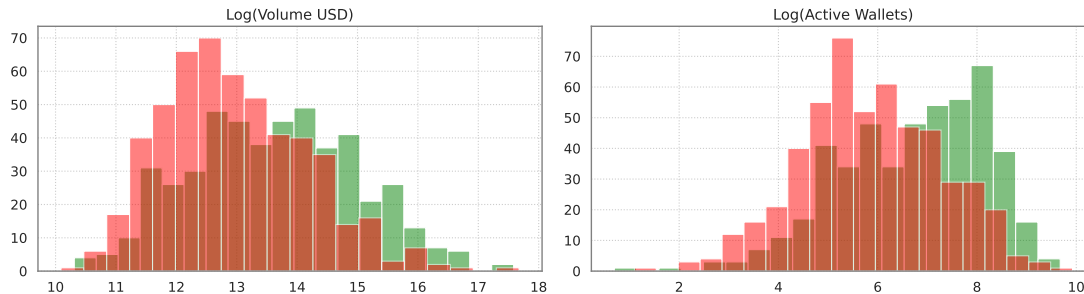


Figure 7. Crashes participation. The figure shows the trading volume (in USD) and the number of active wallets per event split between ex-post crashes (in red) and non-crashes events (in green), in logarithmic scale.

over the entire event window, split between ex-post crashes and noncrashes. It is noticeable that both metrics are lower for ex-post-crash events, as it is natural to expect, given that participation and liquidity are likely to deteriorate after the bubble bursts. Additionally, this suggests that ex-ante liquidity levels may be related to the probability of a crash realizing ex-post. We will test this hypothesis in Section 3.3.

3.3 Predicting Crashes

In this section, we assess the ex-ante predictability of ex-post crashes. More precisely, we check to see if ex-ante variables measured as of $t = 0$ can help predict whether the run-up event will lead to an ex-post crash, that is, to a cumulative return from $t = 1$ to $t = 24$ below -40% . We start by constructing *aggregate* event-level variables that are specific to each NFT collection but do not require agent-level information. These are estimated on the ex-ante period for $t \in [-24, 0]$. Building on the results of Greenwood et al. (2019), who identified a specific set of predictors for bubbles in the equity domain, our study centers on a specific set of variables, namely:

Volatility: The standard deviation of hourly returns.

Turnover: The number of sales divided by the circulating supply.

Age: The number of hours since the collection launch.

Acceleration: The difference between the cumulative return realized on the interval

$[-24, 0]$ and the one realized on $[-24, -12]$.

We use these variables as predictors in an event-level cross-sectional regression where the dependent variable is a dummy indicating an ex-post crash.

The regression estimates are reported in Table 2, unveiling a significant degree of crash predictability. As for equities (Greenwood et al., 2019), a crash is significantly more likely to happen when the return volatility and acceleration are higher in the ex-ante period.⁹ On the contrary, the probability is negatively correlated with the level of ex-ante turnover. In a context such as NFTs, in which it is not possible to sell short, this result squares with the idea that a higher turnover indicates a lower concentration of optimistic and overconfident investors, so as to decrease the risk of a speculative bubble (Scheinkman and Xiong, 2003). Moreover, high turnover can also predict lower liquidity risk, thus marking a smoother price adjustment process (Nneji, 2015).

In the last specification, where all variables are used as regressors, the coefficient of determination is greater than 21%, unveiling an economically significant degree of crash predictability. We will expand on this finding in Section 4, by including some additional agent-level predictors.

These results are robust to estimating a logistic regression model with heteroskedasticity-robust standard errors. Moreover, similar results, but with opposite signs, are obtained using an alternative specification in which the cumulative return in $[0,24]$ is used as the dependent variable. The R -squared coefficient is even larger in this alternative specification, reaching 36%. Finally, the results are robust to alternative definitions of ex-post returns, adjusting by an equally weighted return index or the NFT index proposed by Borri et al. (2022). In nonreported experiments, we run alternative specifications featuring additional regressors suggested by the literature. For instance, we found the skewness of ex-ante returns to be a positive predictor of crashes, but its

⁹Ex-ante volatility could be mechanically correlated with ex-post crashes since higher volatility increases the chance of realizing $\pm 40\%$ price movement. To rule out this possibility, we rerun the regression using as a dependent variable a dummy indicating ex-post returns higher than 40%. We find a significantly negative sign on volatility, suggesting that the potential mechanical effect, if anything, is of second-order importance.

Table 2. Crash Predictability. The table reports results from cross-sectional linear regressions based on the identified run-up events. The dependent variable is a dummy indicating an ex-post crash, that is, a cumulative return from $t = 1$ to $t = 24$ below -40%. The dummy is regressed onto a set of market-level variables measured as of the price run-up identification, that is, at $t = 0$, including: (i) *Volatility*, the hourly returns standard deviation estimated between $t = -24$ and $t = 0$; (ii) *Turnover*, the average turnover ratio, that is, the number of hourly sales transactions divided by the outstanding supply, estimated between $t = -24$ and $t = 0$; (iii) *Age*, the time from the creation of the NFT collection in hours, measured as $t = 0$; and (iv) *Acceleration*, the difference between the cumulative return realized in $[-12, 0]$ and the cumulative return realized in $[-24, 0]$. *T*-stats are reported in parentheses. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

Dep. Variable	(1) Crash	(2) Crash	(3) Crash	(4) Crash	(5) Crash
Volatility	0.4401*** (7.2525)				0.2714*** (4.5709)
Turnover		-0.3768*** (-5.5371)			-0.3199*** (-5.5626)
Age			0.0001*** (4.0524)		0.0001*** (3.1493)
Acceleration				0.1668*** (11.344)	0.1594*** (10.557)
Intercept	0.2582*** (6.6416)	0.5431*** (30.182)	0.4751*** (24.898)	0.3335*** (15.626)	0.1669*** (3.9723)
Observations	1,017	1,017	1,017	1,017	1,017
R-squared	0.0468	0.0192	0.0217	0.1601	0.2110

statistical significance evaporates once we control for *Acceleration*.

Overall, our results are consistent with the empirical evidence on U.S. and international industry portfolios (Greenwood et al., 2019) and theories of heterogeneous beliefs with shortsale constraints, which predict that volatility and price acceleration during price run-up events are positively related to the probability of a crash. The empirical findings of this study suggest that the emergence of speculative bubbles in conventional financial securities and novel digital securities, such as NFTs, adheres to similar dynamics. This suggests that the NFT market is not inherently more irrational than traditional financial markets, albeit learning about these innovative digital assets and their underlying technology may play a crucial role in this context (Pastor and Veronesi, 2009).

