**DEPARTMENT OF BUSINESS AND MANAGEMENT**

Course of Cybercrime and Fraud Detection

# Multi-Scale Network Analysis for Fraud Detection in NFT Markets

Prof. Gianluigi Me

SUPERVISOR

Elisa Presciutti

ID - 288611

ACCADEMIC YEAR 2024/2025

# ABSTRACT

This thesis investigates fraudulent behaviors in non-fungible token (NFT) markets through a multi-scale, graph-based forensic framework. Traditional approaches to fraud detection—focused on transaction anomalies or isolated assets—often fail to capture the systemic and coordinated nature of manipulation in these markets. To address this gap, the research develops and tests a framework that integrates micro-level (single asset), meso-level (wallet cluster), and macro-level (ecosystem-wide) analyses.

The project begins with the selection of a single NFT, Quirkie #3600, as a controlled entry point. A clean baseline network was constructed around this token, with ownership history verified to be free of suspicious cycles or anomalous trading. From this baseline, the network was expanded in two controlled stages. The first expansion reconstructed the complete NFT portfolios of Quirkie #3600's past and current owners, linking each token to its collection and normalizing identifiers to prevent duplication. The second expansion incorporated the most frequent trading partners of high-activity wallets within a 90-day temporal window, generating a heterogeneous graph of more than 1,465 entities—including wallets, NFTs, collections, and synthetic transactions—while keeping growth computationally tractable. On this expanded network, a set of structural and temporal detection modules was implemented.

Results show that while single-NFT trading histories often appear legitimate, systemic risks emerge only when ownership structures, temporal patterns, and community clusters are analyzed together. Centrality and community detection revealed the disproportionate influence of a few hub wallets, while fraud modules detected wash trading loops, circular trading cycles, and bot-like transaction timings. These findings confirm that manipulation in NFT markets is scale-dependent, becoming visible only when structural and temporal signals are aggregated.

Theoretically, the thesis contributes an operational taxonomy that translates abstract fraud categories—trading manipulation, price manipulation, social engineering, and technical exploits—into graph-theoretic detection logics. Practically, it demonstrates how graph algorithms, cycle detection, and similarity measures can support forensic auditing, regulatory oversight, and marketplace monitoring.

The study acknowledges methodological limitations, including incomplete price data, bounded network expansion, and the ethical risks of deanonymization and profiling. Nonetheless, it outlines clear avenues for future development: integrating money flow analysis with reliable pricing data, extending the framework across multiple blockchains, and incorporating real-time monitoring supported by AI-driven anomaly detection.

Overall, the research shows that NFT fraud is not random or isolated but systemic and multi-layered. By demonstrating the effectiveness of a multi-scale, network-centric approach, the thesis offers both theoretical insights and practical tools for enhancing transparency, accountability, and trust in digital asset markets.

# TABLE OF CONTENTS

# INTRODUCTION

## Context: Growth of the NFT Market and Emerging Risks

The origins of non-fungible tokens (NFTs) date back to early blockchain experiments in the mid-2010s. One of the first widely recognized examples was *Quantum* (2014), created by Kevin McCoy, which later sold at Sotheby's for over 1.4 million USD (Ruan, 2022). Early projects such as *Spells of Genesis* (2015) and *Rare Pepes* (2016) introduced NFTs into gaming and digital art, establishing the notion of verifiable digital scarcity. The real breakthrough came in 2017 with *CryptoPunks* and *CryptoKitties*, which popularized generative collections and blockchain-based gaming, paving the way for mainstream adoption (Wang et al., 2021).

The evolution of NFT markets reached a turning point in 2021, often described as the "NFT boom." Global trading volumes exceeded 25 billion USD, fueled by speculative investment, heightened media attention, and endorsements from celebrities and brands (Nadini et al., 2021). NFTs were increasingly framed not only as technological innovations but also as investment vehicles, capable of transforming digital ownership into a new asset class (Kim, 2023). This rapid expansion brought NFTs from niche blockchain communities into global financial and cultural discourse.

From an economic perspective, NFT markets established new forms of value creation. Artists and creators benefit from royalties automatically enforced through smart contracts, while platforms monetize transaction fees on primary and secondary sales (De Guzman, 2021). Yet, the market remains highly concentrated. OpenSea quickly emerged as the dominant marketplace, at one point accounting for the majority of Ethereum-based NFT trades (White et al., 2022).

Despite significant growth, digital asset markets remain characterized by high volatility and speculative behavior, including rapid turnover, price fluctuations, and cyclical patterns of expansion and contraction. These dynamics reflect both evolving valuation frameworks and structural conditions that may enable various forms of market manipulation, including pseudonymous participation, regulatory ambiguity, and fragmented oversight mechanisms. Certain structural characteristics further challenge market integrity: user anonymity can facilitate artificial trading practices, while regulatory uncertainties may leave platforms operating with limited oversight or consumer protections. Information disparities between different participant types can also influence market dynamics, with some actors potentially having advantages in accessing market data or understanding pricing mechanisms.

These documented cases of fraud, including counterfeit NFTs and volume manipulation on platforms like OpenSea and LooksRare, have resulted in significant financial losses and reputational damage, threatening the long-term sustainability of NFT markets (Elliptic, 2022; Chainalysis, 2023).

## Research Problem

Fraudulent practices have rapidly emerged as a defining challenge in the development of NFT markets. Among the most recurrent forms of misconduct are wash trading, where actors trade the same asset between controlled wallets to inflate prices and volumes; pump-and-dump schemes, involving coordinated buying followed by collective selling; and other types of price manipulation, such as volume inflation (Ruan, 2022; Sifat, 2024). In addition, the use of bots to automate bidding and trading, together with technical exploits exploiting vulnerabilities in smart contracts or metadata, highlight the diversity and sophistication of risks affecting this new asset class. These practices distort price discovery, mislead investors, and threaten the credibility of NFT marketplaces as a whole, with some studies estimating that wash trading accounts for up to 95% of transaction volume in certain collections (Von Wachter et al., 2022).

While these problems are widely recognized, the methods traditionally employed to detect them remain inadequate and often reactive rather than preventive. On the one hand, transaction-based approaches—such as monitoring price and volume anomalies—can identify unusual patterns but are easily circumvented by sophisticated actors capable of fragmenting trades or mimicking organic activity (Ruan, 2022). For instance, traditional volume-based alerts may flag a sudden spike in transactions, but fail to detect when the same entity controls multiple wallets executing coordinated trades over extended periods. On the other hand, platform-level surveillance often relies on rule-based systems or reactive alerts, which lack transparency and consistency across marketplaces (Sifat, 2024). Finally, analyses focusing on isolated NFTs fail to capture broader dynamics, such as coordinated activity across collections or wallet clusters. As a result, fraudulent behavior often escapes detection, eroding investor confidence and hindering the mainstream adoption of NFTs as reliable digital assets.

## Objectives, Contributions, and Thesis Structure

This thesis aims to explore a multi-scale investigative framework for fraud detection in NFT markets, combining micro-level signals from individual NFTs with macro-level network perspectives to highlight manipulation patterns that often remain invisible when examined in isolation. It is guided by two research questions: whether organic network discovery from a single NFT can reveal hidden manipulation patterns, and how fraud indicators differ between NFT-level and ecosystem-level analysis. The framework is implemented through graph-based techniques in Neo4j and tested on a real dataset, beginning with Quirkie #3600 and expanding into its surrounding network.

The research offers contributions at different levels: theoretically, by adapting existing classifications of manipulation patterns into an operational taxonomy suitable for the NFT context; methodologically, by experimenting with a multi-scale approach that integrates organic network discovery and address-only investigation; empirically, by showing how a single NFT case can expand into a network of 1,465 nodes and be analyzed across micro, meso, and macro scales.

The thesis is structured into four chapters: the first provides a review of the NFT market and the academic literature on manipulation and misuse, from which a comprehensive taxonomy of fraud patterns is developed; the second outlines the methodological framework, introducing graph-based analytical tools and detection logics; the third focuses on the empirical case study of Quirkie #3600, applying detection rules and graph algorithms to analyze this single NFT; and the fourth expands the network organically from this seed, and concludes with the discussion of results, implications, limitations, and directions for future research.

# CHAPTER 1: STRUCTURE AND MANIPULATION RISKS OF NFT MARKETS

## 1.1 Technical and Structural Evolution of NFT Markets as Enablers of Manipulative Behavior

The architecture of NFT markets is grounded in blockchain standards and smart contracts, which enable innovation while simultaneously introducing structural vulnerabilities exploitable for manipulation. The most widely adopted token standards, ERC-721 and ERC-1155, govern asset uniqueness, semi-fungibility, and ownership (Ko, 2023; Nadini et al., 2021). Through the minting process, NFTs are created, assigned to wallets, and programmed with features such as royalties. In principle, royalties can be enforced automatically via smart contracts (e.g., ERC-2981), yet enforcement varies across marketplaces, with some platforms bypassing or disabling these mechanisms (De Guzman, 2021; White et al., 2022).

Marketplace mechanics such as listing, bidding, and reselling have contributed to market growth but also expanded the attack surface for exploitation. Moreover, the reliance on off-chain metadata stored on centralized servers—while cost-efficient—creates vulnerabilities to alteration, loss, or manipulation (Nadini et al., 2021). Smart contracts themselves may also contain exploitable flaws, ranging from faulty minting logic to bugs in approvals and transfers.

The 2022 OpenSea exploit illustrates this vulnerability: when NFT owners tried to remove their listings without paying gas fees by transferring their NFTs to another wallet and back, the action removed listings from OpenSea's website but left old listings active in the blockchain contract. Attackers bought these NFTs at the old, lower prices and immediately resold them at current market prices, with one attacker making $200,000 from a single Bored Ape NFT. This case demonstrates how technical flaws between a platform's interface and blockchain can be systematically exploited for profit (Verma et al., 2024).

Finally, the pseudonymous nature of blockchain transactions makes it trivial for actors to operate multiple wallets, facilitating practices such as wash trading, circular trading, and price inflation schemes. Automated bots and opaque platform designs further enable malicious users to fabricate market activity and mislead participants about the true liquidity or value of an asset (Leppla et al., 2022; Von Wachter et al., 2022).

Beyond the technical infrastructure, NFT markets are shaped by diverse actors with distinct incentives that create systemic information asymmetries. Creators profit from royalties and visibility, while collectors are often motivated by cultural value or community belonging but remain disadvantaged by fragmented information. Speculative traders seek short-term arbitrage and flipping opportunities, often amplifying market volatility. Platforms such as OpenSea and Blur monetize fees and implement liquidity rewards that encourage aggressive trading cycles, with OpenSea historically controlling more than 80% of total trading volume, granting it privileged power over listing rules and visibility (Von Wachter et al., 2022).

These structural imbalances are reinforced by uneven access to market intelligence. Although blockchain data is theoretically transparent, interpretation and analysis capabilities vary dramatically across participants. Tools such as rarity rankings provide sophisticated traders with informational advantages that casual participants cannot exploit.

Empirical research shows that the introduction of rarity ranks reduced information asymmetry but paradoxically lowered transaction volumes and prices in low-rarity collections, since informed buyers revised valuations downward (Yuan et al., 2023).
Market concentration further amplifies these disparities, as evidenced by the 2022 insider trading case at OpenSea (Chainalysis, 2023).

Speculative dynamics worsen these structural risks through network effects and feedback loops. Social media platforms and online communities on Twitter and Discord amplify hype, drawing in new participants whose activity generates additional visibility and liquidity.

In sum, the rapid expansion of NFT markets has been fueled by technological innovation, platform incentives, and network effects. Yet these same forces—off-chain metadata reliance, pseudonymity, market concentration, and speculative amplification—have generated fragile structures where transparency does not equate to fairness. The interplay of technical flaws, informational asymmetries, and speculative dynamics has created an ecosystem in which manipulative behavior can thrive, setting the stage for the vulnerabilities examined in the following section.

## 1.2 Structural Vulnerabilities and Regulatory Gaps

A defining characteristic of NFT markets is the pseudonymous nature of blockchain transactions. While every transfer is permanently recorded on-chain, wallet addresses appear only as alphanumeric strings, providing no intrinsic link to real-world identities (Chen & Omote, 2022). This design lowers barriers to entry and enables global participation, but it also generates a structural opacity that can be strategically exploited. In particular, pseudonymity allows single actors to control multiple wallets simultaneously, creating the illusion of independent trading activity when transactions are, in fact, coordinated by the same entity (Niu et al., 2024).

This capacity underpins one of the most pervasive forms of manipulation in NFT markets: wash trading. Empirical evidence confirms the widespread nature of this practice: forensic analyses show that in certain collections, up to 95% of transaction volume is attributable to wash trading (von Wachter et al., 2022). A particularly illustrative case was the launch of LooksRare in 2022, where incentive mechanisms rewarding trading activity inadvertently encouraged large-scale manipulation. To maximize token rewards, participants created elaborate networks of wallets and executed self-trades that simulated liquidity but eroded trust in reported volumes. Automated bots further amplify these effects by executing high-frequency trades across wallets, masking manipulation under the guise of organic liquidity (Leppla et al., 2022).

The risks are compounded by the absence of robust Know-Your-Customer (KYC) requirements on most NFT marketplaces. Unlike regulated financial exchanges, platforms such as OpenSea or LooksRare rarely verify user identities, making it easy for a single actor to create dozens of wallets and simulate market activity (Ruan, 2022). As a matter of fact there are studies that emphasize the fact that the ease of generating pseudonymous addresses, combined with minimal entry costs, not only enables wash trading but also complicates detection efforts (Unveiling Wash Trading in Popular NFT Markets, 2023).

Beyond distorting individual prices, this manipulation undermines market credibility by eroding confidence in reported volumes, misleading participants about liquidity, and skewing valuation metrics built on transaction history (Sifat, 2024). For retail investors, the inability to distinguish genuine from fabricated activity raises entry barriers, while professional traders exploit asymmetries to their advantage. Over time, persistent opacity risks may deter institutional adoption and attracting stricter regulatory scrutiny.

The decentralized architecture of NFT markets adds further complications for oversight. Unlike traditional financial systems, where centralized intermediaries act as identifiable gatekeepers, NFT transactions are executed directly on blockchain networks via smart contracts, bypassing regulated intermediaries.

As a consequence, regulatory divergence has emerged across jurisdictions. In the United States, NFTs may fall under the remit of the SEC if classified as securities, or the CFTC if linked to derivatives, while many remain unregulated as digital collectibles. In the European Union, the Markets in Crypto-Assets Regulation (MiCA) provides a general framework for digital assets but only partially addresses NFTs, excluding those deemed "unique" unless they resemble financial instruments (Sulkis, 2024). Singapore and Hong Kong, meanwhile, emphasize anti-money laundering (AML) compliance but stop short of NFT-specific rules. This results in a patchwork of frameworks in which identical activities may be subject to different—and sometimes conflicting—treatments depending on jurisdiction (De Guzman, 2021).

Such fragmentation enables regulatory arbitrage, where platforms often incorporate in permissive jurisdictions to avoid stringent oversight, exploiting differences in securities law interpretation to minimize exposure (Zarifis & Castro, 2022). This "jurisdiction shopping" undermines the effectiveness of national enforcement and creates an uneven playing field, where compliance becomes a competitive disadvantage rather than a baseline standard. The global nature of blockchain worsens this challenge: regulators can target centralized exchanges, but enforcement remains restricted by national authority.

