How Emerging Technologies Reshape Governance and Market Practices through Disintermediation

A DISSERTATION

SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL

OF THE UNIVERSITY OF MINNESOTA

BY

Agnes Yang

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Dr. Gautam Ray, Advisor

July, 2025

Acknowledgments

To the ones I adore, admire and the ones who inspired me down the road.

First, I want to express my deep gratitude to my academic advisor, **Dr. Gautam Ray**, whom I admire both personally and professionally. He has been a dependable and supportive mentor who helped me endure the many challenges along the way. I once asked him how I could tell if I was doing well as a PhD student, what qualifies as “good enough.” Without hesitation, he offered a piece of wisdom I’ve carried with me ever since: Whatever you do, try your best every day and feel good about yourself. That’s all. From that moment on, I no longer needed validation from elsewhere. That single comment gave a fearful squirrel the guts to keep showing up, even when it felt hard to do so.

I am grateful to **Dr. Joseph Ritter** for his unassuming yet steadfast support on my committee, even long after your retirement. I thank **Dr. Shawn Curley** for his considerate yet straightforward advocacy for students. I leave this school with peace of mind, knowing that our future students are in good hands with him as our coordinator. **Dr. Mani Subramani**, the way he looked me in the eye offering the heartfelt words gave me real comfort when I needed it most. I would also like to thank **Dr. Mochen Yang**, our alumnus and faculty member, who is striving to help others in the most effective ways. I feel lucky to have spent these years learning in the same space with Mochen. Lastly, in moments of hesitation, I found unexpected clarity in the examples set by **Dr. Gwanhoo Lee**, **Dr. Jaehwuen Jung**, **Dr. Antino Kim**, **Dr. Derrick Huang**, **Dr. Idris Adjerid**, and **Dr. Pallab Sanyal**. Taking my first steps into an academic career, encountering such characters reminded me that there's more than one way to thrive in this field.

To my family in Korea, **Ranhee Lee, Jaesook Lee, Youngwoo Kwon, JinJa Kim, Seunghoon Yang, Eunkyung Kwon, Jaewoo Jung**, thank you for always rooting for me across the miles. Your belief in me has been my greatest resource. I'm also grateful to my lifelong friends, **Gayoung Ryu, Sooyoung Ko, and SinHeun Kim** for snapping me out of my self-doubt and pushing me to keep going.

In Minnesota, I was fortunate to befriend beautiful individuals, including **Jieyi Chen, Namji Eum, Ramin Zandvakili, and Nasim Eshgarf**. I look forward to continuing to grow together in the years to come. Special thanks to **Reiko Hirai, Nisha Fernandes and Anne**, extraordinary women who empowered me with wisdom and grace. And, **Rosie**, every moment with you was a joy. You'll always be my favorite corgi in spirit.

Last and most importantly, I dedicate this work to my dearest husband, friend, and colleague, **Youngjin Kwon**. What began as a first-sight crush turned into an endless, sweaty adventure with your clueless but ever-grateful partner. Words fall short of capturing the depth of your presence in my life and the quiet strength I draw from your grounded spirit. I'll stand by you through every challenge, just as you've always been there for me. I look forward to all the chapters ahead in our adventure as the maker and the breaker.

Abstract

This dissertation investigates how emerging technologies such as blockchain disintermediate traditional gatekeepers, thereby transforming governance and market dynamics. Across three essays, I examine how disintermediation reshapes organizational resilience, market access, and participatory governance in blockchain-based systems. The first essay finds that DAOs with more concentrated token ownership are better able to recover from cyber-attacks, suggesting that certain forms of centralization can strengthen collective resilience. The second essay reveals that while NFT art platforms eliminate traditional curators, underrepresented artists still face valuation disparities, yet strategic self-curation can partially reduce these gaps. The third essay explores how the structure and content of DAO forum discussions shape voter turnout, showing that reasoned, inclusive discourse increases participation. Together, the essays argue that disintermediation alone does not guarantee equity or efficiency; rather, the quality of coordination, communication, and curation in decentralized settings becomes crucial. These findings contribute to a deeper understanding of community governance and market transformation in the age of decentralized technologies.