3.4 Market Factors

We now test whether run-ups and bubbles are driven by common factors, for example, by returns on a broad index of the NFT market. First, we apply principal component analysis to the panels of daily and weekly returns for our sample of collections. We find that the first five principal components explain less than 1% of the variation in returns. Second, we construct a market return factor by computing the daily average transaction prices across all collections of our sample and taking their percentage changes. We then run time-series regressions of collections' daily returns on the market factor. We find that the absolute value of the estimated market betas is, on average, 0.07 (median 0.05), and the average R -squared coefficient is 0.60% (median 0.05%). Similar results are obtained by using returns at the hourly or weekly frequency. More interesting results are obtained using as a market factor the NFT index proposed by [Borri et al. \(2022\)](#), with the absolute value of estimated market betas averaging 0.41 (median 0.19), and slightly better R -squared averaging 3.34% (median 0.40%). We conclude that, consistent with [Borri et al. \(2022\)](#), returns of NFT collections are highly idiosyncratic and poorly explained by the market index alone.

Next, to test the hypothesis that bubbles are clustered around specific time periods and autocorrelated, we run a regression similar to that reported in [Table 2](#), but including as additional regressors a variable counting the number of run-up events recorded in the previous 1 to 10 days, and another variable measuring the likelihood of crashes for those previous events. In unreported regression estimates, we found that both variables are not significant predictors of crashes. This is in line with the pattern displayed in [Figure 5](#), showing that bubbles are evenly distributed at a weekly frequency.

All in all, we conclude that bubbles are not clustered in specific time periods and that the market factor is not the main driver of price run-ups and their continuation. The fact that *outer* variables proxying for market conditions cannot predict crashes highlights the idiosyncratic nature of the run-up events in our sample. Motivated by this observation, we will take an *inner* look at the events in [Section 4](#), testing the

prediction power of event-specific, agent-level variables.

3.5 Crashes and Liquidity Dry-Ups

In this section, we evaluate the role of liquidity in speculative bubbles. Specifically, we explore two hypotheses: first, that liquidity decreases during bubble bursts, in line with the narrative about speculative bubbles (Kindleberger, 1978). Second, limited liquidity increases the likelihood of a speculative bubble bursting, as predicted by models with rational and financially constrained agents demanding and supplying liquidity (Farhi and Tirole, 2011). To do so, we regress different measures of liquidity and trading activity in the ex-post period on the crash indicator dummy. The measures include: (i) the average turnover ratio in the ex-post window; (ii) the Amihud (2002) illiquidity ratio, defined as the ratio between the absolute cumulative return from $t = 1$ to $t = 24$ and the volume traded over the same period; and (iii) the hourly volatility estimated in the ex-post window. Table 3 presents the regression results of the first three specifications (1) to (3), indicating significantly lower liquidity in the ex-post period for crash events, which is consistent with the liquidity dry-up predicted by the asset bubble narrative. Next, to explore the predictability of ex-post liquidity by ex-ante variables, we incorporate an inverse measure of the change in ownership concentration during the price run-up — referred to as *Unique Owners* — as the explanatory variable in specifications (4) to (6). This variable is extensively defined and analyzed in Section 4.2, and a higher number of unique owners is interpreted as a larger and more diverse trading community that facilitates the resale of financial securities, reduces search costs, and increases risk sharing. The regression results suggest that a decreasing number of owners ex-ante positively predicts ex-post illiquidity. Finally, specifications (7) to (9) employ the ex-ante wash-trading volume as a regressor, as defined and estimated in Section 4.3. The estimated coefficients indicate that higher ex-ante wash trading is linked to aggravated ex-post illiquidity. These findings corroborate the notion that the fictitious volume generated by wash trades during the run-up phase may culminate

Table 3. Liquidity Dry-Up. The table reports results from cross-sectional linear regressions based on the identified run-up events. In specifications (1), (4), and (7), the dependent variable is the turnover ratio, that is, the number of hourly sales transactions divided by the outstanding supply, estimated in the ex-post window [1, 24]. In specifications (2), (5), and (8), the dependent variable is the Amihud ratio, defined as the ratio between the absolute cumulative return from $t = 1$ to $t = 24$ and the volume traded over the same period. In specifications (3), (6), and (9), the dependent variable is the hourly volatility estimated as the standard deviation of hourly returns in the ex-post window [1, 24]. In specifications (1) to (3), the dependent variables are regressed onto the dummy indicating an ex-post crash, that is, an ex-post cumulative return lower than -40% . In specifications (4) to (6), the dependent variables are regressed onto the change in the ratio of unique owners and the supply of NFT tokens from $t = -24$ to $t = 0$. In specifications (7) to (9), the dependent variables are regressed onto the natural logarithm of the amount of wash-trading volume identified for the collection before $t = 0$. T -stats are reported in parentheses. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

Dep. Variable	(1) Turnover	(2) Amihud	(3) Volatility	(4) Turnover	(5) Amihud	(6) Volatility	(7) Turnover	(8) Amihud	(9) Volatility
Crash Dummy	-0.03*** (-6.04)	0.0*** (3.23)	0.11*** (7.09)						
Unique Owners				0.01*** (5.58)	-0.0*** (-4.93)	-0.02*** (-6.43)			
Wash Trading							-0.01*** (-8.63)	0.01 (1.18)	0.02*** (2.85)
Intercept	0.04*** (9.53)	0.01*** (5.10)	0.48*** (35.57)	0.02*** (8.51)	0.01*** (12.14)	0.56*** (50.46)	0.04*** (10.89)	0.01*** (8.25)	0.51*** (40.57)
Observations	1,017	1,017	1,017	1,017	1,017	1,017	1,017	1,017	1,017
R-squared	0.04	0.01	0.05	0.12	0.01	0.04	0.06	0.00	0.01

in more illiquid trading during the correction phase, underscoring the importance of heterogeneous and genuine investors in invigorating the market.