Transparency issues add a further dimension to these vulnerabilities. Metadata provides essential information—such as provenance, description, and associated media—yet the majority of it is stored off-chain, often on centralized servers or semi-decentralized systems like IPFS (Prakash et al., 2023). Centralized storage exposes NFTs to censorship, alteration, or outright disappearance, while even IPFS depends on voluntary data persistence and uneven accessibility (Bhanu Prakash et al., 2023). These weaknesses undermine the very immutability that NFTs are supposed to guarantee.

Another critical challenge is presented by cross-platform tracking. NFT trading is fragmented across marketplaces like OpenSea, Blur, and LooksRare, each with different listing systems, incentive models, and data aggregation practices. The absence of interoperability complicates ownership tracking and opens loopholes—such as outdated prices on one platform while listings remain active on another—exploited in past attacks (Verma et al., 2024). Moreover, competition for liquidity has led some platforms to overstate activity or inadequately filter suspicious trades, reducing the reliability of aggregated market statistics.

These transparency issues directly impair price discovery. Unlike fungible assets traded on centralized exchanges, NFTs lack standardized valuation mechanisms. Liquidity fragmentation creates asynchronous pricing, while royalties, speculative flipping, and thin trading volumes further distort floor prices. In low-volume collections, even a handful of trades can disproportionately shift benchmarks, making prices unreliable indicators of genuine market consensus (Kang & Lee, 2025).

In sum, pseudonymity, decentralized governance, and transparency limitations interact to create systemic vulnerabilities. Weak KYC enforcement and multi-wallet control enable manipulative behaviors; fragmented regulation fosters jurisdiction shopping and regulatory arbitrage; and reliance on off-chain metadata, cross-platform fragmentation, and unstable price discovery obscure reliable signals. Taken together, these dynamics generate an ecosystem where fraudulent practices are not isolated anomalies but systemic risks embedded in the market structure itself.

Building on these enabling conditions, the following section reviews the theoretical taxonomy of fraud types documented in the literature—ranging from trading and price manipulation to social engineering and technical exploits—providing the foundation for the detection-oriented operational taxonomy developed in Chapter 2.

## 1.3 Theoretical Taxonomy of NFT Fraud

The vulnerabilities outlined in the previous section manifest concretely in a range of fraudulent practices that have been documented across NFT marketplaces. Academic research and forensic investigations typically distinguish four broad categories of misconduct: trading manipulation, price manipulation, social engineering, and technical exploits. Taken together, these patterns form the backbone of fraudulent activity in NFT markets and provide the foundation for the taxonomy discussed in the following subsections.

### 1.3.1 Trading Manipulation

Among the different categories of fraudulent behavior in NFT markets, trading manipulation occupies a central position, as it directly undermines the mechanisms through which value is signaled and perceived. Building on the structural vulnerabilities discussed previously, this manipulation takes three primary forms: wash trading, circular trading, and coordinated flipping.

Wash trading is the action of artificially inflating trading activity by buying and selling the same NFT between wallets controlled by the same entity to create false demand signals. Beyond aggregate data, specific high-profile cases have brought public attention to the manipulative potential of wash trading. Perhaps the most famous is CryptoPunk #9998, which in October 2021 appeared to have sold for over $500 million. In reality, the trade was engineered by transferring the NFT between two wallets controlled by the same entity, using a self-financed loan to create the illusion of an unprecedented record sale.

Although the transaction was quickly reversed, the inflated sale momentarily distorted price benchmarks and received widespread media coverage, illustrating how wash trades can be weaponized to generate hype and legitimacy.

More subtle but equally impactful were wash trading cycles observed in mid-tier collections such as Meebits, where repetitive self-dealing created sustained illusions of market activity that proved attractive to less sophisticated investors.

Closely related to wash trading is circular trading, in which NFTs circulate within small, closed clusters of wallets in repeated loops. Unlike wash trades between two addresses, circular trading often involves a group of wallets collaborating to give the appearance of diverse buyer interest. For example, an NFT may be sold from Wallet A to Wallet B, then to Wallet C, and eventually back to Wallet A, all under the control of the same actor or colluding parties. This behavior not only inflates transaction counts but also creates misleading indicators of "unique buyers" and market popularity. Transaction graph analyses reveal these schemes as cyclical ownership patterns, where NFTs appear to change hands frequently but ultimate control remains unchanged. Cases of circular trading were particularly visible in incentive-driven markets in 2022, when clusters of addresses cycled assets at scale to farm liquidity rewards. Collections such as Meebits and Mutant Ape Yacht Club displayed abnormal looping turnover during this period, raising red flags in blockchain forensics studies (Unveiling Wash Trading, 2023).

A further variant of trading manipulation is coordinated flipping, which exploits the speculative nature of NFT markets. Flipping, in its legitimate form, refers to buying an NFT and reselling it quickly for profit, often based on short-term shifts in hype or visibility. However, when groups of actors synchronize their flips, the practice evolves into a manipulative strategy. Coordinated flipping typically involves concentrated bursts of buying, immediately followed by rapid reselling at incrementally higher prices, creating a short-term surge in perceived demand. These operations are amplified by social signaling—such as announcements in Discord channels or viral posts on Twitter—that attract uninformed buyers once the price momentum is visible. Research into opportunistic trading has shown that sophisticated wallets often initiate these cycles around events that boost visibility, such as platform front-page listings or influencer endorsements, extracting profits before late entrants realize the underlying demand is artificially manufactured (Exploiting Unfair Advantages, 2022).

Coordinated flipping displays unusually short holding times and synchronized listing behaviors. In certain PFP collections, NFTs change hands multiple times within hours, at incrementally higher prices. They are then sold to external buyers at peak valuations. While individual flipping is speculative, coordinated flipping becomes fraudulent when it deliberately misleads buyers. The impact is then significant: they distort floor prices, create artificial bubbles, and leave uninformed participants at disproportionate risk of losses.

Collectively, wash trading, circular trading, and coordinated flipping highlight the fragility of price discovery in NFT markets. By manufacturing apparent interest, simulating buyer diversity, and staging demand proxies, these practices compromise fair-value inference. Their prevalence, documented in both high-profile cases and large-scale forensic analyses, demonstrates that NFT markets remain highly susceptible to manipulation, with systemic implications for credibility and long-term sustainability.

### 1.3.2 Price Manipulation

Price manipulation in NFT markets differs from pure trading manipulation in that it primarily targets perceptions of value rather than transaction counts alone. Instead of fabricating trades between wallets, these schemes rely on hype cycles, distorted signals, and supply-side tactics to inflate perceived worth.

Pump-and-dump schemes are among the most visible examples. In this dynamic, organizers accumulate NFTs from a targeted collection, use social media channels such as Twitter, Telegram, or Discord to generate hype, and then liquidate their holdings once demand—and therefore prices—surge. Some NFT-specific cases have been documented. For instance, Chainalysis (2022) reported that small collections with thin liquidity are especially vulnerable: even a handful of coordinated buyers can trigger price spikes that attract unsuspecting retail participants, who are then left exposed when insiders exit. Studies of Telegram-based NFT groups show that campaigns often promise "exclusive alpha" on upcoming drops, only to collapse within days as organizers dump assets onto latecomers (Rajaei & Mahmoud, 2023).

A second form is market signaling manipulation, which operates not through self-trading but by broadcasting distorted indicators of demand. Rather than inflating volumes mechanically, manipulators use tactics such as coordinated floor-price bidding, sudden mass listings, or fake celebrity endorsements to suggest momentum. Elliptic (2022) documents how coordinated buying groups have temporarily raised the floor prices of niche collections, creating the illusion of organic growth. Once external buyers entered at inflated prices, insiders reversed their positions, pocketing profits

Finally, artificial scarcity is a supply-side form of price manipulation. While scarcity is inherent to NFTs, it becomes manipulative when creators or intermediaries deliberately restrict supply in misleading ways. Examples include releasing only a small fraction of a minted collection to create FOMO before flooding the market, or artificially burning assets to exaggerate rarity. As Nadini et al. (2021) note, such tactics may inflate demand in the short term but often undermine trust once buyers realize scarcity was engineered rather than natural. Crucially, artificial scarcity is problematic only when these practices are hidden or vaguely communicated; if supply strategies are disclosed transparently, participants can decide with full awareness, and the scarcity, though engineered, is not deceptive.

In sum, price manipulation in NFT markets encompasses demand-side schemes (pump-and-dumps), perception-based distortions (market signaling), and supply-side constraints (artificial scarcity). Unlike wash trading, which fabricates activity directly, these practices recalibrate.

### 1.3.3 Social Engineering

Social engineering in NFT markets refers to manipulative strategies that exploit the community-driven dimension of these ecosystems. Unlike technical exploits, these practices target trust, hype, and collective perception. By simulating legitimacy and popularity, malicious actors influence market behavior in ways that are difficult to trace yet highly effective.

Bot networks are one of the most common techniques. Automated accounts inflate activity by simulating bidding wars, spamming Discord chats, or amplifying visibility on Twitter.

The goal is to stage demand proxies that lead buyers to believe an asset is scarce or highly sought after. Studies show that such bots can distort both on-chain activity and social media metrics, making it difficult to separate organic liquidity from manufactured hype (Leppla et al., 2022).

Fake endorsements represent another powerful tool. Fraudsters impersonate well-known figures or exploit compromised accounts to give credibility to fraudulent collections. A notable case occurred in 2021, when an NFT purportedly linked to Banksy was promoted through the artist's hacked website and sold to collector Pranksy before being exposed as fraudulent. Similarly, the Big Daddy Ape Club scheme mimicked the branding of Bored Ape Yacht Club and boosted engagement with bots and fake influencer accounts, misleading thousands of buyers before disappearing. These cases highlight how quickly social proof, once manufactured, can drive investor decisions.

On a larger scale, entire coordinated campaigns are often orchestrated to build the illusion of authentic communities. Unlike pump-and-dump schemes, which are short-lived and price-focused, these campaigns emphasize long-term credibility. Fraudulent teams set up Discord servers with bots posing as active members, publish detailed roadmaps, and organize staged "Ask Me Anything" events. The Frosties NFT scam (2022) illustrates this dynamic: the project cultivated a large following with promises of a play-to-earn game and regular community engagement, raising over $1.3 million before the founders abandoned it in a rug pull. Here the manipulation relied less on immediate price inflation and more on sustaining the illusion of a committed community.

The consequences of these practices are significant. They corrode trust in community-based projects, making it harder for genuine creators to attract long-term participants. Retail buyers, often guided by often driven by market excitement and social validation and social validation, are disproportionately exposed, while professional traders may capitalize on or even orchestrate these manipulations. Moreover, regulation struggles to address social engineering: unlike technical exploits, these tactics operate in a grey area between aggressive marketing and outright fraud, leaving little clear evidence for enforcement.

In sum, social engineering in NFT markets includes bot networks that inflate apparent demand, fake endorsements that hijack trust in public figures, and coordinated campaigns that fabricate entire communities. Unlike pump-and-dump events, which are transient and price-centered, these campaigns simulate legitimacy over time, producing artificial trust that either precedes or conceals later manipulation of value.

### 1.3.4 Technical Exploits

A distinct category of fraudulent practices in NFT markets arises from technical exploits, which target flaws in the code or infrastructure rather than relying on market behavior or social deception. These incidents expose how weaknesses in smart contracts or metadata storage can directly compromise ownership, security, and the long-term reliability of NFTs.

Smart contract vulnerabilities represent some of the most visible technical risks. Since NFTs are governed by self-executing smart contracts, any flaw in their design or implementation can be exploited to bypass permissions or hijack assets. A well-documented case occurred in April 2022 on Rarible, one of the largest NFT marketplaces.

Researchers at Check Point discovered that attackers could upload NFTs containing hidden malicious code. When users clicked on these assets within the marketplace, the code was automatically activated, creating fake approval requests. To victims, these looked like normal marketplace prompts, but by clicking "approve" they unknowingly gave attackers complete access to their digital wallets and all their NFTs.The most notorious incident was the theft of the Bored Ape NFT #3738 belonging to singer Jay Chou, later resold for around $500,000. The exploit worked not because of flaws in the blockchain itself, but because Rarible's web interface failed to validate approval requests, turning a trusted environment into a vector for attack. Following responsible disclosure in early April, Rarible fixed the issue, but the case illustrated how even established platforms with millions of users can expose participants to devastating losses through overlooked technical weaknesses.

Metadata manipulation constitutes another critical avenue of technical exploitation. While ownership records are stored permanently on the blockchain, the actual content of NFTs—such as titles, images, and associated files—is often stored on separate servers. This creates a weakness: if a server goes offline, or if a file is deliberately changed, the NFT may lose its intended appearance or meaning, even though the ownership record remains valid. Studies have shown that a significant proportion of NFTs rely on centralized storage, leaving them vulnerable to disappearance, unauthorized modification, or censorship. Even with distributed storage systems, content preservation depends on voluntary participation, meaning that unpopular or abandoned assets can effectively vanish once servers stop hosting the files.

Real-world cases highlight the impact of this fragility. Some projects have engaged in so-called "metadata rug pulls", where collections initially displayed polished previews but were later replaced with crude or irrelevant images after the sale, leaving buyers with tokens stripped of their promised value. In other cases, server failures have caused assets to become inaccessible altogether, effectively voiding the practical ownership of the NFT. These practices reveal that the cultural and financial value of NFTs is only as secure as the infrastructure hosting their metadata—a dimension often invisible to retail buyers.

In summary, smart contract exploits and metadata manipulation demonstrate how technical weaknesses undermine the very mechanisms that are supposed to guarantee trust in NFT markets. Unlike wash trading or pump-and-dump schemes, these attacks weaponize flaws in code or infrastructure, bypassing economic signals and striking directly at the integrity of ownership. The immutability of the ledger preserves transaction traceability but not the permanence or authenticity of referenced content—an immutability paradox that fraudsters can exploit.

The review reveals that NFT fraud spans four main categories: trading manipulation, price manipulation, social engineering, and technical exploits. Each exploits distinct structural weaknesses, demonstrating that NFT markets are vulnerable to both opportunistic behaviors and systematic schemes. These categories often interconnect in practice. Wash trading combines with social media campaigns to support pump-and-dump operations, while technical exploits may be preceded by staged community engagement. Fraudulent strategies are rarely isolated but part of hybrid schemes that maximize deceptive impact.

Retail participants face disproportionate exposure. While professional traders can identify manipulations, less sophisticated buyers rely on signals—trending lists, floor prices, influencer endorsements—that fraudulent practices systematically distort. This erodes trust and limits institutional adoption. The diversity of these fraud patterns necessitates a structured taxonomy serving both descriptive and operational purposes. By categorizing misconduct into four dimensions, researchers can map specific signals to appropriate analytical tools.

The next chapter transitions this descriptive taxonomy into an operational framework, shifting from cataloguing fraud types to developing detection instruments—supporting the thesis's central aim of systematic misconduct detection and mitigation.