Contents

List of Tables

List of Figures

1 Introduction

1.1 Research Background

Emerging digital technologies such as blockchain and Web3 platforms are radically altering how organizations and markets operate by removing traditional intermediaries. Decentralized systems allow users to connect directly: for example, non-fungible token (NFT) marketplaces give creators "direct access to global audiences without relying on intermediaries such as galleries or auction houses," and decentralized autonomous organizations (DAOs) enable community members to collectively own and control the organization, bypassing traditional centralized leadership structures. In general, digitization lowers search costs and entry barriers and reallocates control away from entrenched gatekeepers. These shifts can democratize access and empower participants in sectors long governed by intermediaries (e.g., art curators or corporate managers), but they also raise new challenges in governance and equity. For instance, while disintermediation promises broader inclusion, new informational frictions may persist in decentralized settings.

The overarching theme of this dissertation is disintermediation; how emerging technologies remove intermediaries and thereby reshape governance and market practices. Each of the three essays examines a different context in which digital platforms restructure relationships among stakeholders. The first essay studies blockchain-based DAOs, which fuse ownership and control through tokenization; it asks how distributing tokens among participants affects a DAO's ability to withstand shocks. The second essay focuses on a decentralized NFT art marketplace, asking whether eliminating traditional art-world gatekeepers affects the market success of underrepresented artists. The third essay analyzes DAO governance forums, exploring how the quality of online discussion influences collective decision-making. Together, these essays illustrate the unifying concept of disintermediation.: blockchain-backed organizations and markets enable peer-to-peer coordination (whether among investors, creators, or voters) and thereby change both organizational governance and market outcomes. In what follows, Section 1.2 summarizes each essay's research question, empirical setting, methods, and key findings, highlighting its contribution to understanding this theme.

1.2 Overview of Three Essays

Essay 1 (DAO Resilience and Token Ownership). The first essay investigates how token ownership structure influences the resilience of blockchain-based organizations. In particular, it examines decentralized autonomous organizations (DAOs) in the decentralized finance (DeFi) sector and asks whether having large token holders, blockholders, helps a DAO recover from adverse events. To study this, we gathered a panel of 115 DAOs and identified 42 cybersecurity attacks that hit these organizations. We employ a staggered difference-in-differences design, using each cyberattack as an exogenous shock to a DAO's governance and market performance. We also augment these data with a comprehensive dataset of 2.9 million governance votes from Snapshot, a representative off-chain voting platform for DAOs. In regression models, token ownership concentration (measured by a Theil's T index) interacts with the attack indicator to predict outcomes such as cumulative abnormal returns.

Our findings show that DAOs with more concentrated token ownership, i.e., more blockholders, are more resilient to shocks. In other words, ownership concentration significantly enhances a DAO's recovery after an attack. Empirically, the interaction term between cyberattack and ownership concentration is positive and significant, implying that a 1-unit increase in token concentration raises post-attack returns by over 10%. In practical terms, DAOs with a few committed blockholders suffer smaller losses when attacked than more diffuse DAOs. Further analysis suggests a mechanism behind this: DAOs with blockholders exhibit higher governance participation in terms of voter turnout. This is consistent with a "voice" channel, where engaged large holders actively steer the organization through the crisis. We also document an inverted-U relationship between ownership concentration and participation, indicating that turnout is highest when ownership is neither too dispersed nor too concentrated. Overall, Essay 1 shows that in DAO governance, a certain level of concentrated stakeholding can bolster collective resilience, a novel insight into how decentralization and blockholder theory interact in practice.

Essay 2 (NFT Markets and Artist Inclusion). The second essay examines the effects of disintermediation in a decentralized art market. This study centers around whether eliminating intermediaries in a curated art market may level the playing field for historically underrepresented artists. Using transaction-level data from SuperRare, one of the most representative curated NFT art market, we analyze around 27,000 sales and 90,000 offers for 40,000 NFTs created by 2,500 artists. The research questions are: (1) Are artworks by female or non-White artists less likely to sell, or do they fetch lower prices, than those by White male artists? (2) Can allowing artists to disclose additional information (a form of self-curation) mitigate any disparities? (3) Which types of information signals are most effective at reducing gaps?