4 Agent-level Analysis

The objective of this section is to develop a set of metrics for analyzing individual investor behavior in the NFT market, and to explore their predictive power for identifying crashes. To this end, we construct three agent-level variables based on economic arguments. The first variable identifies a group of sophisticated investors, sorting by the realized profits during run-up events. We validate their superior abilities in the wider decentralized finance (DeFi) space by analyzing their blockchain transactions, and examine the relation between their involvement in run-ups and the occurrence of

crashes. The second variable is a measure of ownership concentration, which is determined by counting the unique owners of a collection. We test whether this metric denoting investor heterogeneity is related to crash probability and ex-post illiquidity. The third variable involves identifying suspicious transactions that are likely used to artificially inflate the trading volume of a collection (wash trading) and testing whether the degree of this phenomenon is a predictor of ex-post crashes and illiquidity.¹⁰

4.1 Sophisticated Agents

The motivation underlying this analysis stems from the narrative surrounding speculative bubbles, as documented in the literature (e.g., [Kindleberger, 1978](#); [Brunnermeier, 2016](#)). This narrative posits that market participants can be broadly classified into two groups: speculators and sophisticated investors. While the former are primarily interested in the resale price, the latter possess the ability to link the market value of a financial security to its fundamental value. However, in the case of NFTs, identifying the fundamental value is challenging due to the absence of cash flow streams, the unique characteristics of these digital assets, and the innovative nature of the underlying technology. When speculators become the dominant force in the market, the likelihood of a bubble occurring increases. The existing literature offers various aspects that differentiate economic agents, including but not limited to asymmetric information (e.g., [Allen et al., 1993](#); [Hong and Stein, 2007](#)), arbitrage (e.g., [Shleifer and Vishny, 1997](#); [Abreu and Brunnermeier, 2003](#)), heterogeneous beliefs (e.g., [Scheinkman and Xiong, 2003](#)), the ability to time the market (e.g., [Brunnermeier and Nagel, 2004](#)), and risk premia combined with learning about new technologies (e.g., [Pastor and Veronesi, 2009](#)). Our underlying assumption is that these factors may give rise to heterogeneity among investors, which, in turn, can affect their performance consistently over time.

As a preliminary step, we aim to investigate whether certain investors consistently

¹⁰It is important to acknowledge that individual investors may possess multiple wallets, which could introduce errors into the agent-level variables constructed using data at the wallet-address level. However, we contend that such additional noise is likely to hinder the identification of statistically significant predictors of crashes.

outperform others, that is, whether there exists a group of sophisticated investors who possess the ability to generate gains by riding successive price run-ups while avoiding price corrections. This would serve as a first indication of their sophistication. To this end, we compute the percentage profits obtained by individual investors during the entire run-up event window of $[-24, 24]$. We find that the profitability of investors is strongly autocorrelated across different run-up events, with a significant autoregressive coefficient of 0.21 (t -stat = 27.59). These results imply that some investors consistently excel at riding price run-ups and avoiding crashes, while others consistently perform poorly. Motivated by these findings, we define a group of *sophisticated* investors based on their past performance. Specifically, for each event e and investor a , we consider the investor’s profitability in the previous five events in which they actively participated. We require sophisticated investors to have participated in at least five run-up events and to have realized an average profit of more than 25% on their invested capital during the previous five events. Profitability is calculated over the entire event window, from $t = -24$ to $t = 24$. If an NFT is not sold before $t = 24$, its profit is based on the average transaction price at $t = 24$. We define the first agent-level variable as the percentage of sophisticated investors actively participating in the ex-ante period of a run-up event, relative to the total number of active wallets during the same period.

Before moving to the definition of the remaining two agent-level variables, the next subsections provide more information about the identified sophisticated agents, focusing on their skills in the broader DeFi ecosystem and their ability to properly time NFT bubbles.

4.1.1 Characteristics of Sophisticated Agents

We now conduct a detailed analysis of the behavior of sophisticated agents, so as to conjecture whether they have superior abilities even outside of the NFT scope. As discussed earlier, factors that might distinguish them are (a) superior information extracted from trading or learned by observing the price discovery process or (b) honed

skills in conducting their trading activities. To this end, we have augmented our existing data set by procuring a full record of all transactions on the Ethereum blockchain conducted by each wallet in our sample from the platform [Etherscan](#). These transactions encompass a wide range of activities, including not only the straightforward transfer of ether or ERC-20 tokens, but also interactions with diverse smart contracts, even those that are not directly related to NFTs.

Drawing on our extensive data set, we proceed to investigate two key determinants that are potentially related to the level of sophistication exhibited by market agents. We explore the impact of two factors, namely, the frequency and value of transactions and the agents' proficiency in certain trading activities, on the performance of sophisticated traders in NFT markets. First, we contend that increased levels of trading activity enable sophisticated agents to extract a greater volume of *information* from the market, thereby enhancing their capacity to acquire a more nuanced understanding of fundamental value compared to market price evolution during the price discovery process. In this vein, we scrutinize the extent to which the superior performance of sophisticated agents correlates with more frequent and valuable transactions that are primarily attributable to informed traders, according to microstructure theories of market dynamics (e.g., [Easley and O'Hara, 1987](#)). Moreover, we surmise that the extended presence of sophisticated agents in the market provides a distinct advantage, especially in the context of nascent assets like NFTs, whose fundamental value remains obscure. Second, we posit that sophisticated traders exhibit superior *ability* in at least three domains, namely, trading with precise timing to maximize gains, leveraging their positions to exploit profit opportunities and access arbitrage capital, and engaging in cross-market trading.

Table 4 presents an overview of the average characteristics of agents who are identified as *sophisticated* in at least one event, taking into account NFT-related metrics, such as profits, trading volume, holding period, and the number of NFT transactions, in conjunction with supplementary metrics derived from the blockchain data. Specifically,

we construct these metrics by counting interactions with a set of DeFi-related smart contracts, including those unrelated to NFTs, over the entire lifespan of each wallet account and measured as of December 2022.

The resulting metrics comprise the number and total value of transactions conducted in ETH, the wallet’s age measured in days, as well as swaps and liquidity provisions on decentralized cryptocurrency exchanges (DEX) such as Uniswap or Curve, and lending or borrowing operations on decentralized lending platforms such as AAVE or Compound. We note that a similar set of results can be obtained by focusing only on blockchain transactions conducted before the start of our sample period, which is January 2021.

Table 4 presents evidence that agents who outperform, as defined above based on their realized profits during price run-up events, display greater market presence both in the time series and across different DeFi applications, as well as higher levels of financial leverage. In particular, these agents exhibit greater sophistication, as evidenced by the fact that they created their wallets earlier and have conducted more transactions on the blockchain in terms of both size and frequency. These findings are consistent with previous empirical evidence on the formation of bubbles in traditional financial securities (Greenwood and Nagel, 2009) and laboratory experiments (Smith et al., 1988). Furthermore, outperforming agents interact more frequently with DeFi protocols, such as decentralized exchanges, which suggests cross-market ability and information extraction. Most importantly, they engage in more decentralized lending platform activities, which indicates a more proactive use of leverage and arbitrage capital. Taken together, these results suggest that our definition of sophisticated investors, based on NFT trading performance, captures individual retail investors who are more experienced and sophisticated agents in the broader DeFi space.