# CHAPTER 2: ANALYTICAL FOUNDATIONS FOR NFT FORENSICS

## 2.1 Blockchain Forensics

### 2.1.1 Address Clustering and Heuristics in Bitcoin Forensics

As already stated, a central challenge in blockchain forensics is the pseudonymous nature of cryptocurrencies. While every Bitcoin or Ethereum transaction is publicly visible on a blockchain, the entities controlling individual addresses are not directly identifiable. To overcome this opacity, investigators have developed systematic approaches for linking addresses and detecting suspicious patterns, forming the methodological foundation for modern blockchain forensics.

The cornerstone of blockchain forensics is address clustering—grouping addresses likely controlled by the same real-world entity. Early studies by Ron and Shamir (2013) and Meiklejohn et al. (2013) introduced practical heuristics that remain widely used. The multi-input heuristic infers that addresses spending together in a single transaction belong to the same user, since coordinated control of multiple private keys is required. The change-address heuristic exploits the technical design of Bitcoin: when transactions generate "change" outputs, these can often be identified and linked back to the original sender.

These clustering methods have proven effective in major investigations, such as the Silk Road case, where linking Bitcoin addresses enabled authorities to trace illicit proceeds. However, their reliability depends on assumptions that sophisticated users can deliberately break. Privacy-enhancing mechanisms like CoinJoin invalidate multi-input heuristics by mixing inputs from multiple users, while privacy-focused cryptocurrencies prevent clustering entirely (Moser et al., 2017; Böhme et al., 2015).

Complementing clustering techniques, forensic analysis relies on volume and temporal pattern detection to identify anomalies. Commercial forensic tools integrate anomaly detection algorithms that assess whether transaction flows deviate significantly from historical baselines, providing investigators with leads for deeper analysis.

The evolution of forensic methods reflects an ongoing competitition between investigators and malicious actors. As traditional heuristics become less reliable due to privacy enhancements and sophisticated obfuscation, the field has increasingly turned toward graph-based approaches that model the full network structure of transactions rather than isolated patterns. In the NFT context, where manipulation often involves coordinated activity across multiple wallets and collections, these network-centric methods offer particular promise for detecting systematic fraud.

### 2.1.2 Limits of NFT Analysis

Extending blockchain forensics from fungible cryptocurrencies to NFTs introduces unique characteristics that weaken or even break classical methods. Three elements in particular—non-fungibility, reliance on off-chain metadata, and partial transparency—create an environment where traditional heuristics become less effective.
Non-fungibility is the most fundamental difference. In Bitcoin, all units of the same denomination are interchangeable, which makes it possible to study systemic patterns by combining transaction data across thousands of addresses.

NFTs, by contrast, are fundamentally different from each other: each token has a unique identifier and represents a distinct digital asset, even within the same collection. This uniqueness makes it difficult to generalize findings or apply large-scale statistical analysis. As Nadini et al. (2021) emphasize, the movement of each NFT must be interpreted individually. Consequently, patterns detectable in fungible systems—such as layering in money laundering—are more opaque in NFT contexts, where each transfer could plausibly be explained by unique trading patterns rather than coordinated manipulation.

A second limitation stems from NFTs' reliance on off-chain metadata. The fact that the descriptive and visual components of a token are almost always hosted somewhere different than on the blockchain, for example on centralized servers or semi-decentralized systems, creates opacity: the blockchain proves that ownership was transferred, but it offers no guarantee of the permanence or authenticity of the associated asset.

The third limitation is the partial transparency of NFT markets. Although blockchains guarantee that every transfer is visible, understanding what's happening becomes difficult due to low trading activity, varied pricing, and the ability of single actors to operate multiple wallets. In active cryptocurrency markets, unusual transaction patterns can reliably signal illegal activity. In NFTs, however, low trading volumes mean that just a handful of transactions can distort key measures such as floor prices. When one person controls multiple wallets, it creates the false appearance of many different buyers, making transaction counts or holder diversity unreliable indicators of genuine demand (von Wachter et al., 2022).

These characteristics show why forensic methods developed for regular cryptocurrencies cannot be directly applied to NFT markets without changes. In practice, NFT investigations increasingly rely on combined approaches that merge blockchain data with external signals such as marketplace listings, metadata verification, and social media analysis. NFTs do not provide the same level of investigative clarity as regular cryptocurrencies, and their unique features create blind spots that malicious actors can exploit.

The evolution from heuristic clustering in Bitcoin to the unique challenges posed by NFTs illustrates both the progress and the limits of traditional blockchain forensics. While early methods provided powerful tools for de-anonymizing transactions, NFTs' structural features—non-fungibility, off-chain metadata dependency, and low-liquidity environments—undermine their effectiveness. This shift points toward the importance of graph-based analysis, which models wallets, tokens, and transactions as interconnectednodes and edges. By focusing on structural interaction patterns rather than isolated metrics, network analysis enables the detection of clusters, loops, and behavioral similarities that would otherwise remain invisible. The next section explores these techniques in detail, outlining how network analysis provides the foundation for detecting fraudulent patterns in NFT markets.

## 2.2 Graph-based analysis

### 2.2.1 Basic Concepts of Network Analysis

Network analysis offers a powerful perspective for studying blockchain systems because it shifts attention from individual transactions to the broader structures they form. A network is composed of nodes, which can represent wallets, contracts, or tokens, and edges, which capture relationships such as transfers, bids, or sales.

In traditional finance, networks have long been used to trace flows between accounts or institutions; in blockchain forensics, this approach is particularly effective since transaction records are public and can be systematically translated into graphs.

In cryptocurrency networks like Bitcoin, nodes are typically addresses or clusters of addresses, with edges denoting transfers of value. In Ethereum and NFT markets, however, the network perspective is richer: nodes may correspond to wallets, contracts, or individual NFTs, and edges capture a wide range of interactions. For example, if Wallet A sells an NFT to Wallet B for 5 ETH, this can be represented as a directed edge from A to B; if the token later returns to A, the cycle reveals a closed loop of ownership.

This distinction between linear and structural analysis is central to NFT forensics. A linear approach focuses on individual transfers—who paid whom, when, and how much—while structural analysis reveals higher-order patterns that emerge from multiple interactions. By moving from isolated exchanges to graph-based representations, investigators gain a systematic toolset for identifying manipulative behaviors and detecting suspicious wallet clusters.

### 2.2.2 Centrality Metrics

One of the most widely used approaches in network analysis is the study of centrality, which seeks to identify the most "important" or influential nodes in a network. Each centrality measure captures a distinct dimension of influence or structural position, valuable for detecting wallets that play unusual or disproportionately impactful roles in trading networks.

Degree centrality, the simplest measure, counts the number of direct connections a node has. In NFT markets, wallets with unusually high degree values often correspond to highly active traders who buy and sell across multiple collections. While some may be legitimate collectors or arbitrageurs, forensic studies show that wallets engaged in wash trading frequently appear as hubs with abnormally high degree centrality, since they orchestrate rapid back-and-forth exchanges to fabricate liquidity.

Betweenness centrality captures a different aspect by measuring how often a node lies on the shortest paths between others. Wallets with high betweenness act as intermediaries, controlling flows of assets or value across the network. In NFT trading graphs, these often correspond to addresses that serve as “bridges” between collusive clusters and external participants. For instance, in coordinated flipping schemes, certain wallets purchase NFTs from insiders at elevated prices and quickly resell them to outsiders, channeling profit from the manipulation. These addresses may not have the highest number of direct connections, but their brokerage role makes them critical to the operation.

Closeness centrality measures how quickly a node can reach all others in the network. In financial systems, this is often associated with efficient access to liquidity. Within NFT markets, wallets with high closeness can interact with many others through only a few steps, sometimes spanning across collections. In legitimate contexts, such wallets may belong to aggregators or liquidity providers. In forensic settings, however, high closeness values can reveal actors strategically positioned to exploit arbitrage opportunities or to spread manipulative behavior across multiple projects.

PageRank, originally developed to rank web pages, assigns more weight to links with already important nodes. This helps investigators highlight wallets disproportionately tied to "prestigious" collections or influential actors. In NFT markets, this can reveal addresses engaged in artificial boosting of high-prestige collections, where ties to highly ranked nodes are manufactured rather than organically accumulated. PageRank is therefore useful in distinguishing genuine high-value collectors from actors attempting to simulate prestige.

Applied to NFT markets, these measures expose structural asymmetries between ordinary participants and manipulative actors, revealing wallets that act as hubs, brokers, or amplifiers of artificial trends.

### 2.2.3 Community Detection

While centrality metrics highlight the importance of individual wallets, another crucial question in blockchain forensics is how wallets organize into groups or clusters. This is where community detection becomes central. In graph theory, communities are subsets of nodes that are more tightly connected to each other than to the rest of the network. Identifying these groups is particularly important in NFT forensics, since many manipulative behaviors—especially wash trading and collusive schemes—are rarely carried out by a single wallet, but instead by coordinated clusters.

The Louvain method is the most widely used community detection algorithm in blockchain analysis. It works by maximizing modularity, a measure of how well a network is divided into communities compared to a random distribution of edges. In practice, this means the algorithm highlights clusters where wallets trade more frequently among themselves than with the broader marketplace.

In NFT markets, this technique has revealed manipulation patterns where wallets formed closed groups and engaged in repetitive transactions to inflate trading volumes or simulate demand. For example, during reward-driven phases on certain marketplaces, clusters of wallets interacted almost exclusively with each other to farm incentives, creating the illusion of high liquidity. In other cases, like trading observed in mid-sized collections, community detection isolated groups of addresses that circulated NFTs within their cluster but rarely interacted with outsiders—a strong indication of fabricated liquidity rather than genuine demand.

Suppose Wallets A, B, and C trade the same NFT back and forth in a tight loop, while only rarely engaging with external wallets. When this network is processed with the Louvain algorithm, A, B, and C would be grouped into a single high-modularity community. To an investigator, this signals a suspicious collusive cluster: the wallets' activity is abnormally inward-focused, consistent with wash trading, rather than resembling legitimate market participation, where collectors trade with a more diverse pool of actors.

Beyond wash trading, community detection can also reveal pump-and-dump groups. In network terms, this produces a dense internal cluster connected to the broader market by just a few bridge wallets—the intermediaries who channel fabricated demand into real profit.

The value of community detection in NFT forensics lies in its ability to scale from individual suspicion to systemic patterns, highlighting collective behaviors that deviate from organic market structures.

### 2.2.4 Cycle Detection

Cycle detection focuses on identifying loops of repeated transactions among a small set of wallets. In graph theory, a cycle occurs when a token begins with one wallet and, after passing through others, eventually returns to its starting point. These loops are significant because they provide a structural footprint of circular trading—a manipulative practice where assets circulate within a closed group to simulate liquidity or rising demand.

For example, consider the sequence: Wallet A → Wallet B → Wallet C → Wallet A. On the surface, it appears that three independent buyers are exchanging an NFT, while in reality all wallets may be controlled by the same actor. Longer loops, such as Wallet A → Wallet B → Wallet C → Wallet D → Wallet A, follow the same principle: the token ends up where it started, without genuine ownership change.

The forensic value lies in separating normal reselling from artificial repetition. In healthy markets, NFTs may change hands often, but transactions involve diverse participants. Collusive cycles, by contrast, exhibit a closed pattern restricted to the same wallets. By highlighting these short, closed loops, cycle detection provides investigators with clear structural evidence of collusion, even in markets with low or fragmented trading volumes.

Studies have documented such patterns in various NFT collections, showing how short loops were used to inflate trading activity and manipulate floor prices (von Wachter et al., 2022). It captures micro-level manipulation patterns that complement broader community detection, making it one of the clearest tools for identifying circular trading in NFT ecosystems.

### 2.2.5 Similarity Measures

Similarity measures quantify how closely the behavior of two wallets resembles one another by examining overlap in transaction partners, timing, or asset flows. This makes them especially useful for identifying wallets that operate in tandem, either because they are controlled by the same actor or because they are part of a coordinated scheme.

The Jaccard index measures the degree of overlap between two sets. In blockchain analysis, these sets might correspond to the counterparties with whom two wallets have traded, or the collections in which they have been active. A Jaccard score of 1 indicates perfect overlap (all connections in common), while 0 means no overlap at all. For instance, if two wallets both trade with the same three wallets in a short period, their high Jaccard similarity suggests coordination rather than coincidence.

Cosine similarity compares the proportionality of two transaction histories. Each wallet's activity can be represented as a vector, where dimensions capture interactions such as the number of trades with a specific counterparty or the volume of activity in a collection. By measuring the angle between these vectors, cosine similarity captures not just overlap but also how balanced the two patterns are. For example, two wallets that both trade heavily in the same collection will show higher similarity than two wallets that overlap only occasionally.

Applied to NFTs, similarity measures have clear forensic value. Wash traders frequently operate multiple wallets that mirror one another. Groups preparing pump-and-dump schemes often show unusually high similarity in their buying phase, as they simultaneously accumulate the same assets before exiting together.

The strength of similarity measures lies in their ability to formalize "patterned behavior." Instead of requiring direct proof of linkage, they provide quantitative evidence that wallets are acting "too similarly" to be independent, uncovering collusive substructures hidden within heterogeneous NFT trading.

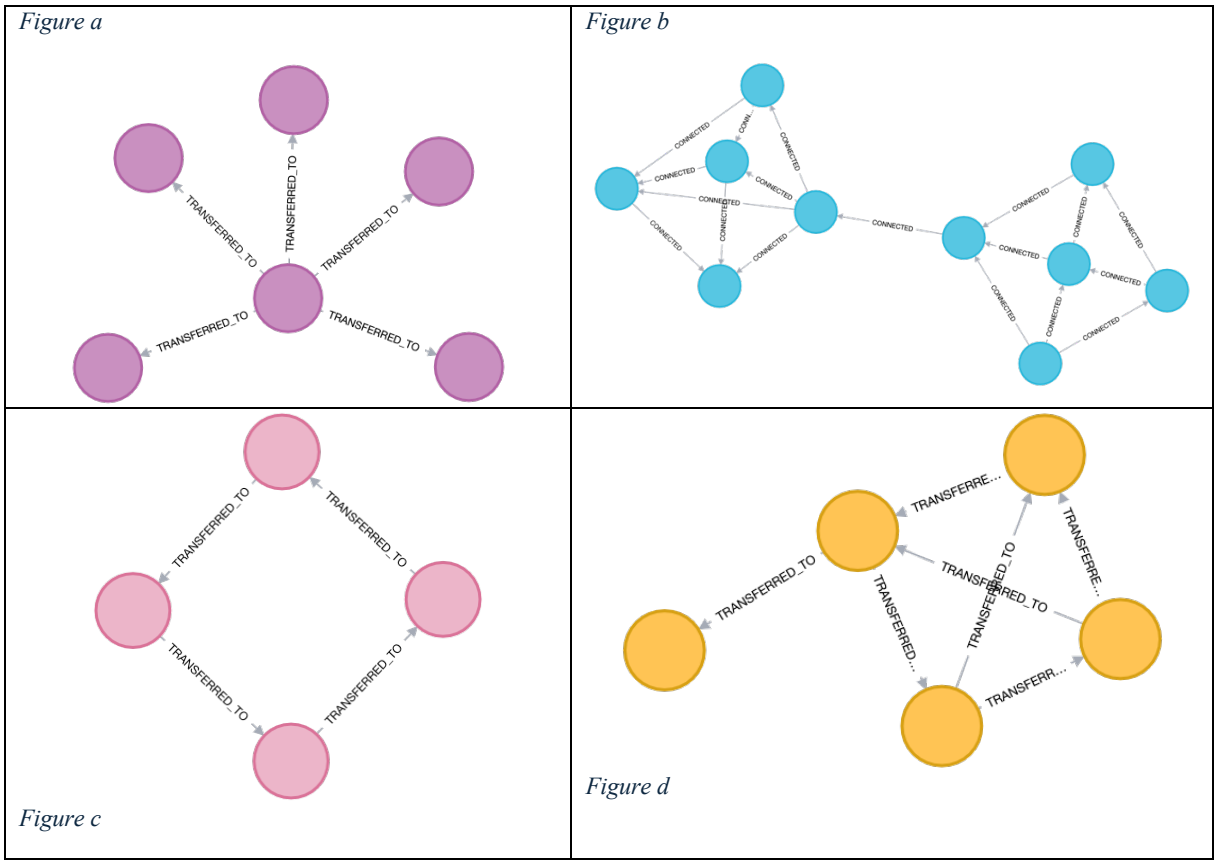


*Illustrative graph patterns in NFT forensics: (a) hub wallet, (b) bridge nodes connecting clusters, (c) closed loop representing circular trading, and (d) collusive cluster indicative of coordinated manipulation.*

## 2.3 Cybersecurity and AML

One of the most consolidated successes of blockchain forensics has been the use of graph analysis to detect money laundering and illicit financial flows. Early research showed how network-based approaches could uncover entities linked to darknet markets and mixing services (Meiklejohn et al., 2013; Ron & Shamir, 2013). By applying structural analysis to transaction networks, investigators could move beyond isolated transactions to reveal broader criminal ecosystems. In practice, these methods have been used to identify key intermediaries within ransomware networks, such as those linked to the WannaCry outbreak, where ransom payments were traced through complex laundering chains before attempts at cash-out (Chainalysis, 2018).