We estimate regression models of sale probability and sale price on artist demographics and control variables. The results reveal persistent disparities. Female and non-White artists face significant disadvantages on the NFT platform: compared to White male artists, non-White artists in particular have a much lower likelihood of sale and lower prices for their work. Notably, the racial gap is even larger than the gender gap. In other words, disintermediated NFT markets are not free of minority discount; the market still mirrors societal prejudices in how artworks are valued. However, we find that these outcome gaps shrink when artists provide credible quality signals. Specifically, information that signals artistic merit (such as verified exhibition history or detailed credentials) significantly attenuates the disadvantages faced by minority artists. For example, including exhibition venue and prior client list in the artist's profile narrows the racial price gap, while narrative descriptions and past commissions help female artists. These findings are consistent with the idea that self-curation on NFT platforms can empower creators. In sum, Essay 2 shows that blockchain-enabled disintermediation expands market access for creators but does not automatically erase gender and racial discount; to secure fair valuation, creators must actively engage in strategic self-curation.

Essay 3 (DAO Discourse and Voter Turnout). My third essay explores how the structure and content of online discourse in DAOs influence democratic participation. It asks: What features of community discussion threads predict higher or lower voter turnout in decentralized governance? To answer this, we assemble the largest linked DAO dataset to date, combining 248k forum posts (pre-vote replies on Discourse) with 11,489 governance proposals and 4.2 million individual votes on Snapshot.

To guide our analysis of DAO community discourse, we introduce a structural-content framework that distinguishes between the structure of interaction, both relational and semantic, and the content of communication, both informational and emotional. The structural dimension includes patterns of participant interaction and thematic convergence, while the content dimension captures the substance and tone of discourse, such as argumentative or cognitive moves. This framework provides a layered view of deliberative processes and serves as the basis for quantifying how discussion quality relates to governance outcomes.

The empirical results highlight the importance of deliberation quality. Threads characterized by semantic alignment and reasoned debate see significantly higher voter turnout.

In particular, when discussion posts converge on a common semantic space and emphasize premise-driven, fact-based arguments, the voter turnout is likely to increase. Moreover, having more active participants, in both terms of the quantity of the responses communicated with others and the relative influence size captured by Pagerank, is associated with higher engagement.

These findings bridge open-source coordination and social media engagement literature, showing that in blockchain governance, the relational and semantic quality of community discourse materially influences subsequent voting participation. Practically, this suggests that DAO designers and moderators should foster clear, consensus-building communication, e.g., by encouraging reasoned, premise-based discussion and distributed interaction, to broaden and deepen voting participation.

In conclusion, each essay together contributes to understanding how digital technologies transform traditional governance and market practices through disintermediation. The first essay reveals how token-based governance reorganizes control in organizations, blurring the classic ownership-control separation, and how that reorganization affects resilience. The second essay demonstrates that eliminating traditional gatekeepers can widen market access, yet underscores the need for new curatorial mechanisms to safeguard equity. The third essay proposes that both the structural dynamics and substantive content of DAO community discussions materially influence subsequent governance participation in decision-making. Collectively, these essays offer a multi-dimensional view of how blockchain technology enables disintermediation in ways earlier technologies never could.

2 The Role of Token Ownership Structure on the Resilience of Decentralized Autonomous Organizations (DAOs)

2.1 Introduction

"DAO is a powerful term that captures many of the hopes and dreams that people have put into the crypto space to build more democratic, resilient and efficient forms of governance."

Vitalik Buterin, Co-founder of Ethereum

How decentralized should power within a collective be to prevent exploitation? Historically, this question has been a key driver behind movements like Technological Utopianism in the 20th century, which advocated for decentralized societies built around internet, but ended as premature visions. Although these efforts were never fully realized, the relevance of decentralized governance has only grown, especially as the centralization by technology giants has led to even more severe forms of exploitation. Centralized digital platforms, such as Google and Facebook, now hold mass influence over access to individual data and pricing, with platforms systematically gathering and analyzing data, which can be used to shape user behavior for their own benefit (Tirole, 2021, Sockin and Xiong, 2023).