4.1.2 Timing the Bubble

Another trait usually associated with sophisticated traders is the ability to *time the bubble*. The objective of this section is to test this hypothesis by examining whether

Table 4. Sophisticated Agents. The table displays the average characteristics of the wallet accounts participating in our run-up events, split between sophisticated and nonsophisticated investors, defined in Section 4.1. The first set of characteristics (rows 1 to 4) is constructed using data from our run-up events only and includes the following: the number of trades, the generated trading volume, the hourly percentage profits accrued by trading NFTs, and the holding period for single NFTs. The second set of characteristics (rows 5 to 11) is constructed using blockchain data comprising all transactions performed by those wallet accounts on the Ethereum blockchain, including those unrelated to NFTs, across their entire existence, as of December 2022. This set includes: the total number of transactions (including simple ETH transfers and interactions with smart contracts), the total ETH value attached to those transactions, the number of days that elapsed since the first transaction performed by the wallet, the number of swap transactions in decentralized exchanges (DEX; specifically, Uniswap v2, Uniswap v3, Sushiswap v2, Metamask Swap, Ox Exchange, OneInch v2, OneInch v3, OneInch v4, Balancer, and Curve), the number of liquidity deposits and withdrawals in the same list of exchanges, and the number of lending and borrowing operations in decentralized lending platforms (AAVE v1, AAVE v2, and Compound). Qualitatively similar differences are obtained by using only transactions performed before the beginning of our samples, as of January 2021. The t -stats of the differences are reported in the last column. Asterisks denote significance levels (***= 1%, ** = 5%, * = 10%).

	Sophisticated	Others	Difference	t-stat
NFT Market				
Number of Trades	22.41	3.05	19.36	136.26***
Traded Volume (ETH)	8.49	1.43	7.06	98.37***
Percentage Profits	94.92	69.41	25.51	44.59***
Holding Period (hours)	7.76	5.72	2.04	62.25***
Blockchain				
Number of Transactions	1,088.48	273.17	815.30	123.40***
Transacted Value	631.61	220.70	410.91	13.83***
Age (days)	519.13	429.89	89.24	50.24***
DEX Swaps	75.24	24.40	50.83	34.59***
DEX Liquidity	3.58	1.25	2.32	19.64***
DeFi Lending	0.30	0.08	0.22	6.17***
DeFi Borrowing	0.79	0.27	0.52	6.99***

they skillfully time the bubble during ex-post crash events, constructing an agent-level measure of market timing, TS , standing for *Timing Score*. To define the measure, for each ex-post crash, we first identify the peak of the bubble, that is, the event time at which the average sale price reaches its maximum in the ex-post window following the price run-up identification. Denote the peak time as t^* , expressed in event time units (at the hourly frequency). We then define the distance to the peak as a translation of the event time as $d = t - t^*$. The timing scores for buy and sell orders are two piece-wise linear functions of d , with a kink at $d = 0$, defined on the window $d \in [-24, 24]$ centered around the peak time. Let us denote these two functions by $TS_{buy}(d)$ and $TS_{sell}(d)$,

respectively, and define them as

$$\begin{cases} TS_{buy}(d) = -d - 12 & \text{for } d \leq 0 \\ TS_{buy}(d) = d - 12 & \text{for } d > 0 \end{cases} \quad (1)$$

and $TS_{sell}(d) = -TS_{buy}(d)$. In other words, the timing score for buy orders decreases linearly while moving closer to the peak, and it increases linearly after the peak. This is in line with the intuition that it is better to buy early in order to fully profit from the price run-up. The opposite holds for a sell order, where perfect timing involves selling as close as possible to the peak. We compute the timing score for an investor participating in a given event as

$$TS = \sum_{d=-24}^{24} B(d) \cdot TS_{buy}(d) + S(d) \cdot TS_{sell}(d), \quad (2)$$

where $B(d)$ and $S(d)$ count the number of buy and sell trades performed by the investor at time d , respectively. The specified definition of the timing score implies that any investor who buys and sells an NFT within the same hour will not receive a positive contribution to their timing score. This restriction is necessary to ensure the accuracy and usefulness of the metric.

We compute the timing score TS for all investor-event pairs, restricting to crash events. Additionally, we compute the *buy* and *sell* legs separately. Next, we compute percentile rankings by event for the obtained metrics to properly compare investors' performance within each event. We then run a panel regression of the TS metrics on a dummy indicating sophisticated investors. Results are reported in Table 5 and show that sophisticated investors achieve better timing scores across all specifications. The coefficients on specifications (1) and (4) imply that sophisticated agents enjoy, on average, a higher percentile ranking of the within-event timing score by 2 percentage points. We thus conclude that sophisticated investors are more skilled in timing the bubble during ex-post crash events by both buying early into the run-up phase and selling close to

Table 5. Market Timing. The table reports results from linear regressions at the agent-event level, restricting to the sample of crash events. The dependent variables are proxies for the investor’s ability to time the bubble, whose definition is described in the text. The proxies are regressed onto a dummy variable indicating investors’ sophistication. In specifications (1) to (3), the explanatory variable is the time-varying sophistication indicator defined in 4.1, while in specifications (4) to (6) the dummy is time-invariant, that is, an agent is considered sophisticated if she is marked as such at least once in the time-varying version. Standard errors are clustered by event. *T*-stats are reported in parentheses. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

Dep. Variable	(1) Timing	(2) Buy Timing	(3) Sell Timing	(4) Timing	(5) Buy Timing	(6) Sell Timing
Sophisticated (time-varying)	0.0209*** (12.791)	0.0100*** (6.0698)	0.0107*** (14.365)			
Sophisticated (constant)				0.0204*** (16.356)	0.0080*** (6.3915)	0.0118*** (22.123)
Intercept	0.4971*** (839.90)	0.4992*** (846.56)	0.4991*** (2077.1)	0.4930*** (766.76)	0.4979*** (778.76)	0.4965*** (1998.2)
Observations	214,855	214,855	214,855	214,855	214,855	214,855
R-squared	0.0011	0.0003	0.0015	0.0016	0.0003	0.0028

the peak of the bubble.