The success of these techniques lies in their ability to exploit the fungibility and scale of cryptocurrencies. Cryptocurrencies are divisible, interchangeable, and generate vast datasets where statistical irregularities stand out clearly. Although challenges have emerged with privacy-oriented coins such as Monero and Zcash, or through increasingly complex obfuscation strategies like cross-chain swaps (Böhme et al., 2015), graph-based approaches succeed by capturing structural patterns of behavior rather than relying solely on individual transactions.

This principle becomes complex when moving into the NFT context, where the forensic challenge extends beyond tracing illicit value flows to identifying manipulative behaviors that distort perceptions of demand and market value. While anti–money laundering (AML) forensics in cryptocurrencies has focused on hiding the origin and destination of large-scale flows, NFT manipulation follows a different logic. Instead of fungible units exchanged at massive scales, NFTs involve unique assets often manipulated at the level of individual tokens or small clusters. The goals also diverge: laundering seeks to remain invisible, while NFT manipulation often aims to generate visibility, fabricate confidence, and attract external buyers through the illusion of liquidity.

These differences expose a clear methodological gap. Classical AML approaches, optimized for fungible flows, are poorly suited to detecting manipulation patterns rooted in token uniqueness and market signaling. There is no comprehensive multi-scale framework that connects behavior at the level of a single asset to coordinated wallet clusters and, ultimately, to ecosystem-wide distortions. Moreover, NFT forensics requires integration of on-chain evidence with off-chain signals such as metadata integrity, marketplace activity, and even social media promotion.

## 2.4 Operational Taxonomy for NFT Fraud Detection

This section presents an operational taxonomy that translates the fraud patterns from Chapter 1.3 into detectable network signatures and quantifiable analytical methods. While the previous chapter documented fraudulent behaviors through descriptive cases and empirical examples, the following taxonomy shifts focus to the structural and temporal patterns that make these behaviors identifiable within transaction networks.

Each category specifies not only the conceptual nature of the manipulation but also the graph-theoretic markers, centrality anomalies, and clustering patterns that enable systematic detection.

This operational framework serves as the bridge between theoretical understanding of NFT fraud and the practical implementation of detection algorithms, providing investigators with concrete analytical tools for identifying manipulative activity across the four primary categories of misconduct.

### 2.4.1 Trading Manipulation

In NFT markets, trading manipulation designates practices that artificially inflate transaction activity to create misleading signals of liquidity, demand, or popularity. Unlike price manipulation, which primarily targets valuation metrics, this category operates directly at the transactional layer by simulating market momentum through orchestrated transfers between wallets under common control or collusion. Its main manifestations—wash trading, circular trading, and coordinated flipping—are central to the erosion of market integrity.

Patterns of abuse typically surface in repetitive transactions between the same wallets, short cycles of ownership where tokens return to the original holder, or sudden bursts of activity in otherwise illiquid collections. These signals, combined with abnormally short holding times and minimal counterparty diversity, diverge sharply from the irregular, dispersed behavior associated with genuine collectors.

Network analysis provides the most effective means of capturing these anomalies: wash trading materializes as recurring bidirectional edges between wallets over compressed intervals, while circular trading appears as closed loops of three to five nodes in which control of the asset never truly changes. At a higher level, collusive clusters are uncovered when wallets concentrate activity within dense subgraphs rather than engaging broadly with the market, and coordinated flipping emerges from the interplay between temporal markers and structural ones (small groups repeatedly reselling).

These signals reveal how manipulation can be distinguished from speculation by analyzing structural motifs and temporal rhythms rather than isolated trades, exposing the underlying coordination that sustains artificial activity in NFT ecosystems.

### 2.4.2 Price Manipulation

Price manipulation changes how value is perceived rather than inventing trading activity. It targets the signals traders rely on—price movements, scarcity, and demand—and often takes the form of pump-and-dump schemes, false market signaling, or artificial scarcity.

In graph terms, pump-and-dump schemes look like star-shaped bursts in ownership networks. At first, a small group of wallets buys many tokens, which appear as many incoming edges into a tight cluster. When the “dump” happens, this cluster suddenly sends tokens out to many different wallets, creating a wave of outgoing edges. Detecting this involves looking for sharp increases in a wallet’s degree centrality and star-like motifs. A suspicious case can be defined as when edge counts grow more than five times the usual baseline, with most transfers (for example, 80%) happening in just a few hours.

Market signaling appears in bidding graphs. For a short time, many wallets place bids on the same collection, producing a sudden increase in edges around one node.

The pattern then quickly disappears. Detection here means checking if edge density triples compared to the past 24 hours and then collapses. Short-lived clusters can be spotted using temporal community detection methods such as sliding-window Louvain.

Artificial scarcity leaves traces in ownership graphs. A small number of wallets may hold an unusually large share of tokens and then release them all at once. This can be measured with concentration metrics like the Gini coefficient. A warning sign is when the top five wallets control more than 40% of a collection, followed by a redistribution where over 70% of supply is sold within a single day. Such patterns indicate deliberate withholding and staged release.

The real difficulty is telling manipulation apart from genuine market shifts. Organic redistribution usually happens gradually and often follows visible events like new project announcements or market growth. Manipulation, by contrast, tends to appear suddenly and without clear outside causes. To improve detection, analysts can cross-check network signals with off-chain information such as marketplace activity or social media.

By turning these behaviors into measurable criteria—like spikes in edges, compressed time windows, and unusual ownership concentration—price manipulation becomes easier to detect. Graph analysis not only shows the shape of these behaviors but also allows them to be tested and verified systematically, making manipulation distinct from normal speculation.

### 2.4.3 Social Engineering

Although social engineering starts outside the blockchain, its effects reshape transaction networks in ways that can be measured. The graph often shows patterns that differ from organic collector activity, making manipulation visible.

When many wallets join at once and interact with the same project wallet, the graph takes on a star shape: one central node connected to many new peripheral nodes. This can be measured by a sudden spike in the degree centrality of the central wallet and by an abnormal ratio of one-to-many edges within a short time window.

Bot-driven campaigns produce radial clusters of small, low-value transactions. In the graph, these appear as several nodes with very low overall degree, all connected to the same project node. Detection relies on filtering wallets with limited transaction history and testing whether their interactions are concentrated in a narrow timeframe.

Fake endorsements create sudden centrality jumps. Wallets tied to the endorsed collection move from low to unusually high degree or betweenness centrality in a very short period, without the gradual increase expected from genuine adoption.
Analysts can measure this by monitoring relative changes in centrality over sliding windows and flagging outliers where growth is abrupt and unsupported by broader activity.

Coordinated campaigns leave dense subgraphs where wallets trade heavily with each other and with the project at nearly the same time. Graphically, this looks like a compact community with strong internal connectivity but weak links to the wider market. Community detection algorithms such as Louvain can isolate these clusters, and temporal analysis can confirm their synchronization by showing that a large share of their edges occur within the same short interval.

In short, social engineering becomes visible through recognizable graph motifs—stars, radial bot networks, and dense synchronized clusters. Each can be captured with standard network measures: centrality for stars, degree distributions for bots, and modularity with temporal compression for coordinated groups. These features turn off-chain deception into measurable, on-chain anomalies.

### 2.4.4 Technical Exploits

Technical exploits target weaknesses in the infrastructure that underpins NFTs rather than the social or transactional layers of the market. The two primary vectors are vulnerabilities in smart contracts and manipulation of off-chain metadata, both of which compromise the very mechanisms meant to guarantee ownership and authenticity.

Exploited contracts often appear in transaction graphs as nodes with abnormally high out-degree within compressed intervals, reflecting mass outflows of NFTs to attacker-controlled wallets. Transfers may occur without the corresponding payment edges, creating "value gaps" inconsistent with standard market behavior. Metadata manipulation, though occurring off-chain, can be inferred through sudden structural shifts: multiple tokens linked to altered metadata sources, or abrupt updates producing inconsistencies between token IDs and their descriptive attributes.

These anomalies highlight a fundamental difference from other categories. The issue is not inflated demand or deceptive perception but structural breaks in the integrity of transactions and asset representation. Detection requires combining anomaly analysis with external verification of metadata and contract integrity, revealing cases where the foundational promises of NFTs—immutability and authenticity—are directly undermined.

## 2.5 Multi-Scale Framework for NFT Fraud Detection

The detection of fraud in NFT markets cannot rely on a single analytical level. Traditional forensic methods often focus on either individual transactions or aggregate flows, but manipulative strategies in NFTs operate simultaneously at multiple scales. A wash trade may be visible at the level of a single token, yet only becomes suspicious when contextualized within a cluster of wallets or when seen as part of a wider marketplace campaign. To capture this layered complexity, a multi-scale framework is required—one that moves from the micro perspective of single-asset analysis, through meso-level wallet networks, to macro-level ecosystem dynamics, integrating findings into a coherent operational pipeline.

### 2.5.1 Micro Level – Single Asset Analysis

At the micro level, analysis focuses on individual NFTs and their immediate transaction histories. This scale is essential for detecting manipulative practices that target specific tokens, often with the goal of inflating their visibility or value how long wallets hold tokens, how often they trade, and sudden price changes. Short and repetitive ownership cycles are particularly informative: an NFT that returns to its original owner within a handful of transactions raises a strong signal of wash trading or circular trading. Similarly, abnormally short holding times compared to the collection average suggest speculative flipping designed to create artificial momentum.

Cycle detection provides the primary tool at this level. By scanning transaction paths for short closed loops, investigators can expose cases where the apparent diversity of buyers conceals repeated transfers between linked wallets. Temporal anomaly detection complements this by highlighting sudden bursts of trading activity that break with the asset's prior history or with baseline collection-level patterns.

Individual NFTs receive a risk score that integrates structural anomalies (loops, repeated counterparties) with temporal ones (unusual bursts, rapid flips), providing a quantifiable measure of suspicion for single tokens. While not sufficient to prove manipulation on its own, this score flags assets that warrant deeper investigation within wallet clusters or cross-collection contexts.

### 2.5.2 Meso Level – Wallet Cluster Analysis

The meso level shifts focus from individual tokens to the wallets transacting them. Fraud in NFT markets rarely occurs in isolation; it is typically coordinated by groups of addresses acting in concert. Detecting these collusive structures requires tools that capture relational and structural features of wallet networks.

Community detection algorithms, such as Louvain modularity optimization, reveal dense subgraphs where wallets transact disproportionately with each other while limiting interaction with the broader marketplace. Centrality measures further refine this picture: addresses with unusually high betweenness centrality may serve as brokers that channel manipulated assets toward unsuspecting buyers, while high degree hubs can indicate wallets orchestrating repetitive trades. Similarity metrics, including Jaccard and cosine measures, capture behavioral resemblance across wallets, exposing cases where multiple addresses mirror each other's trading profiles—an indicator of sybil wallets or coordinated flipping groups.

Analysis at this level centers on identifying collusive clusters. These may take the form of tightly knit rings engaged in wash trading, groups staging synchronized flips of a collection, or radial structures where multiple low-activity wallets funnel trades into a central operator. By mapping suspicious groups, investigators can move beyond individual trades to uncover organized manipulation, distinguishing isolated speculation from systemic collusion.

### 2.5.3 Macro Level – Cross-Collection Ecosystem Analysis

At the macro scale, the focus broadens to patterns spanning entire collections and marketplaces. Many manipulative strategies aim not only to distort individual token prices or collection floors but to shape ecosystem-wide perceptions of activity and legitimacy. This requires tracking correlations and anomalies across projects, identifying when unusual patterns spread across multiple collections.

Metrics at this level include cross-collection trading correlations, suspicious timing of volume increases across separate collections and platform-wide concentration of activity among a handful of wallet clusters. For example, a sudden spike in trading across multiple low-liquidity collections, all involving overlapping wallets, may indicate a coordinated campaign to inflate marketplace volume or to exploit incentive mechanisms.

Analysis integrates structural and temporal perspectives. At the structural level, overlapping communities of wallets that span multiple collections signal attempts to manufacture credibility across projects. At the temporal level, synchronized bursts of trading across collections suggest orchestrated strategies rather than independent market dynamics.

Ecosystem health indicators include the proportion of marketplace volume attributable to identified collusive clusters, the degree of wallet overlap between collections, and measures of systemic concentration. Such indicators provide a broader diagnostic of whether manipulative activity is localized or indicators of system-wide manipulation.

### 2.5.4 Integration Framework

The strength of the multi-scale approach lies in its integration. Micro-level NFT risk scores feed into meso-level wallet cluster analysis, enabling investigators to connect suspicious tokens with the addresses responsible. These wallet clusters are then traced at the macro level, where their impact on collections and marketplaces can be quantified.

This multi-scale framework ensures that NFT fraud detection does not stop at the surface of isolated anomalies but instead reconstructs the full spectrum of manipulative strategies—from single-asset cycles to marketplace-wide distortions—providing a robust foundation for forensic investigation and regulatory oversight.

# CHAPTER 3: METHODOLOGY AND FRAMEWORK IMPLEMENTATION

## 3.1 Dataset and Data Pipeline

This chapter outlines the methodological foundation and the technical implementation of the proposed framework for fraud detection in NFT markets. The architecture of the system was designed to provide a consistent and scalable foundation for detecting fraudulent behavior.

The empirical foundation of the framework relies on Ethereum blockchain data combined with transaction records and metadata extracted from leading NFT marketplaces such as OpenSea and Rarible. The resulting dataset captures both the atomic events (wallet-to-wallet transfers) and the contextual signals (marketplace activity and user behavior) necessary for fraud detection.