Amid this growing concern over centralization, blockchain technologies have envisioned a promising alternative. In particular, as smart contracts [1] have unlocked blockchain's potential beyond the means of payments, decentralized platforms are increasingly seen as a potential counterpart to the current centralized models (Tsoukalas and Falk, 2020). Decentralized platforms have rapidly expanded across sectors, from decentralized finance (DeFi) and non-fungible tokens (NFTs) to social media and marketplaces (Zhao et al.,

[1] Smart contracts are self-executing programs that run on blockchains. They consist of the terms of an agreement written directly into code, ensuring automatic execution when predefined conditions are met. This automation reduces the need for trust among parties, decreasing the risk of fraud and enhancing security.

2022). Many of these platforms are governed by their associated decentralized autonomous organization (DAO), a novel form of organization where members collectively own and control the organization, bypassing traditional centralized leadership structures (Ellinger et al., 2024).

A DAO operates under a distinctive governance structure where members, as a community, collectively initiate, discuss, and reach consensus on proposals for platform improvement (Sockin and Xiong, 2023). By relying on its owners for decision-making, a DAO challenges the traditional norm in large profit-oriented entities, where ownership and control are separated, a principle long established in corporate governance literature (Shleifer and Vishny, 1986, Fama and Jensen, 1983). Even in their early stages and with an unconventional structure, DAOs have demonstrated their strength by overcoming numerous governance challenges and extreme market volatility. As of mid-2024, there are approximately 12,000 DAOs globally, with a network of 10 million governance token holders collectively managing over $21 billion in assets.[2]

One fundamental feature that effectively coordinates DAO owners in community governance is the dual-value tokenization (Ellinger et al., 2024). Governance tokens within DAOs serve a dual purpose: they grant holders the right to participate in decision-making processes while also providing potential monetary gains through token appreciation, typically linked to the success of the platform (Howell et al., 2020a, Sockin and Xiong, 2023). However, this system (re)introduces the classic problem of free-riding, where participants may benefit from the platform's success without actively contributing to its governance—a challenge that traditionally led to the separation of ownership from control in large corporations (Fama and Jensen, 1983, Holderness, 2009, Tsoukalas and Falk, 2020, Beck et al., 2018, Ellinger et al., 2024).

[2] https://deepdao.io/organizations/

In this study, we investigate the pivotal role of token ownership structure in DAO governance through the lens of blockholder theory in corporate governance. In the literature on blockholders, which examines the influence of large shareholders, scholars generally view that their presence benefits corporations by improving governance effectiveness (Shleifer and Vishny, 1986, Edmans and Holderness, 2017). In DAOs, while token-weighted voting resembles shareholder voting, it's unclear whether blockholders benefit governance in the same way as in corporations (Tsoukalas and Falk, 2020). Given greater governance power for DAO token holders and lower participation costs, blockholders may not play as central a role as they do in corporate settings (Murray et al., 2021). Practitioners also raise concerns about the widely observed concentration of token ownership, where a small number of participants hold a significant portion of tokens, potentially undermining the core ethos of decentralization (Sockin and Xiong, 2023).

Against this background, we examine the impact of token ownership concentration on DAO performance, using cyberattacks as exogenous shocks, which pose a significant threat to blockchain-based platforms. Drawing on a comprehensive dataset that includes 42 cyberattacks and 2.9 million governance proposal votes, we assess how varying levels of ownership concentration influence a DAO's ability to recover in the aftermath of cyberattacks. Our findings reveal that DAOs with more blockholders, indicated by higher token ownership concentration, demonstrate greater resilience to cyberattacks. Furthermore, we explore the potential drivers of this relationship, as informed by blockholder theory. Analysis of community participation shows that DAOs with more blockholders have higher participation rates in governance voting, indicating the possible presence of a voice mechanism. In doing so, this research provides valuable insights into DAO governance by exploring the unique role of dual-value tokenization. In Section 2.8, we take a step further by proposing a new perspective on the ongoing discourse around decentralization, a term that often causes confusion, and discuss its theoretical and practical implications.