4.2 Ownership Concentration

To construct a second agent-level variable for each run-up event, we compute the *fraction of unique owners* of the relevant collection, at an hourly frequency, defined as the ratio between the number of unique wallet addresses holding at least an NFT of the collection and the total number of NFTs available. This metric, displayed on OpenSea and other prominent marketplaces for NFTs, can be interpreted as an inverse measure of ownership concentration or investor heterogeneity. It is expected to be a positive predictor of future performance for a number of reasons. First, in the same spirit as the stock price fragility in [Greenwood and Thesmar \(2011\)](#), a high fraction of unique owners indicates the lack of large holders that may harm the collection’s value by putting a large number of NFTs on sale, potentially motivated by a liquidity shock. Second, a larger and more diverse community of traders should facilitate the resale of financial securities, reduce search costs, and increase risk sharing. Thus, it positively correlates

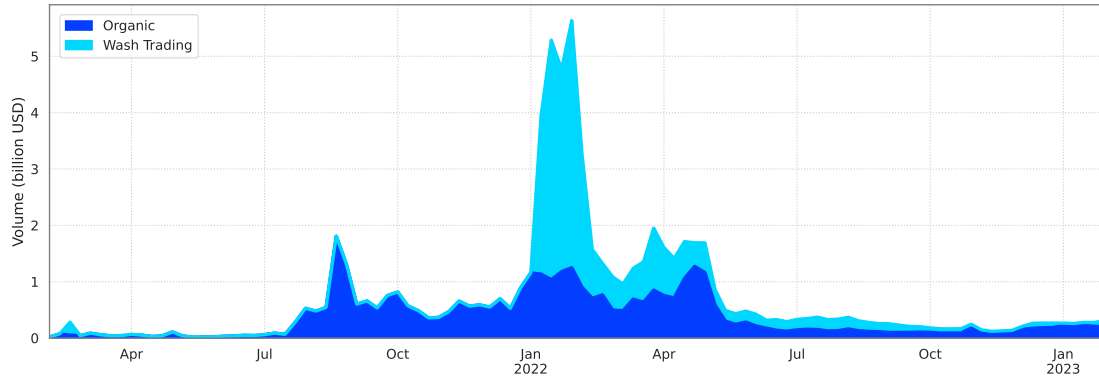


Figure 8. Wash trading across all exchanges. The figure displays the daily historical trading volume in USD across all NFT marketplaces, split between *wash-trading* and organic trading volume. The former is estimated using the Hildobby method, as proposed in Hildobby (2022), which incorporates four transaction-level conditions to identify potentially suspicious trades.

with the level of liquidity of the collection, indicating a large number of potential buyers and a high level of interest in the project. Finally, it proxies for the size of the community supporting the NFT project. During a price-run up, an increase in unique holders may signal that the rally comes from a genuine interest in the project. On the contrary, a run-up accompanied by a decrease in unique holders signals that the price dynamic is driven by a limited number of buyers who could potentially influence or even manipulate prices in their favor, an issue we study below. We thus use the change in the fraction of unique owners during the price run-up (from $t = -24$ to $t = 0$) as an additional variable to predict the ex-post price evolution.

4.3 Wash Trading

According to the American Commodity Futures Trading Commission, wash trading is the practice of “entering into, or purporting to enter into, transactions to give the appearance that purchases and sales have been made, without incurring market risk or changing the trader’s market position” (?). Wash trading could be especially prevalent on cryptocurrency (unregulated) exchanges, with previous research estimating that it may account for more than 70% of total trading volume on unregulated exchanges (Cong et al., 2021). In the NFT market, wash trading is usually conducted to artifi-

Table 6. Wash Trading by Exchange. The table presents statistics for the top five NFT marketplaces by trading volume during the period from February 1st, 2021, to February 9th, 2023, based on data provided by Dune Analytics. The first column reports the total traded volume in USD, while the second column reports the volume marked as *wash trading* according to the Hildobby method, proposed in Hildobby (2022), and the third column displays the percentage of wash-trading volume relative to the total. Columns (4) to (5) repeat the exercise using transaction count instead of transaction volume.

	Total Volume	Wash Volume	Wash %	Total Trades	Wash Trades	Wash %
OpenSea	32,797,993,136 \$	761,535,142 \$	2.32%	27,270,308	181,410	0.67%
LooksRare	26,590,247,528 \$	26,064,869,084 \$	98.02%	115,288	27,247	23.63%
X2Y2	4,025,756,413 \$	3,432,460,213 \$	85.26%	587,407	125,241	21.32%
CryptoPunks	2,495,977,352 \$	28,220,202 \$	1.13%	22,427	328	1.46%
Blur	1,311,868,278 \$	142,010,668 \$	10.83%	1,766,529	27,044	1.53%

cially inflate prices within a collection by repeatedly trading tokens and thus increasing the total volume of the collection. These potential dynamics are relevant to our work for two reasons: First, wash trading may foster artificial increases in collections’ valuations, thus increasing the probability of a subsequent crash. Second, neglecting such a practice could harm the identification of sophisticated investors based on agent-level profitability. We therefore proceed and identify wash-trading transactions, with the two-sided purpose of (i) computing an additional agent-level predictor of crashes and (ii) testing the robustness of our results to the presence of wash trading.

Figure 8 shows the evolution of estimated wash-trading volume over time. There is a noticeable increase in volume concentrated in early 2022, but much of the time the volume remains below \$1 billion. Table 6 shows that most of the wash trading is performed on exchanges other than OpenSea. While wash-trading volume on OpenSea — across all existing collections — is estimated at around 2% of the total, other exchanges feature much higher shares of fraudulent transactions. For instance, this applies to LooksRare and X2Y2 for which rates of 98% and 85%, respectively, have been estimated. When the number of wash transactions is considered, the same pattern holds, with only 0.67% of the transactions in OpenSea being identified as wash trades. Since our data set covers only transactions performed in OpenSea, we expect a relatively low share of wash volume. Nevertheless, we use two distinct approaches to estimate the magnitude of such a phenomenon quantitatively. The first is inspired by Cong et al.