Data was imported in two ways:

(i) CSV files for the single-NFT baseline analysis, offering full control and easy replication;

(ii) an API connection through Alchemy for multi-NFT analysis at scale, maintaining the same structure.

The dataset was structured into five modular CSV files, each with a well-defined role:

| File name | Content / Description | Purpose |
|---|---|---|
| wallet.csv | Unique Ethereum wallet addresses involved in NFT activity. | Identification of actors in the network. |
| transactions.csv | Event-level records of transfers and sales, including transaction hash, timestamp, and counterparties. | Captures temporal and structural dynamics of trades. |
| nft.csv | Token-level identifiers (contract address and token ID). | Links transactions to unique digital assets. |
| collection.csv | Metadata connecting tokens to the Quirkies collection. | Groups NFTs under their parent collection. |
| social_profiles.csv | Optional enrichment linking wallets to OpenSea, Etherscan, or ENS profiles. | Provides additional behavioral and identity context. |

Table 1: Dataset structure and role of modular CSV files

The Quirkies collection underwent a migration from its original contract (0x9303…) to a new contract (0xD4B7…). This generated duplicated ownership paths and logically identical transactions with different hashes, resulting in redundant wallet nodes and inflated relationship counts. To preserve data integrity, the analysis adopts a post-migration-only policy, treating the new contract as the authoritative state of the collection.

The graph model implemented in Neo4j captures this structure through five node types, corresponding to the CSV files, and six relationship types:

| Relationship | Description |
|---|---|
| (:Wallet)-[:OWNS]-(:NFT) | Current reconstructed ownership. |
| (:Wallet)-[:TRANSFERRED_TO]-(:Wallet) | Directional link between transaction counterparties, central for detecting structural manipulation patterns. |
| (:NFT)-[:PART_OF]-(:Collection) | Collection membership. |
| (:Transaction)-[:EXCHANGED]-(:NFT) | Association between an event and the asset exchanged. |
| (:Wallet)-[:INVOLVED_IN]-(:Transaction) | Wallet participation in a given event (buyer/seller roles). |
| (:Wallet)-[:HAS_PROFILE]-(:SocialProfile) | Optional connection to public identities. |

Table 2: Graph schema: relationships and meaning

The ingestion pipeline followed a sequential process to ensure consistency. First, nodes representing wallets, NFTs, transactions, collections, and social profiles were imported into Neo4j. Relationships were then created to capture ownership, transfers, transaction participation, and collection membership. This order reduced the risk of incomplete links or invalid references.

Data integrity was ensured through constraints and uniqueness rules. Wallet addresses, transaction hashes, and token IDs were defined as unique identifiers, preventing duplicates from being created during ingestion. Referential consistency was also explicitly validated by verifying that every transaction's “from_wallet” and “to_wallet” addresses matched an existing wallet node, thereby avoiding orphaned relationships.

Data quality was further enhanced through targeted cleaning steps. Wallet addresses were normalized by converting them to lowercase and trimming whitespace to prevent case-sensitive duplicates. Transaction data was checked for missing or malformed hashes, which were either excluded or corrected. In addition, timestamps were normalized into a consistent Unix epoch format—storing each time point as the number of seconds since 1 January 1970. This convention, native to Ethereum, ensured a uniform temporal scale across the dataset and allowed straightforward conversion into ISO 8601 format when required for visualization or reporting.

Additional enrichment was applied by linking NFTs to their respective collection and, where available, associating wallets with verified public profiles from platforms such as OpenSea or Etherscan. These adjustments ensured that the dataset not only mirrored the raw blockchain records, but was also optimized for analytical tasks such as time-window detection, cycle identification, and cross-platform validation.

The schema captures both the historical transaction sequence and the reconstructed ownership—an essential distinction for detecting manipulative behaviors such as wash trading. Neo4j was chosen for its efficiency in detecting complex patterns like cycles, repeated wallet pairs, and dense clusters. This data foundation supports the detailed analysis of Quirkie #3600 as the baseline case study. The structured ingestion and validation approach ensures that detected patterns reflect genuine behavioral signals rather than artifacts, while the network expansion design allows scaling from single-asset to ecosystem-wide analysis.

## 3.2 Single-NFT Baseline

Quirkie #3600 serves as the reference case study for validating the proposed detection framework. Its clean ownership history, moderate transaction volume, and the presence of multiple wallet types make it an ideal benchmark for distinguishing legitimate market activity from manipulative behavior. By focusing on this single asset, the analysis establishes both the methodological approach and the baseline metrics against which broader network dynamics can later be compared.

The reconstructed network for Quirkie #3600 consists of 26 nodes distributed across five categories: seven wallets, six transactions, one NFT, one collection, and eleven social profiles. These entities are connected by 37 relationships, the majority reflecting wallet participation in transactions and wallet–profile linkages.

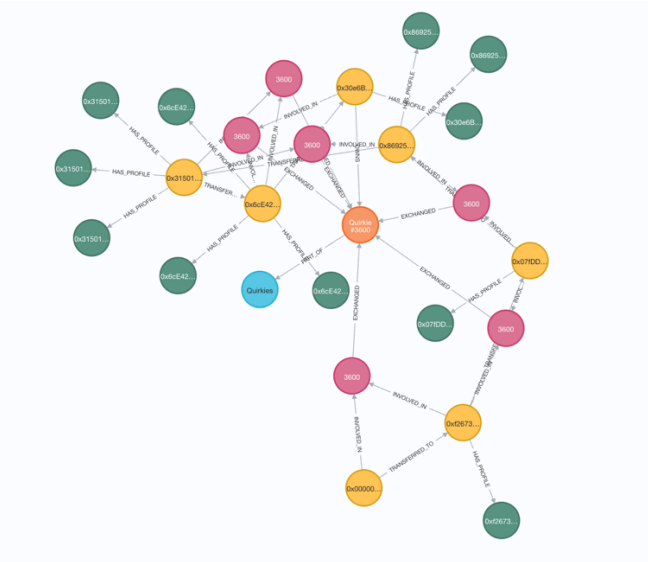

*Figura 3.1. Network representation of Quirkie #3600*

Within the wallet-to-wallet subgraph, the network achieves a density of approximately 14.3%, consistent with a sequential ownership chain rather than the dense clusters that would typically suggest collusive trading. Degree distribution supports this structural interpretation: transaction nodes display the highest connectivity due to their role as intermediaries, while among wallets, frequent traders such as Trader4 and Trader5 exhibit greater centrality than long-term holders.

Homophily analysis tests whether wallets of the same type are more likely to connect than expected by chance. In the Quirkie #3600 network, about 50% of transfers occur between wallets of the same type, compared to a baseline of 28.6% expected under random mixing. This suggests a moderate tendency for similar wallets to interact, but the effect is limited by the small size of the network. Diadicity ratios show a slight overrepresentation of trader-to-trader connections, though not at a level that would indicate collusion.
The local clustering values also reflect the chain-like structure of ownership. Whenever a wallet has both a predecessor and a successor, those neighbors are automatically connected through the transfer sequence, which produces a clustering coefficient of 1.0.

Table 3: Structural Properties of the Quirkie #3600 Wallet-to-Wallet Subgraph

| **Metric** | **Observed Value** | **Baseline / Expected Value** | **Interpretation** |
|---|---|---|---|
| Network Density | 14.3% | High (>40%) indicates dense clusters | Sparse structure consistent with sequential chain; no collusive concentration. |
| Degree Distribution | Low, skewed | Manipulation shows hub nodes with > 8–10 connections | Transaction nodes naturally highest; Trader4/Trader5 more central but within normal range. |
| Homophily | 50% | 28.6% under random mixing | Moderate within-type preference; not strong enough to suggest manipulation. |
| Diadicity (Trader–Trader) | Slightly overrepresented | Random mixing baseline | More trader-to-trader edges than expected, but density too low for collusion. |
| Local Clustering Coefficient | 1.0 (wallets with ≥2 neighbors) | 1.0 expected in linear chains; 0.6+ indicates dense clustering | Normal for sequential ownership; high clustering would indicate manipulation. |
| Connectivity Validation | Robust (no isolates) | Fragmented networks often show orphans | Fully connected, no missing links; consistent with clean transfer sequence. |

Finally, connectivity validation confirms the robustness of the reconstruction, with no isolated wallets, orphaned transactions, or disconnected components present.

Building on this structural foundation, the temporal analysis of ownership patterns reveals a clear chronology. The ownership history of Quirkie #3600 spans from its mint in April 2023 to its current holder in August 2025, encompassing six transfers across seven distinct wallets. The first owner retained the asset for nearly two years, reflecting the behavior of a long-term collector. In contrast, subsequent transfers in 2025 occurred in rapid succession, compressing into a four-month window. Holding durations therefore vary substantially: while the InitialOwner held the NFT for 670 days, later traders maintained ownership for only days or even hours, with Trader5 transferring the asset less than eight hours after acquisition. Despite this acceleration, the sequence follows a natural progression without evidence of circular trading or rapid flipping designed to simulate liquidity. Transfers occurred through both direct wallet exchanges and established marketplaces such as OpenSea, suggesting organic usage of legitimate transaction channels.

To validate these behavioral interpretations, social profile data provides additional context. Eleven profiles were identified across OpenSea, Twitter, and ENS, with active traders maintaining the most extensive multi-platform presence. In particular, Trader4 and Trader5 linked their wallets to all three platforms, while collectors generally limited their presence to marketplace profiles. Cross-platform consistency was verified by comparing usernames, biographical information, and activity timelines across platforms, reinforcing the credibility of these identities. The presence of ENS registrations adds further weight, as such domains require costs and continuity, reducing the likelihood that they are associated with disposable fraud accounts.

Taken together, these results show that the network of Quirkie #3600 is structurally simple and consistent with legitimate trading. The descriptive checks covered the essential dimensions of network analysis: node and relationship counts, global density, degree distribution, reciprocity and diadicity of transfers, wallet-type homophily, local clustering coefficients, and overall connectivity validation.

Across all metrics, the picture remains that of a straightforward transfer sequence, without anomalies such as cycles, dense clusters, or disconnected components.

## 3.3 Graph Algorithms

Building on the diagnostics from Section 3.2, graph algorithms were applied through Neo4j's Graph Data Science (GDS) library to analyze wallet interactions and test for structural anomalies. The GDS toolbox provides a scalable set of algorithms—ranging from centrality to community detection and similarity analysis—designed to quantify influence, grouping tendencies, and behavioral overlap in trading networks. For Quirkie #3600, the algorithms serve not only to characterize its ownership structure but also to establish a methodological baseline for later comparison with expanded datasets.

To support different analytical needs, two projections of the wallet network were created. The undirected version (walletUndirected) simplified transfers into mutual links, useful for studying overall cohesion and community structure. The directed projection (walletDirected) preserved the original transaction flow, which is essential for algorithms such as PageRank where the direction of influence matters.

Centrality results highlight the distribution of importance within the chain. Degree centrality was uniformly low, with most wallets connected to only one or two others, consistent with a stepwise progression of ownership. Betweenness scores identified Trader3, Trader2, and Trader4 as the main intermediaries bridging earlier and later parts of the chain. Closeness centrality confirmed this picture, ranking Trader3 as the most "accessible" node, positioned at the center of the transfer flow. PageRank, by contrast, assigned higher scores to the CurrentOwner (0.679) and Trader5 (0.623) due to their positions in the directed link structure, where they receive authority from preceding nodes in the ownership sequence. In a collusive cluster, such values could indicate suspicious influence; here they simply reflect the natural flow of authority through a linear transfer chain.

Structural measures reinforced the lack of anomalous patterns. No cycles of length 3–5 were detected, ruling out circular trading loops. The global clustering coefficient was exactly 0.0, consistent with a network where wallets connect in sequence but not to one another's neighbors. Average path length was 2.67, with a shortest path of 1 and a longest of 6, confirming moderate sparsity but uninterrupted connectivity across the sequence. Community detection using Louvain assigned wallets into three partitions (sizes 3, 2, and 2), yielding a modest modularity of 0.319. However, these groups merely reflected consecutive segments of the chain rather than meaningful clusters. This absence of dense groupings provides strong evidence against coordinated substructures.

Overlap metrics offered further perspective. Jaccard indices were moderate (e.g., 0.5 between CurrentOwner and Trader4; 0.333 for InitialOwner and Trader3), while Cosine scores reached as high as 0.707 between adjacent wallets. On the surface, these values might appear significant, but they are mechanically induced by the simplicity of the dataset: when two wallets each perform a single transfer, their behavioral overlap is artificially perfect.

Temporal and behavioral homophily reinforced this observation—many adjacent pairs scored 1.0 because they mirrored one another in transaction counts, with the only deviation occurring at the final step, where the CurrentOwner still retains the asset. In larger datasets, such strong alignments would be suspicious; here they reflect trivial symmetry in a linear chain.

Collectively, the algorithmic results confirm the absence of manipulative structures within the network. Uniformly low degrees, the absence of cycles, zero clustering, modest modularity, and mechanically induced similarities converge toward the same conclusion: Quirkie #3600's history exhibits the structural footprint of legitimate trading. As expected for a single NFT, its sequence of ownership forms a linear chain that serves as a negative control, clarifying which algorithmic signals are trivial in a clean baseline and which become meaningful indicators only in more complex multi-token networks.

In larger trading graphs, by contrast, red flags would typically appear as clustering coefficients above 0.6, reflecting the unusually dense neighborhoods often associated with collusive subgroups (von Wachter et al., 2022), modularity scores above 0.4 in small partitions, consistent with the community detection thresholds widely used in network science (Newman & Girvan, 2004), or repeated Jaccard similarities above 0.7 across wallet pairs, values that prior blockchain forensics research has linked to coordinated or common-control behavior (Nadini et al., 2021; Chainalysis, 2018).

## 3.4 Pattern Detection Modules

To complement the structural graph analysis, a set of operational modules was applied to detect specific patterns of fraudulent trading. These modules combine structural rules, temporal thresholds, and scoring logic, enabling the risk assessment as zero, low, medium, or high depending on the strength of detected signals.

The analysis relies on authentic blockchain data, including wallet addresses, transaction hashes, timestamps, and transfer relationships extracted directly from the Ethereum blockchain. This ensures that the detection of transaction-based and timing-based manipulation is grounded in verifiable evidence. Within this scope, the methodology can reliably capture wash trading, circular trading, coordinated buying and selling, volume inflation, bot-like activity, and smart contract manipulation.

One limitation concerns price-dependent fraud patterns such as pump-and-dump schemes. NFT transactions recorded on-chain typically contain zero ETH values, since financial settlement is often handled through marketplace smart contracts or off-chain mechanisms. As a result, actual sale prices cannot be reconstructed directly from blockchain data. To address this constraint, simulated prices were temporarily introduced in controlled tests. This allowed the system to demonstrate how it would respond to pump-and-dump dynamics if reliable pricing information were available.

The first detection approach focused on wash trading patterns, examining repeated transfers of the same NFT between identical wallet pairs. The scoring methodology distinguishes between isolated transfers (low risk), bidirectional exchanges without temporal clustering (medium risk), and rapid back-and-forth transfers occurring within 24 hours (high risk). Analysis of Quirkie #3600 revealed no such repeated exchanges, resulting in a zero risk classification.