2.2 Institutional Background

2.2.1 Making Sense of DAOs: Transforming Governance in Decentralized Platforms

While early blockchains like Bitcoin and Litecoin primarily served as the means of payments, the introduction of smart contracts, particularly through Ethereum, has significantly expanded the scope of blockchain applications (Lumineau et al., 2021). Among others, decentralized platforms have emerged as a compelling alternative to existing centralized platform models like Google and Uber (Tsoukalas and Falk, 2020, Sockin and Xiong, 2023). Unlike their centralized counterparts, decentralized platforms offer users the promise of more direct control and, consequently, a larger share of profits (Sockin and Xiong, 2023, Gan et al., 2023). Expanding on a massive scale, decentralized platforms span various sectors from decentralized finance (DeFi), non-fungible tokens (NFTs), to social media and marketplaces (Zhao et al., 2022).

Many of these platforms are governed by their designated community, also known as decentralized autonomous organization (DAO). A DAO is a collectively-owned, autonomously-governed organization managed through smart contracts on a blockchain (Lumineau et al., 2021). A DAO forgoes the principle of separating of ownership from control, which is a modern governance norm among for-profit organizations (Fama and Jensen, 1983, Shleifer and Vishny, 1986). Instead, it integrates ownership and control through token-based ownership, achieving organizational goals without centralized authorities like managers or boards of directors (Ellinger et al., 2024, Appel and Grennan, 2023). As a community, DAO members initiate, discuss, and reach consensus on platform-improvement-protocols, all within the framework of decentralized token ownership (Sockin and Xiong, 2023, Ellinger et al., 2024).

2.2.2 Dual-value Tokenization and Token Ownership Structure

Decentralized platforms often rely on the issuance of cryptocurrencies or tokens that help raise capital for initial development and foster ongoing participation (Gan et al., 2023). Tokens are digital assets operating on a blockchain network that can be broadly categorized into three types: coins, security tokens, and utility tokens (Sockin and Xiong, 2023, Howell et al., 2020a).[3] Coins primarily function as a medium of exchange and a store of value, with Bitcoin being the most prominent example (Hsieh and Vergne, 2022). Security tokens are designed to represent cash flow rights, functioning similarly to traditional securities, with the key difference being that ownership is recorded on a blockchain (Howell et al., 2020a, Sockin and Xiong, 2023).

Lastly, utility tokens, which make up the majority of initial coin offerings, are central to most studies on decentralized platforms, including our own.[4] By issuing utility tokens, a platform grants users consumptive rights to its services and potential governance rights (Gan et al., 2023). In addition, as the governance proceeds smoothly, token holders may benefit from increased token value, which often rises with the overall value of the decentralized platform (Howell et al., 2020a). This dual-value tokenization incentivizes the users to actively engage in governance, highlighting the crucial role of governance tokens in understanding the unique structure of DAO governance. (Ellinger et al., 2024).

Nevertheless, the design of dual-value tokenization comes with innate challenges. Since every holder benefits from token appreciation resulting from good decision-making, a risk

[3]These three categories are not always mutually exclusive. For instance, Ether (the token of the Ethereum network) falls into both the first and third categories, while USDT (a stablecoin issued by Tether Limited Inc.) fits into both the first and second categories. For a detailed taxonomy of tokens, see (Howell et al., 2020a).

[4]See Howell et al. (2020a), Sockin and Xiong (2023), Appel and Grennan (2023), Gan et al. (2023), Tsoukalas and Falk (2020), Zhao et al. (2022), Ellinger et al. (2024)

of free-riding emerges, where diffuse joint owners lack the same level of incentive that a few large owners would have (Tsoukalas and Falk, 2020, Holderness, 2009). Conversely, practitioners seem rather concerned about the high token ownership concentration observed in DAOs, which could not only lead to decision-making dominance by a few large holders but also undermine the core vision of decentralization (Sockin and Xiong, 2023, Ellinger et al., 2024).[5] All in all, the impact of ownership concentration on DAO governance and performance, whether it is detrimental, and the reasons behind it, remains unclear, forming the central focus of this study. In the following sections, we contextualize our study by outlining typical decision-making processes in DAOs and by providing key industry background related to cyberattacks.