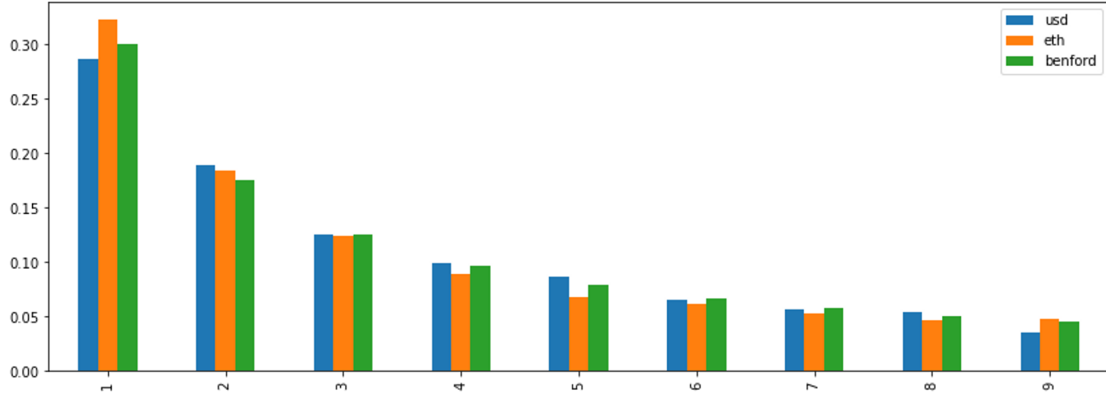


Figure 9. Benford’s Law. The observed distribution of the leading digits of USD (in blue) and ETH (in orange) transaction prices and the theoretical distribution predicted by Benford’s Law (in green).

(2021) and provides an aggregate test for the presence of wash-trading transactions. It involves testing whether the distribution of transactions’ amounts follows patterns that are typical of such distributions. Benford’s Law, frequently used to identify fraud, states that the probability of a digit from 1 to 9 appearing as the leading digit of a number adheres to the formula:

$$\text{Prob}(N \text{ is the first significant digit}) = \log_{10}(1 + N^{-1}), \quad (3)$$

with $N \in \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. We test Benford’s Law applied to transaction prices in our sample, both in USD and ETH denominations. The leading digits’ distribution in our sample shows no significant deviation from Benford’s Law, with Chi-statistic tests’ p -values approaching unity.¹¹ Based on transaction prices, Figure 9 shows the difference between the theoretical distribution predicted by Benford’s Law and the observed distribution. The second approach that we employ is based on the method proposed by Hildobby (2022), which is tailored for blockchain transactions. By capitalizing on the transparency of the decentralized ledger, this approach enables a transaction-level identification of wash trading. The methodology involves implementing four filters on

¹¹We also estimate the power-law cumulative density function exponent. Research has shown that such an exponent lies in the range from 1 to 2 for traditional assets and cryptocurrencies (Cong et al., 2021). In our case, the exponent value amount to approximately 2.16, which constitutes a small deviation from the aforementioned range.

our database of transactions to identify potentially dubious trading behavior. Those are: (i) the buyer and seller wallet addresses are identical, (ii) a pair of trades with seller and buyer inverted for the same NFT, (iii) the same address has bought three or more times the same NFT, and (iv) the addresses that first funded the wallets of buyer and seller with ETH are identical or are one another.¹² We apply this methodology restricting to our sample of collections and the time period we cover in the analysis. The process results in the identification of 33,293 transactions (0.23% of the sample), corresponding to a trading volume of 63 million \$ (0.35% of the total). The fraction of wash trading is insignificant for the median collection in our sample (0.22%). Only the *Og-Crystals* collection stands out, with the estimated wash-trading volume accounting for 73.6% of the total. As a robustness check, we also apply the methodology employed in [La Morgia et al. \(2022\)](#), based on graph theory. The transactions identified by the two methods are highly correlated, but the Hildobby method is less conservative, and the identified wash volume is more than two times larger. The methodology can thus be thought of as an upper bound for the real amount of detectable wash trading.

We use the resulting set of wash-trading transactions to construct a variable at the event level. Namely, for each price run-up event, we define the *Wash Volume* metric as the natural logarithm of one plus the total wash volume transacted (in ETH units) on the relevant collection before the run-up identification hour ($t = 0$). This variable will be used as a regressor in the cross-sectional regression of [Section 4.4](#).

4.4 Agent-level Crash Predictability

We now turn our attention to the question of whether agent-level variables can serve as predictors of asset bubbles. To investigate this, we incorporate the three agent-level variables studied above as additional regressors in a regression specification similar to that of [3.3](#), but with the crash indicator as the dependent variable. First, we add

¹²While applying condition (iv), a list of roughly 500 addresses belonging to centralized crypto exchanges and all existing tornado cash smart contracts are excluded.

Table 7. Agent-level Crash Predictability. The table reports results from cross-sectional linear regressions based on the identified run-up events. In specifications (1) to (3), the dependent variable is a dummy indicating an ex-post crash, that is, a cumulative return from $t = 1$ to $t = 24$ below -40%. In specifications (4) to (6), the dependent variable is the cumulative return from $t = 1$ to $t = 24$. The dependent variables are regressed onto the same set of market-level variables as in Table 2 and a set of agent-level variables, both measured as of the price run-up identification, at $t = 0$. The agent-level variables include: (i) *Sophisticated*, the ratio between the number of sophisticated investors actively trading the collection in the ex-ante period $[-24, 0]$, defined in Section 4.1, and the number of total investors actively trading the collection in the same period; (ii) *Unique Owners*, the change in the fraction of unique owners divided by the supply of NFT in the collection recorded during the ex-ante window $[-24, 0]$; (iii) *Wash Trading*, the natural logarithm of one plus the total wash-trading volume in ETH identified by the method described in 4.3 before $t = 0$. T -stats are reported in parentheses. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

Dep. Variable	(1) Crash	(2) Crash	(3) Crash	(4) Ex-Post Return	(5) Ex-Post Return	(6) Ex-Post Return
Sophisticated	-0.5114*** (-3.4262)		-0.4155*** (-2.8422)	2.1993*** (4.8323)		1.7250*** (3.8923)
Unique Owners	-0.0114* (-1.8545)		-0.0070 (-1.0926)	0.0585*** (3.1115)		0.0403** (1.9849)
Wash Trading	0.0770*** (6.6193)		0.0613*** (5.0739)	-0.2718*** (-7.7302)		-0.1988*** (-6.4099)
Volatility		0.2714*** (4.5709)	0.1785*** (2.9338)		-1.3846*** (-9.0643)	-1.0269*** (-5.9488)
Turnover		-0.3199*** (-5.5626)	-0.2224*** (-3.9693)		1.0092*** (4.8642)	0.6125** (2.1825)
Age		0.0001*** (3.1493)	0.0001 (0.0698)		-0.0002*** (-4.9529)	-0.0002 (-0.3302)
Acceleration		0.1594*** (10.557)	0.1613*** (10.781)		-0.5686*** (-12.267)	-0.5767*** (-12.848)
Intercept	0.5183*** (15.389)	0.1669*** (3.9723)	0.2400*** (4.8339)	-0.6695*** (-7.1665)	0.9322*** (8.5946)	0.5826*** (4.0961)
Observations	1,017	1,017	1,017	1,017	1,017	1,017
R-squared	0.0700	0.2110	0.2391	0.1263	0.3641	0.4101

the fraction of sophisticated agents active during the event in the ex-ante window $[-24, 0]$, denoted as *Sophisticated*. The underlying sophistication indicator is based on the realized returns on the previous five events, as discussed earlier. Second, we include the change in the fraction of unique owners recorded during the ex-ante window $[-24, 0]$, denoted as *Unique Owners*. Third, we incorporate the total wash-trading volume realized for the collection before $t = 0$, denoted as *Wash Trading*.