Building on this foundation, circular trading detection extended the analysis to larger wallet groups, systematically searching for closed loops involving three to six participants. Risk assessment escalates proportionally with cycle frequency: moderate concern arises when one or two loops are present, while high risk indicators emerge when repeated loops suggest sustained coordination among participants. The Quirkie baseline demonstrated no closed ownership cycles, reinforcing the straightforward nature of its ownership progression.

Coordinated trading surveillance employed 15-minute temporal windows to identify synchronized market behavior. Within each interval, the system simultaneously evaluated the number of distinct participating wallets and measured price convergence patterns. Low risk scores emerge when two wallets operate within the same timeframe, medium scores when three or more wallets execute transactions with partially aligned pricing, and high scores when four or more participants converge on highly similar valuations. When applied to Quirkie #3600, this module generated a low risk score, supporting the interpretation of natural trading dynamics rather than orchestrated market manipulation.

To test price manipulation detection capabilities, a specialized module simulated pump-and-dump dynamics through artificial price trajectory modeling. Given that raw NFT transfers typically contain zero ETH values, synthetic pricing data was temporarily introduced, incorporating gradual appreciation followed by sharp appreciation and subsequent market crash patterns. The detection algorithm triggers alerts when price spikes exceed 30% followed by declines of 25% or more, particularly when combined with limited wallet diversity during active trading periods. Under these controlled conditions, the system correctly identified high risk indicators, validating its capacity to detect such manipulations when comprehensive pricing data becomes available. Following testing, all synthetic values were removed to maintain dataset authenticity.

Moving beyond price-focused analysis, the evaluation encompassed both volume inflation and automated trading detection. Volume inflation assessment monitored transaction clustering within one-hour periods, with risk scores increasing when high trade volumes concentrate among limited wallet participants. The Quirkie analysis yielded low risk scores, as trading activity never approached levels typically associated with artificial liquidity manipulation. Concurrently, bot activity detection analyzed inter-transaction timing gaps for individual wallets, where ultra-short intervals measured in seconds would indicate automated control systems. The algorithm prioritizes detection of burst patterns featuring sub-minute transaction gaps, though none were identified in this case, producing zero risk scores while reinforcing evidence of human-controlled trading behavior.

The final component addressed smart contract manipulation through systematic anomaly detection. This module identifies irregularities including abnormally low ETH transfer values, repetitive unusual cross-contract interactions, and deviations from standard marketplace operational logic. Comprehensive analysis of the Quirkie network revealed no such technical irregularities, maintaining the consistent pattern of zero risk classifications across technical exploitation vectors.

The comprehensive testing across all operational modules validates the detection framework's readiness for deployment on larger, more complex datasets. While Quirkie #3600 consistently produces minimal risk indicators—confirming its value as a clean baseline—the successful identification of synthetic pump-and-dump patterns demonstrates the system's sensitivity to genuine manipulation signals.

This dual validation establishes both the framework's specificity in avoiding false positives and its capacity to detect authentic fraud patterns when present.
With the methodological foundation now established through single-asset analysis, the framework can scale to multi-token networks where manipulative behaviors are more likely to manifest. The expanded analysis will reveal whether the patterns observed in individual NFTs aggregate into systematic fraud schemes, and whether the detection thresholds calibrated on clean baselines can effectively distinguish organic market activity from coordinated manipulation across diverse collections and trading contexts.

| **Detection Module** | **Risk Thresholds** | **Quirkie #3600 Observed Pattern** | **Risk Level** |
|---|---|---|---|
| Wash Trading | Rapid back-and-forth transfers within 24 hours (High); Bidirectional exchanges without clustering (Medium) | No repeated exchanges between wallet pairs | Zero |
| Circular Trading | Repeated loops suggest coordination (High); 1–2 loops present (Medium) | No closed ownership cycles detected | Zero |
| Coordinated Activity | 4+ wallets converging on similar values in 15-minute windows (High); 3+ wallets with aligned pricing (Medium) | Natural timing patterns observed | Low |
| Pump-and-Dump | 30% price spike followed by 25% decline with low wallet diversity (High) | High risk detected under synthetic test conditions | High |
| Volume Inflation | High trade volumes concentrated among limited participants in 1-hour windows | Normal activity levels; no artificial liquidity patterns | Low |
| Bot Activity | Sub-minute transaction gaps indicating automated control (High); Burst patterns (Medium) | Human-like inter-transaction intervals | Zero |
| Smart Contract Manipulation | Abnormally low ETH transfers, unusual cross-contract interactions | No technical irregularities detected | Zero |

Table 4: Pattern Detection Module Results for Quirkie #3600

# CHAPTER 4: EXPANDED NETWORK ANALYSIS

## 4.1 Network Expansion Methodology

Building on the baseline established in Chapter 3, this study expands from a single NFT—Quirkie #3600—to a controlled, manageable network designed for graph analysis and fraud-pattern detection. The process combines API-driven organic discovery with limited data expansion so that the graph reflects real trading structure without growing uncontrollably.

The first step in the expansion process involved extending the Quirkie #3600 baseline network through a holdings-based approach. Specifically, the analysis retrieved the complete set of NFTs currently owned by the wallets directly tied to the Quirkie #3600 ownership chain, including the InitialOwner, the CurrentOwner, and four high-activity adjacent traders (Trader2–Trader5). This used Alchemy's getNFTsForOwner endpoint to extract contract addresses, token IDs, collection names, token standards, and associated metadata. To guarantee cross-collection uniqueness and prevent identifier collisions, every NFT was assigned a compound UID in the form contract:tokenId. This ensured that tokens from different collections but with identical numeric IDs could coexist in the graph without duplication. At the same time, wallet addresses were normalized to lowercase to remove inconsistencies due to checksum formatting, thereby avoiding fragmented degree counts and ensuring that all transfers aggregated to the correct entity.

For each wallet–NFT pair, the graph database recorded a wallet node, an NFT node keyed by the UID, a directed ownership relationship labeled with source metadata (e.g., retrieval origin and timestamp), and a link from the NFT to its collection. Where available, additional attributes such as token type (ERC721 or ERC1155), media links, and collection metadata were harmonized across different API versions to maintain consistency in the dataset.

The holdings expansion created a structured subgraph that incorporated not only the Quirkie token itself but also the broader portfolio of its historical owners. By capturing these cross-collection ties, the expansion provided a more complete view of how wallets operate across the market, highlighting latent connections between collections that would not be visible in a single-token view. At the same time, the process remained controlled: only wallets directly connected to the Quirkie ownership chain were included, preventing uncontrolled growth while still enriching the research context. The result was a first expanded layer of the dataset that laid the groundwork for subsequent fraud-pattern modules by embedding ownership in its broader portfolio context.

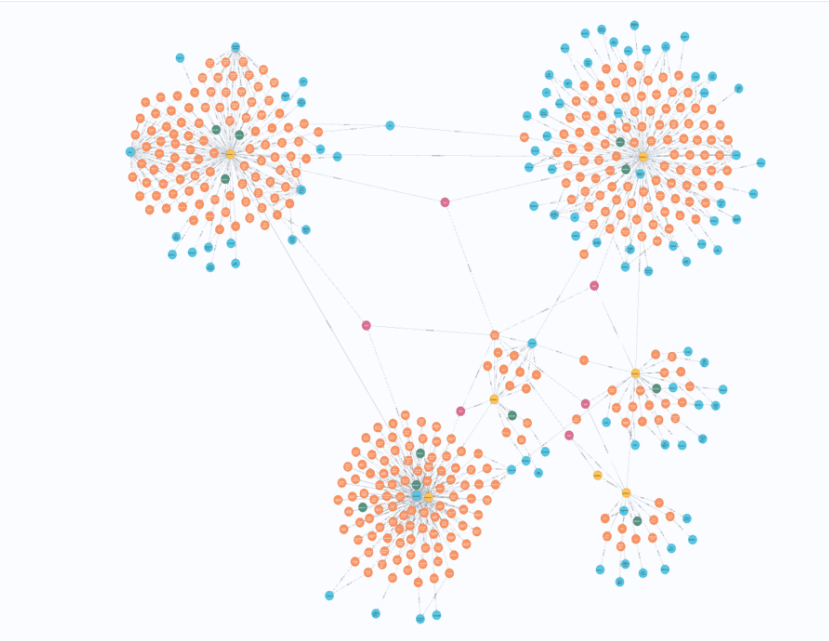

*Figura 4.1. Wallet Holdings Expansion of Quirkie #3600 Ownership Chain*

The second stage of the network growth moved from static wallet holdings to the dynamic layer of trading interactions. Whereas the first expansion described which assets were held by the Quirkie #3600 ownership chain, this phase focused on identifying and incorporating the most relevant trading partners of its core traders. The objective was to capture wallet-to-wallet connections that reveal patterns of exchange, while maintaining a limited and analyzable dataset.

The procedure focused on the four most active wallets—Trader2, Trader3, Trader4, and Trader5—identified from the initial expansion. For each, historical transfer data within a ninety-day period surrounding the Quirkie #3600 transaction was analyzed to extract their five to six top trading partners, ranked by frequency of interactions. These counterparties were then introduced into the graph together with the transactions and NFTs involved, ensuring that relationships remained grounded in specific digital assets and collections.

To avoid uncontrolled growth, the expansion relied on capped synthetic transactions rather than full historical imports. Each trader was assigned a maximum number of transactions in proportion to its overall activity, with high-volume actors such as Trader5 represented by twenty-five transactions and lower-volume traders such as Trader2 limited to ten. Within these caps, synthetic links were allocated through proportional distribution across partners according to their observed share of activity and preserved the original direction of transfers. Every synthetic transaction was connected to an NFT node, guaranteeing that wallet-to-wallet flows remained tied to tokenized objects relevant for detecting manipulation schemes such as wash trading or circular trading.

The outcome of this expansion was an expanded trading network that bridges the asset-holding view with the transaction-driven dynamics of the NFT market. By combining frequency-ranked partner inclusion, capped synthetic transactions, proportional distribution, and rigorous validation rules, the graph captures the core motifs of NFT trading behavior—dense dyads, short trading cycles, and partner clusters—without overwhelming the research framework. The resulting structure supports centrality analysis, similarity scoring, clustering, and anomaly detection, while remaining manageable. As such, the trading partner expansion demonstrates how a scalable yet disciplined methodology can extend fraud detection frameworks from isolated NFT holdings to the multi-wallet ecosystems where manipulation schemes typically unfold.

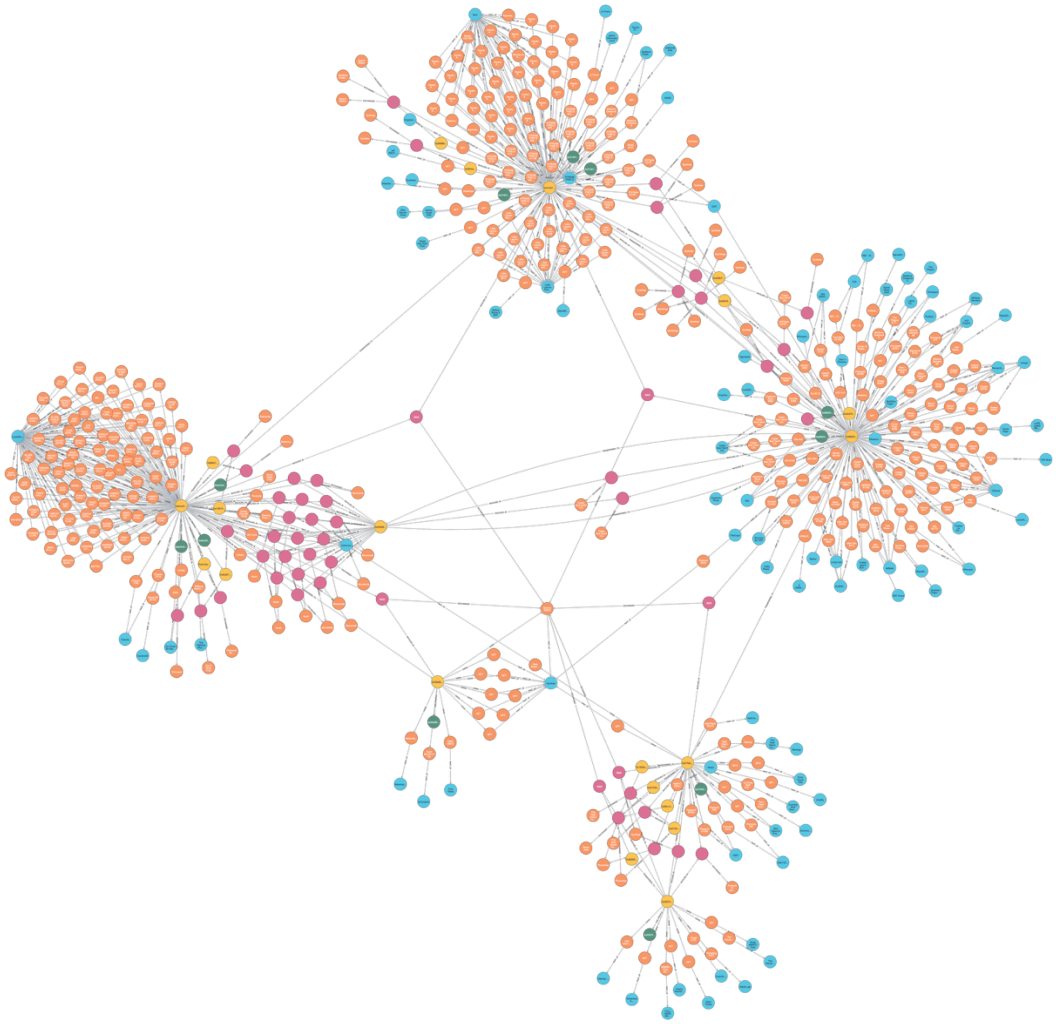

*Figura 4 2. Trading Partner Expansion of the Quirkie #3600 Network*

## 4.2 Expanded Network Characteristics

The expansion process transformed the network from a minimal baseline of 26 nodes into a heterogeneous structure comprising over 1,465 entities. These include wallets, NFTs, collections, and transactions, each contributing to the representation of different layers of marketplace activity. Such growth illustrates how a single-asset starting point can evolve into a multi-scale environment where structural patterns, trading behaviors, and asset flows can be observed simultaneously.

The increase in scale not only expands the analytical horizon but also introduces the complexity necessary for testing fraud detection mechanisms under more realistic market conditions.

In terms of connectivity, the expanded network exhibits a directed density of roughly 6% and an undirected density of about 11%. This confirms that the graph remains sparse, as expected in open blockchain systems, yet the density is sufficient to highlight concentrated regions of activity. The degree distribution follows a heavy-tailed profile: while the majority of wallets connect to only one or two peers, a small subset reaches much higher degrees, with two wallets showing eight direct connections and several others ranging between four and six. This skew is consistent with power-law tendencies often documented in digital transaction networks, where a minority of highly active actors drive a disproportionate share of the interactions.

Path length and clustering measures further illuminate the structure. The global clustering coefficient reveals strong local cohesion around specific triads, with some wallets displaying coefficients close to 1.0, meaning their neighbors are densely interconnected. At the same time, other actors present coefficients near zero, reflecting star-like interaction patterns where trades are routed through central hubs without reciprocal connections. Shortest paths between wallets remain relatively small due to the presence of these hubs, while longer sequences are observed in more peripheral segments of the network. This coexistence of tightly knit clusters and extended paths suggests a layered topology combining both local trading circles and broader market diffusion.