Table 7 presents the outcomes of the aforementioned exercise, indicating that agent-level variables considerably enhance the predictive power of the model in terms of forecasting both crashes and ex-post returns. Specifically, three primary observations emerge from the results: First, the presence of sophisticated investors during the ex-ante period is a negative indicator of ex-post crashes. Similarly, the fraction of sophisticated agents positively predicts ex-post returns, exhibiting high levels of statistical significance. Second, the change in unique holders exhibits a negative association with the ex-post crash probability, albeit with a lower level of statistical significance. The relation with future returns has a positive point estimate, though it is not statistically distinct from zero. Third, a higher quantity of wash trading in the period prior to the run-up identification date is correlated with an increased crash probability and lower ex-post returns. These associations are highly statistically significant in all specifications. Taken together, the three additional agent-level predictors elevate the predictive R -squared coefficient by nearly 3% for the prediction of crashes and almost 5% for the forecast of ex-post returns, thus increasing the predictive power of the model by more than 10%. As we will document in Section 4.5, this constitutes an economically significant prediction improvement.

These results indicate a correlation between the probability of crashes following price run-ups and different genuine (or nonmanipulative) sophisticated investors, but they do not imply a causal relation. One may conjecture that sophisticated agents are able to estimate ex-ante the probability of ex-post crashes and, thus, self-select into non-crash events. Although additional analyses are needed to establish causality, the results reported here are consistent with the asset bubble narrative that the predominance of pure speculation is a major factor in the creation of bubbles destined to burst. Further, these findings suggest that wash trading may play a role in the formation of bubbles in the NFT market. A high amount of wash-trading volume may increase the collection's price artificially, thus fostering the formation of a price run-up and increasing the probability of a subsequent crash.

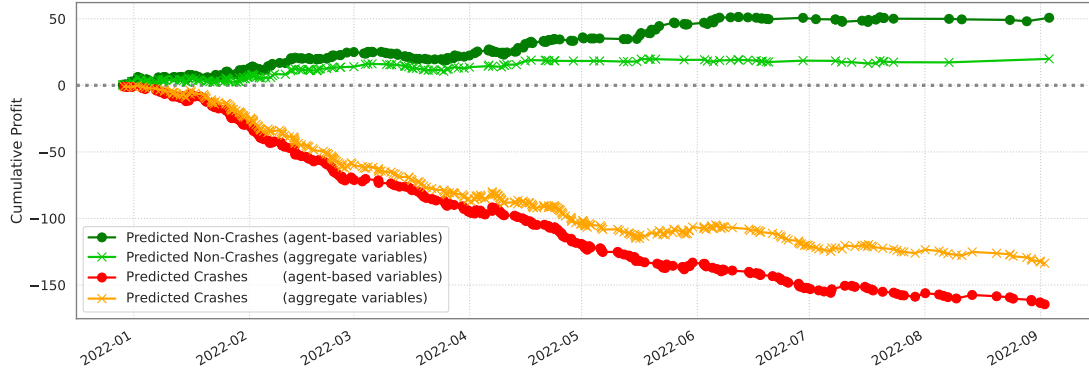


Figure 10. Out-of-sample performance. The figure displays the out-of-sample cumulative returns of two portfolios, constructed using a prediction model estimated with agent-level and market-level variables (red and green lines), and two portfolios constructed using a prediction model estimated with market-level variables only (orange and light green lines). The prediction models are estimated using data from the 478 run-up events realized in 2021, while the displayed portfolio returns are based on the 539 run-up events realized in 2022. The returns of the “Predicted Noncrashes” portfolios result from buying into a collection at time $t = 1$ when the model prediction is *below* the median of the estimation sample predictions, and selling at $t = 24$. Similarly, the returns of the “Predicted Crashes” portfolios result from buying into a collection at time $t = 1$ when the model prediction is *above* the median of the estimation sample predictions, and selling at $t = 24$. Entry and exit prices are based on average sale prices for the collection in the relevant hour.

We finally rerun the full agent-level analysis by excluding the 33,293 identified wash-trading transactions. Our purpose for this exclusion is to assess the impact on the computation of profits for each agent and, consequently, the identification of sophisticated investors. After conducting this analysis, we find that our results remain virtually unchanged and exhibit a slight increase in statistical significance. Based on these findings, we conclude that wash trading does not significantly influence our results, but rather contributes only a minimal amount of noise to our data set.

4.5 Economic Value of Agent-level Variables

To economically assess the level of predictability and, in particular, the value of agent-level variables, we propose an out-of-sample simple trading strategy based on our regression results. More precisely, we fit regression models similar to those displayed in columns (2) and (3) of Table 7, but using only data from run-up events realized in 2021. These account for 478 events, about half of the sample. We then use the fitted

models to produce predictions on the likelihood of an ex-post crash for the remaining 539 events in the sample, those recorded in 2022, and we construct two calendar-time portfolios based on these predictions. For the first portfolio, when the predicted crash probability is below the median value of the 2021 sample, 1 ETH is invested at time $t = 1$, and the position is liquidated at the end of the ex-post window ($t = 24$). Similarly, for the second portfolio, we invest when the predicted crash probability is *above* the median. The profits and losses are computed based on average traded prices. It should be emphasized that these portfolio strategies are not fully implementable and tradable given the impossibility of short selling. Figure 10 presents the results from the exercise, showing the performance of the two portfolios. The orange and lightgreen lines represent the time evolution of the economic value of predictions based *only* on market-level variables. The red and green lines are based on a prediction model estimated with market-level variables *and* agent-level predictors. The plot shows that the ex-ante predictability of crashes following price run-ups can generate economically significant profits, primarily when agent-level variables are used for the estimation of the prediction model. Using only aggregate variables, the portfolio of predicted non-crashes generates a profit of 19 ETH, while the portfolio of predicted crashes loses 133 ETH. The wedge between the two portfolios significantly increases when we also exploit agent-level variables, generating a profit of 50 ETH and a loss of 164 ETH, respectively.

5 Conclusion

This paper presents a comprehensive study of nonfungible tokens (NFTs), which are unique digital assets recorded on the blockchain, providing proof of ownership and authentication. After constructing a representative sample of the total NFT collection market, we systematically analyze all price run-up events and those that lead to ex-post crashes.

The primary contributions of this research are twofold: first, establishing the pre-

dictability of asset bubbles and, second, analyzing the behavior of individual retail investors during asset bubbles. The availability of blockchain technology and the transparency of the NFT market allow us to study these phenomena in a unique way. Our analysis reveals that agent-level variables, such as investor sophistication and heterogeneity, in conjunction with asset-specific variables, such as volatility, price acceleration, and turnover, significantly predict bubble formation and price crashes. Sophisticated investors consistently outperform their peers, with indications that they possess superior information and skills, which aligns with existing narratives surrounding asset pricing bubbles.

References

- Abreu, D. and Brunnermeier, M. K. (2003). Bubbles and crashes. *Econometrica*, 71(1):173–204.
- Allen, F., Morris, S., and Postlewaite, A. (1993). Finite bubbles with short sale constraints and asymmetric information. *Journal of Economic Theory*, 61(2):206–229.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5:31–56.
- Ante, L. (2022). The non-fungible token (nft) market and its relationship with bitcoin and ethereum. *FinTech*, 1(3):216–224.
- Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, 129(2):203–227.
- Barbon, A. and Ranaldo, A. (2023). Non-fungible tokens. In Walker, T., Nikbakht, E., and Kooli, M., editors, *The Fintech Revolution: How Financial Innovations Transform the Banking Industry*, chapter 6, pages 266–290. Springer International Publishing, Cham.
- Borri, N., Liu, Y., and Tsyvinski, A. (2022). The economics of non-fungible tokens. Working paper.
- Bosch-Rosa, C., Meissner, T., and Bosch-Domènech, A. (2018). Cognitive bubbles. *Experimental Economics*, 21(1):132–153.

- Brunnermeier, M. (2016). Bubbles. In Jones, G., editor, *Banking Crises*, chapter 5, pages 268–36. The New Palgrave Dictionary of Economics, New York.
- Brunnermeier, M. and Nagel, S. (2004). Hedge funds and the technology bubble. *Journal of Finance*, 59(5):2013–2040.
- Campbell, J. Y. (2006). Household finance. *Journal of Finance*, 61(4):1553–1604.
- Cong, L. W., Li, Y., and Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155.
- Corgnet, B., Hernán-González, R., Kujal, P., and Porter, D. (2015). The effect of earned versus house money on price bubble formation in experimental asset markets. *Review of finance*, 19(4):1455–1488.
- Creighton, J. (2022). Nft timeline: The beginnings and history of nfts.
- Dowling, M. (2022a). Fertile land: Pricing non-fungible tokens. *Finance Research Letters*, 44:102096.
- Dowling, M. (2022b). Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters*, 44:102097.
- Easley, D. and O’Hara, M. (1987). Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1):69–90.
- Farhi, E. and Tirole, J. (2011). Bubbly Liquidity. *The Review of Economic Studies*, 79(2):678–706.
- Goldberg, M., Kugler, P., and Schär, F. (2021). Land valuation in the metaverse: Location matters. Working paper.
- Greenwood, R. and Nagel, S. (2009). Inexperienced investors and bubbles. *Journal of Financial Economics*, 93(2):239–258.
- Greenwood, R., Shleifer, A., and You, Y. (2019). Bubbles for fama. *Journal of Financial Economics*, 131(1):20–43.
- Greenwood, R. and Thesmar, D. (2011). Stock price fragility. *Journal of Financial Economics*, 102(3):471–490.
- Hildobby (2022). Nft wash trading on ethereum.
- Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, 54(6):2143–2184.
- Hong, H. and Stein, J. C. (2007). Disagreement and the stock market. *Journal of Economic Perspectives*, 21(2):109–128.
- Ito, K., Shibano, K., and Mogi, G. (2022). Bubble prediction of non-fungible tokens: An empirical investigation.

- Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*. Palgrave Macmillan.
- Kindleberger, C. P. (1978). *Manias, Panics, and Crashes. A History of Financial Crisis*. Palgrave MacMillan.
- Kireyev, P. and Lin, R. (2021). Infinite but rare: Valuation and pricing in marketplaces for blockchain-based nonfungible tokens. Working paper.
- Kong, D.-R. and Lin, T.-C. (2021). Alternative investments in the fintech era: The risk and return of non-fungible token (nft). Working paper.
- La Morgia, M., Mei, A., Mongardini, A. M., and Nemmi, E. N. (2022). Nft wash trading in the ethereum blockchain. *arXiv preprint arXiv:2212.01225*.
- Michailova, J. and Schmidt, U. (2016). Overconfidence and bubbles in experimental asset markets. *Journal of Behavioral Finance*, 17(3):280–292.
- Nadini, M., Alessandretti, L., Giacinto, F. D., Martino, M., Aiello, L. M., and Baronchelli, A. (2022). Mapping the nft revolution: market trends, trade networks, and visual features. *Scientific Reports*, f22(11):102096.
- Nneji, O. (2015). Liquidity shocks and stock bubbles. *Journal of International Financial Markets, Institutions and Money*, 35:132–146.
- Ofek, E. and Richardson, M. (2003). Dotcom mania: The rise and fall of internet stock prices. *Journal of Finance*, 58(3):1113–1137.
- Oh, S., Rosen, S., and Zhang, A. L. (2022). Investor experience matters: Evidence from generative art collections on the blockchain. Working paper.
- Pastor, L. and Veronesi, P. (2009). Technological revolutions and stock prices. *American Economic Review*, 99(4):1451–83.
- Scheinkman, J. and Xiong, W. (2003). Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6):1183–1220.
- Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.
- Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151.
- Van Horne, J. C. (1985). Of financial innovations and excesses. *Journal of Finance*, 40(3):620–631.
- Victor, F. and Weintraud, A. M. (2021). Detecting and quantifying wash trading on decentralized cryptocurrency exchanges. In *Proceedings of the Web Conference 2021*, pages 23–32.

Yousaf, I. and Yarovaya, L. (2022). The relationship between trading volume, volatility and returns of non-fungible tokens: evidence from a quantile approach. *Finance Research Letters*, 50:103175.