Finally, relational properties such as reciprocity and homophily offer insight into behavioral tendencies. Approximately 23% of wallet pairs are reciprocal, indicating that a meaningful share of trading relationships are bidirectional rather than one-sided. This aligns with practices such as repeated exchanges or potential collusive arrangements. Homophily, however, remains low: only about 11% of wallet-to-wallet edges connect nodes with the same type label. The majority of links thus cross between different wallet categories, suggesting heterogeneous mixing of roles in the market rather than segmentation into isolated communities.

Taken together, these results portray an expanded network that is simultaneously sparse and heterogeneous, where a few influential actors concentrate activity while the majority participate peripherally. Such characteristics form a critical foundation for fraud detection, since anomalies like wash trading, circular trading, or coordinated flips are most visible when contrasted against the baseline expectations of sparse connectivity, low clustering, and heterogeneous role mixing.

## 4.3 Graph Algorithm Results

The application of centrality, community detection, and similarity algorithms to the expanded network provides a layered perspective on how influence, subgroup formation, and behavioral convergence unfold within NFT trading activity. Centrality metrics establish the first dimension of this perspective by highlighting which wallets anchor the structure. Degree scores identify the most connected hubs: Trader5 with seven direct peers, followed closely by Trader4, Trader2, and Trader3 with six connections each.

These wallets extend their reach to a wider set of counterparts than others, but high connectivity alone does not necessarily confer structural power. Betweenness centrality adds an important nuance, revealing which actors control the shortest paths between others and therefore shape how tokens flow across the network. Here, Trader3 emerges as the leading intermediary (107.0), trailed by Trader2 (98.0), Trader5 (86.5), and Trader4 (76.5). Their positions indicate that they are not only active traders but also brokers of circulation, able to influence how assets move between otherwise weakly connected peers.

Closeness centrality reinforces this picture, with Trader3 displaying the highest score (0.400) and thus occupying a vantage point from which it can reach all others more quickly than most. PageRank, which weights influence by the authority of connected neighbors, again consolidates Trader3's prominence (1.607), placing it ahead of Trader4, Trader5, and a handful of peripheral but strategically linked wallets. Taken together, these measures converge on a compact elite of addresses that, despite their differences in role, collectively anchor connectivity, access, and authority within the expanded network.

The application of community detection highlights the tendency of activity to consolidate into well-defined clusters rather than spread uniformly. Louvain partitioning divides the graph into three sub-communities with a modularity score of approximately 0.513, a moderate value that points to meaningful but not absolute separation. The first cluster is organized around Trader3 and Trader4, whose repeated exchanges and dense overlaps bind a group of closely linked addresses. A second cluster coalesces around Trader5 and the CurrentOwner, forming a hub-and-spoke pattern where most links radiate outward rather than interconnect. The third cluster follows a different logic, linking MintSource, InitialOwner, and Trader2 into a generational chain that reflects the NFT's passage from creation through early circulation.

Directed cycle detection reveals recurrent loops of length three to five, often involving Trader3, Trader4, and wallets such as 0x2946… and 0x5654…. While short cycles are not inherently suspicious—they may arise from collectors repeatedly exchanging assets—their concentration around the same addresses warrants attention. Triangle counts reinforce this impression, showing that Trader3 and Trader4 in particular are embedded in compact motifs where neighbors are strongly interconnected. The outcome is a topology where certain groups form dense, tightly interlinked cliques, while others extend into longer chains, together shaping a network that balances local cohesion with wider diffusion.

A further behavioral characteristic emerges through similarity and homophily analysis. Node similarity measures, based on Jaccard and Cosine indices, reveal cases of perfect overlap (scores = 1.0), where pairs of wallets share identical sets of partners and proportional flows. Such exact symmetry is unusual in open trading systems and may indicate shared control, automated replication, or deliberately coordinated strategies. Homophily analysis extends this view by examining whether connected wallets behave alike in terms of transaction ratios. Certain pairs, notably Trader3–Trader4 and Trader5 with selected partners, show complete convergence in their exchange patterns, while others stabilize at a ratio of of about 0.67 suggesting asymmetric but consistent flows where one wallet persistently supplies another. These convergences are unlikely to arise purely by chance in a heterogeneous environment, and while they do not in themselves prove collusion, they represent structural signals that merit closer examination.

Overall, the algorithmic analyses reveal that the expanded network is highly uneven. A small cluster of wallets—most notably Trader3, Trader4, and Trader5, with Trader2 acting as a bridge—emerges as disproportionately central. These actors not only maintain more connections than others but also occupy the shortest paths through which tokens flow, group together in semi-autonomous communities, and show trading behaviors that at times appear almost mirrored. Such traits make them natural focal points for closer investigation—not because the metrics alone prove manipulation, but because the overlap of structural influence, subgroup cohesion, and behavioral similarity creates the conditions in which fraudulent schemes are most likely to occur. In this sense, combining centrality, community, and similarity measures reveals patterns of possible coordination that a purely descriptive view of transactions would not expose.

## 4.4 Fraud Pattern Detection

To complement the structural diagnostics and centrality analysis, the expanded network was systematically screened for potential fraudulent trading patterns using a set of detection modules. Each module combined structural criteria with temporal thresholds applied to authentic blockchain data (wallet addresses, transaction hashes, timestamps). The resulting classifications should be understood as signals of risk rather than conclusive evidence of manipulation, since definitive attribution would require triangulation with off-chain sources such as marketplace records, wallet metadata, or social traces.

The first set of findings relates to systemic manipulation. Wash trading modules flagged elevated risk, with evidence of rapid, repetitive transfers of the same token between the same wallets within short intervals—sometimes less than an hour. Such patterns are consistent with artificial volume designed to simulate demand. Extending beyond pairs, circular trading analysis uncovered numerous closed loops of transactions: Trader3 alone appeared in 20 cycles of length three to five, while Trader4, Trader5, and another high-activity address also showed frequent participation. These circuits recycle assets through small groups of wallets, giving the impression of liquidity and diverse ownership while effective control remains unchanged. By contrast, volume inflation was less pronounced, surfacing primarily in Trader2, where short bursts of concentrated activity were visible, though overall ratios of transactions to unique partners remained below the levels typically associated with large-scale manipulation. Bot-like behavior was more evident: ultra-short transaction gaps placed Trader2 in the highest risk tier, while Trader3 and Trader5 showed medium-level signals consistent with partial automation.

A second cluster of modules examined coordination across wallets. Here, the focus was on synchronized buying and selling—cases where multiple addresses transacted at similar prices within very narrow time windows.

A handful of wallets, most notably Trader5, were involved in at least eight events, pointing to possible attempts to simulate demand. Price-based fraud patterns such as pump-and-dump could not be directly tested due to the absence of reliable sale price data in the on-chain records. However, synthetic simulations were used to demonstrate how the framework would flag steep coordinated rises of over 30% followed by collective drops exceeding 25%—dynamics in line with pump-and-dump events documented in prior NFT literature. Importantly, no genuine pump-and-dump events were detected in the observed dataset; all flagged cases emerged from simulated scenarios.

These findings underscore how patterns of risk only become visible at scale. At the single-NFT level—such as Quirkie #3600—metrics suggested negligible exposure, with nearly all modules returning “LOW” or “ZERO.” Yet, once activity was aggregated across collections and expanded wallets, recurring loops, reciprocal dyads, and coordinated clusters began to appear. A small set of traders, particularly Trader3 and Trader5, surfaced repeatedly across multiple modules, suggesting that manipulation in NFT markets may operate less through isolated anomalies and more through systematic coordination anchored by hub actors. The combination of structural and temporal modules thus highlights a vulnerability of NFT ecosystems: legitimacy at the individual token level does not necessarily translate into systemic integrity once networks of interaction are considered.

## 4.5 Visualization and Macro-level interpetation

Building on the algorithmic diagnostics and fraud detection modules, the analysis was extended with visual representations to capture macro-level structural patterns. These representations were produced in Neo4j Bloom, which allows mapping node properties and relationships into interpretable layouts through color gradients, edge thickness, and selective labeling. The goal was not to explore interactive dashboards but to produce static figures highlighting the structural patterns of ownership and overlap.

As illustrated in Figure 4.5a, the first mapping focuses on hub wallets. Wallet nodes are colored according to the number of distinct collections they hold, with a gradient from green (few collections) to red (many collections), while node size is proportional to degree centrality. The figure shows how a small subset of wallets structurally dominates the expanded network, functioning as multi-collection hubs capable of mediating activity across otherwise disconnected subgroups. This concentration of ownership highlights systemic leverage points where coordination or manipulation could be amplified.

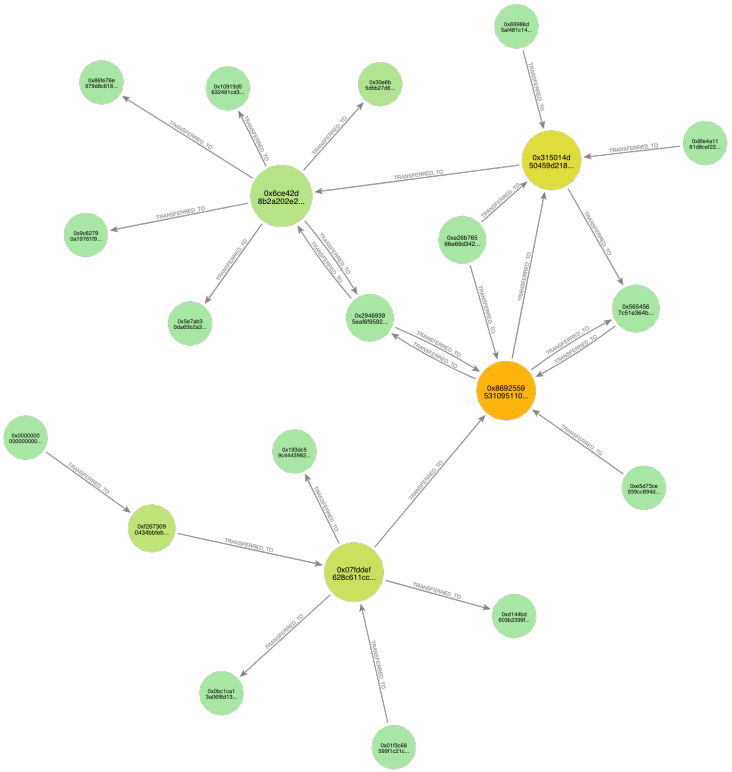


*Figura 4.5a*

Figure 4.5b shows the second visualization, which maps collection-to-collection overlap using the Jaccard similarity index of shared wallet ownership. Edge thickness is proportional to the similarity value, while node colors reflect the degree of connectivity (green = low overlap, red = high overlap). The figure reveals tightly knit clusters of collections that share many common holders, effectively forming “ownership-overlap clusters”. These artificial clusters, while not necessarily fraudulent, demonstrate how structural overlap can blur the boundaries between projects, enabling cross-collection coordination strategies.

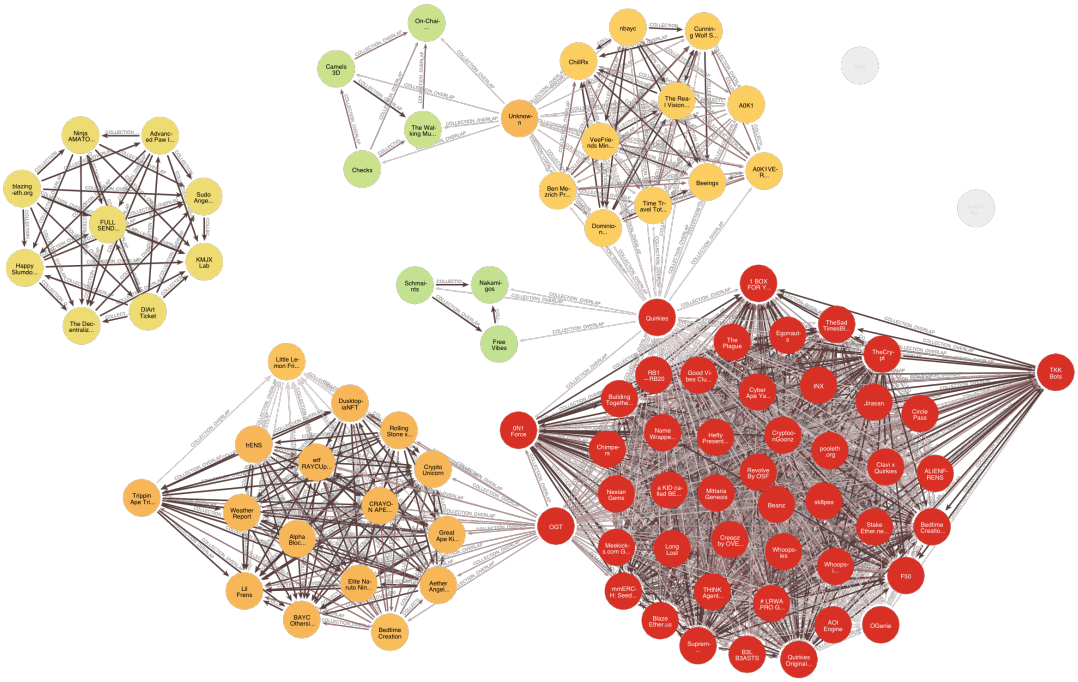

*Figura 4.5b*

When compared with the micro-level results from Quirkie #3600, which showed consistently low or zero manipulation risks, the macro-level visualizations underscore the importance of network-level analysis. While individual NFTs may appear legitimate when analyzed in isolation, network-level perspectives reveal emergent vulnerabilities: hub wallets centralizing influence, and clusters of overlapping collections suggesting opportunities for coordinated activity. This contrast between micro and macro findings stresses that legitimacy at the token level does not necessarily translate into systemic integrity.

Finally, it is important to note a methodological disclaimer: these results are bounded by the Quirkie-connected network. They do not represent absolute market-wide behavior, but rather relative coordination signals within the sampled dataset. Nonetheless, the visual approach demonstrates how graph-based perspectives can surface structural risks invisible in token-level transaction analysis.

# CONCLUSIONS

This thesis examined the detection of fraudulent behaviors in NFT markets through a multi-scale, graph-based framework, progressively expanding the analytical lens from a single token to a broader ecosystem of wallets, collections, and trading interactions. The central insight is that manipulation in NFT markets cannot be fully understood when individual assets are studied in isolation. Instead, fraudulent activity becomes most visible when structural, temporal, and community-level perspectives are brought together and analyzed as part of a connected network.

## Summary of Results

At the micro level, the case of Quirkie #3600 provided a clean benchmark. Its ownership history was linear, free from cycles or anomalous bursts of trading, and it served as a negative control against which manipulation could be measured. By contrast, once the analysis expanded to meso- and macro-level scales, the network revealed a very different picture. A small group of wallets—most prominently Trader3, Trader4, and Trader5—emerged as disproportionately central across multiple measures of connectivity, brokerage, and authority. These actors recurred not only in centrality metrics but also in community partitions and behavioral similarity indices, consistently anchoring activity in ways that set them apart from the wider population of peripheral traders.

Fraud detection modules reinforced this impression. Patterns of repeated back-and-forth transfers flagged instances of wash trading, while closed trading cycles and clustered bursts of activity pointed toward coordinated behavior. Bot-like signatures were also detected, particularly in the form of ultra-short transaction gaps, suggesting the presence of automated strategies. Crucially, these signals were not apparent when the network was confined to a single NFT; they only surfaced once trading interactions were aggregated at scale. This confirms that manipulation in NFT markets is not dispersed evenly but tends to concentrate in clusters of wallets that operate with unusual intensity and synchronization.

## Theoretical and Practical Contributions

The study makes contributions on both theoretical and practical fronts. Theoretically, it adapts descriptive taxonomies of NFT fraud into an operational framework that translates abstract manipulation categories into measurable graph-theoretic signatures. In doing so, it bridges the gap between conceptual discussions of market abuse and the algorithmic methods required to detect them. By formalizing patterns such as wash trading or circular trading as network motifs, the thesis provides a foundation for more rigorous, repeatable investigations of digital asset markets.

On the practical side, the framework demonstrates how investigators, auditors, and regulators can deploy centrality measures, cycle detection, and similarity indices to identify suspicious trading clusters. It illustrates how a structured expansion methodology, combined with clear detection rules, can generate actionable insights for forensic auditing, marketplace monitoring, and policy oversight. Perhaps most importantly, the research shows how fraud indicators escalate across different scales of analysis, offering a replicable blueprint for multi-layered NFT forensics that balances precision with scalability.

## Limitations and Ethical Considerations

Despite these contributions, several limitations must be acknowledged. First, the analysis relied in part on synthetic transactions to supplement sparse data. While these were carefully capped, proportionally distributed, and explicitly flagged, they may still introduce artifacts that diverge from real market behavior. Future studies should aim for full historical datasets where computationally feasible.

Second, the study was bounded to the Quirkie-connected component, meaning the findings reflect relative coordination signals within a subset of the NFT market rather than universal dynamics. Extending the framework across multiple anchors, collections, or longer temporal windows would improve external validity and capture a wider range of manipulative practices.

Third, the absence of reliable on-chain pricing data limited the direct detection of pump-and-dump schemes. Synthetic simulations were used to demonstrate detection logic, but integration with marketplace APIs or oracle services remains necessary for full implementation.

Beyond methodological constraints, the research also raises ethical challenges. Advanced clustering and similarity techniques risk producing detailed behavioral profiles that could compromise pseudonymity when combined with external data. Incorrectly flagging legitimate actors as suspicious may also carry reputational and financial harm.

## Future Directions

Looking ahead, three avenues for development stand out.

- Money Flow Analysis.
  Incorporating transaction values—both in ETH and in stablecoins—would allow the detection of schemes that manipulate not just ownership structures but also economic flows. This would enable direct identification of pump-and-dump events, laundering strategies, and artificial wealth inflation. Linking graph structure with monetary value would also strengthen connections to anti-money-laundering research.

- Multi-Chain Extension.
  As NFT activity increasingly spans Ethereum, Polygon, Solana, and other blockchains, extending the framework across chains would provide a more holistic view. Analyzing bridge transfers and cross-chain liquidity would also expose opportunities for obfuscation and regulatory arbitrage. Comparative studies across ecosystems could further reveal platform-specific vulnerabilities or resilience factors.

- Real-Time Monitoring and AI.
  Embedding the framework into streaming graph systems would allow investigators to detect fraud as it happens rather than retrospectively. Machine learning models trained on confirmed manipulation cases could improve accuracy, while anomaly detection techniques could flag novel schemes not captured by existing rules. Automated alerting systems, combined with predictive models, could even anticipate emerging manipulation campaigns, providing regulators and marketplaces with early warning capabilities.

## Concluding Remarks

In conclusion, this thesis demonstrates that NFT markets exhibit fraudulent behaviors that are systemic, multi-layered, and most effectively captured through graph-based network analysis. The results show how clusters of wallets, rather than isolated actors, drive much of the suspicious activity, and how these behaviors become visible only when the analysis scales beyond the single-asset level.

While challenges of data completeness, scalability, and ethics remain, the proposed framework contributes to both academic debate and practical solutions for strengthening transparency, accountability, and trust in digital asset markets.

# BIBLIOGRAPHY

*[2305.01543] NFT Wash Trading Detection*. (n.d.). Retrieved 30 August 2025, from https://arxiv.org/abs/2305.01543

*[2504.16113] AI-Based Vulnerability Analysis of NFT Smart Contracts*. (n.d.). Retrieved 29 August 2025, from https://arxiv.org/abs/2504.16113

Alizadeh, S., Setayesh, A., Mohamadpour, A., & Bahrak, B. (2023). A network analysis of the non-fungible token (NFT) market: Structural characteristics, evolution, and interactions. *Applied Network Science*, *8*(1), 38. https://doi.org/10.1007/s41109-023-00565-4

Ante, L. (2022). The Non-Fungible Token (NFT) Market and Its Relationship with Bitcoin and Ethereum. *FinTech*, *1*(3), 216–224. https://doi.org/10.3390/fintech1030017

Bolz, M., Bründler, K., Kane, L., Patsias, P., Tessendorf, L., Gogol, K., Kim, T., & Tessone, C. (2024). *Machine Learning-Based Detection of Pump-and-Dump Schemes in Real-Time* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2412.18848

Bose, P., Das, D., Gritti, F., Ruaro, N., Kruegel, C., & Vigna, G. (2023). *Exploiting Unfair Advantages: Investigating Opportunistic Trading in the NFT Market* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2310.06844

Chen, S., Chen, J., Yu, J., Luo, X., & Wang, Y. (2024). The Dark Side of NFTs: A Large-Scale Empirical Study of Wash Trading. *Proceedings of the 15th Asia-Pacific Symposium on Internetware*, 447–456. https://doi.org/10.1145/3671016.3674808

Chen, Z., & Omote, K. (2022). Toward Achieving Anonymous NFT Trading. *IEEE Access*, *10*, 130166–130176. https://doi.org/10.1109/ACCESS.2022.3228787

Cho, E., Jensen, G., Yoo, Y., Mahanti, A., & Kim, J.-K. (2024). Characterizing the Initial and Subsequent NFT Sales Market Dynamics: Perspectives From Boom and Slump Periods. *IEEE Access*, *12*, 3638–3658. https://doi.org/10.1109/ACCESS.2023.3333897

Dae-Yong, K., Meryam, E., & Hongtaek, J. (2020). Examining Bitcoin mempools Resemblance Using Jaccard Similarity Index. *2020 21st Asia-Pacific Network Operations and Management Symposium (APNOMS)*, 287–290. https://doi.org/10.23919/APNOMS50412.2020.9237033

Das, D., Bose, P., Ruaro, N., Kruegel, C., & Vigna, G. (2022). Understanding Security Issues in the NFT Ecosystem. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 667–681. https://doi.org/10.1145/3548606.3559342

De Guzman, G. A. (2021). *NFT Marketplaces design impact: Comprehensive analysis of NFT market and ecosystem* [Master's degree thesis]. Politecnico di Milano.

Hufnagel, S., & King, C. (2025). *Criminal Law and Technology: The Complex Case of Non-Fungible Tokens (NFTs)*. SSRN. https://doi.org/10.2139/ssrn.5237723

Jy Tan, L. (2024). NFT Security. In K. Huang, C. Parisi, L. J. Tan, W. Ma, & Z. W. Zhang (Eds), *Web3 Applications Security and New Security Landscape* (pp. 19–34). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-58002-4_2

Kang, H.-J., & Lee, S.-G. (2025). Market Phases and Price Discovery in NFTs: A Deep Learning Approach to Digital Asset Valuation. *Journal of Theoretical and Applied Electronic Commerce Research*, *20*(2), 64. https://doi.org/10.3390/jtaer20020064

Kim, J. S. (2023). *Beyond the Hype: NFT Art and Its Future* [Master of Arts Thesis]. Pratt Institute.

Ko, K., Jeong, T., Woo, J., & Hong, J. W.-K. (2024). Survey on blockchain-based non-fungible tokens: History, technologies, standards, and open challenges. *International Journal of Network Management*, *34*(1), e2245. https://doi.org/10.1002/nem.2245

La Morgia, M., Mei, A., Sassi, F., & Stefa, J. (2020). Pump and Dumps in the Bitcoin Era: Real Time Detection of Cryptocurrency Market Manipulations. *2020 29th International Conference on Computer Communications and Networks (ICCCN)*, 1–9. https://doi.org/10.1109/ICCCN49398.2020.9209660

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A fistful of bitcoins: Characterizing payments among men with no names. *Proceedings of the 2013 Conference on Internet Measurement Conference*, 127–140. https://doi.org/10.1145/2504730.2504747

Mondoh, B. S., Johnson, S. M., Green, M., & Georgopoulos, A. (Aristeidis). (2022). NFT Legal and Regulatory Compliance: Connoisseurship and Critique. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4275613

Morgia, M. L., Mei, A., Mongardini, A. M., & Nemmi, E. N. (2023). A Game of NFTs: Characterizing NFT Wash Trading in the Ethereum Blockchain. *2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)*, 13–24. https://doi.org/10.1109/ICDCS57875.2023.00018

Nabilou, H. (2019). How to regulate bitcoin? Decentralized regulation for a decentralized cryptocurrency. *International Journal of Law and Information Technology*, *27*(3), 266–291. https://doi.org/10.1093/ijlit/eaz008

Nadini, M., Alessandretti, L., Di Giacinto, F., Martino, M., Aiello, L. M., & Baronchelli, A. (2021). Mapping the NFT revolution: Market trends, trade networks, and visual features. *Scientific Reports*, *11*(1), 20902. https://doi.org/10.1038/s41598-021-00053-8

Niu, Y., Li, X., Peng, H., & Li, W. (2024). Unveiling Wash Trading in Popular NFT Markets. *Companion Proceedings of the ACM Web Conference 2024*, 730–733. https://doi.org/10.1145/3589335.3651580

Prakash, I. B., Tiwari, A. K., & Hariharan, U. (2023). Decentralized Metadata Storage for Non-Fungible Token Collections Using Interplanetary File System. *2023 7th International Conference on Electronics, Materials Engineering & Nano-Technology (IEMENTech)*, 1–6. https://doi.org/10.1109/IEMENTech60402.2023.10423404

Rajaei, M. J., & Mahmoud, Q. H. (2023). A Survey on Pump and Dump Detection in the Cryptocurrency Market Using Machine Learning. *Future Internet*, *15*(8), 267. https://doi.org/10.3390/fi15080267

Ron, D., & Shamir, A. (2012). *Quantitative Analysis of the Full Bitcoin Transaction Graph* (No. 2012/584). Cryptology ePrint Archive. https://eprint.iacr.org/2012/584

Saha Roy, S., Das, D., Bose, P., Kruegel, C., Vigna, G., & Nilizadeh, S. (2024). Unveiling the Risks of NFT Promotion Scams. *Proceedings of the International AAAI Conference on Web and Social Media*, *18*, 1367–1380. https://doi.org/10.1609/icwsm.v18i1.31395

Scharfman, J. (2023). Non-Fungible Token (NFT) Fraud. In J. Scharfman, *The Cryptocurrency and Digital Asset Fraud Casebook* (pp. 69–80). Springer International Publishing. https://doi.org/10.1007/978-3-031-23679-2_5

Serneels, S. (2023). Detecting wash trading for nonfungible tokens. *Finance Research Letters*, *52*, 103374. https://doi.org/10.1016/j.frl.2022.103374

Shi, R., Cheng, R., Han, B., Cheng, Y., & Chen, S. (2024). A Closer Look into IPFS: Accessibility, Content, and Performance. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, *8*(2), 1–31. https://doi.org/10.1145/3656015

Sifat, I., Tariq, S. A., & Van Donselaar, D. (2024). Suspicious trading in nonfungible tokens (NFTs). *Information & Management*, *61*(1), 103898. https://doi.org/10.1016/j.im.2023.103898

Sulkis, A. (2024). The Future of Non-Fungible Tokens (NFTs): An Analysis of Regulatory and Compliance Challenges and Opportunities. *Digital Repository of Theses - SSBM Geneva*. https://repository.e-ssbm.com/index.php/rps/article/view/646

Upadhyay, N., & Upadhyay, S. (2025). The dark side of non-fungible tokens: Understanding risks in the NFT marketplace from a fraud triangle perspective. *Financial Innovation*, *11*(1), 62. https://doi.org/10.1186/s40854-024-00684-6

Verma, R., Chandrawanshi, K., Soni, G., Jain, G., Nigam, S., & Jain, N. (2024). Unveiling Security Vulnerabilities in NFTs: A Comprehensive Risk Assessment. *2024 IEEE 4th International Conference on ICT in Business Industry & Government (ICTBIG)*, 1–6. https://doi.org/10.1109/ICTBIG64922.2024.10911117

Volosovych, S., Nezhyva, M., & Napadovskyi, I. (2025). NFT UNDER CONDITIONS OF CRITICAL TRANSFORMATIONS IN THE DIGITAL ASSET MARKET. *Baltic Journal of Economic Studies*, *11*(2), 27–34. https://doi.org/10.30525/2256-0742/2025-11-2-27-34

Von Wachter, V., Jensen, J. R., Regner, F., & Ross, O. (2022). NFT Wash Trading: Quantifying Suspicious Behaviour in NFT Markets. In S. Matsuo, L. Gudgeon, A. Klages-Mundt, D. Perez Hernandez, S. Werner, T. Haines, A. Essex, A. Bracciali, & M. Sala (Eds), *Financial Cryptography and Data Security. FC 2022 International Workshops* (Vol. 13412, pp. 299–311). Springer International Publishing. https://doi.org/10.1007/978-3-031-32415-4_20

Wang, Q., Li, R., Wang, Q., & Chen, S. (2021). *Non-Fungible Token (NFT): Overview, Evaluation, Opportunities and Challenges* (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2105.07447

Wen, X., Wang, Y., Yue, X., Zhu, F., & Zhu, M. (2023). NFTDisk: Visual Detection of Wash Trading in NFT Markets. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–15. https://doi.org/10.1145/3544548.3581466

White, B., Mahanti, A., & Passi, K. (2022). Characterizing the OpenSea NFT Marketplace. *Companion Proceedings of the Web Conference 2022*, 488–496. https://doi.org/10.1145/3487553.3524629

Yang, S., Chen, J., & Zheng, Z. (2023). Definition and Detection of Defects in NFT Smart Contracts. *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*, 373–384. https://doi.org/10.1145/3597926.3598063

Yuan, L., Gao, C., Leung, A., & Ye, Q. (2024). *Unraveling Information Asymmetry in Blockchain-Enabled Nft Marketplaces: The Impact of Rarity Rank on Consumer Behavior*. SSRN. https://doi.org/10.2139/ssrn.4720944

Zarifis, A., & Castro, L. A. (2022). The NFT Purchasing Process and the Challenges to Trust at Each Stage. *Sustainability*, *14*(24), 16482. https://doi.org/10.3390/su142416482

Zhang, H., Zheng, Z., & Mehra, A. (2023). Information Transparency and Market Efficiency in Blockchain-enabled Marketplaces: Role of Traders' Analytical Ability. *ICIS 2023 Proceedings*. https://aisel.aisnet.org/icis2023/blockchain/blockchain/17