2.2.3 Contextualizing DAO Governance

Historically, the rise of DAOs has been driven by the increasing demand for decentralized governance, spurred by the massive growth of decentralized finance (DeFi) protocols, alternative financial service providers operating without traditional intermediaries (Zhao et al., 2022).[6] While DAOs can be established for virtually any objective, such as operating social network platforms, DeFi DAOs still represent the majority, and thus form the focus of this paper (Appel and Grennan, 2023).[7] DeFi DAOs, like MakerDAO and Uniswap, govern their protocols by allowing their governance token holders to vote on key decisions such as adjusting interest rates, adding new financial products, or updating governance rules (Zhao et al., 2022, Gan et al., 2023). The governance token not only serves as proof of governance power but also holds monetary value, similar to a share price, illustrating the

[5] See https://vitalik.eth.limo/general/

[6] DeFi protocols are decentralized financial platforms that democratize access to financial services by eliminating intermediaries, such as banks and brokers. They support various financial activities, including peer-to-peer lending and borrowing, stablecoins, and staking digital assets.

[7] The boundary condition of our study is DAOs governing DeFi protocols on public, permissionless blockchains, such as Ethereum, ensuring a coherent and focused discussion on DAO governance.

dual-value nature discussed in Section 2.2.2.

A decision making process in DAOs often begins with community members initiating discussions about potential improvements or strategic changes on a forum or Discord chat (Appel and Grennan, 2023). These discussions, covering a wide range of topics from daily operations to long-term directions, allow members to voice concerns and collaboratively develop ideas.

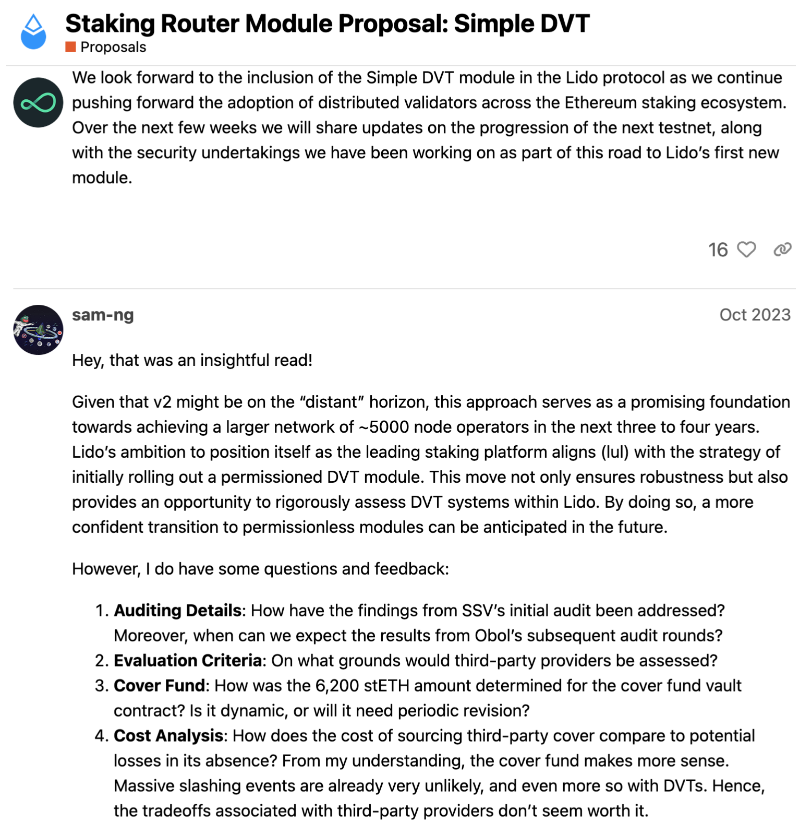
Staking Router Module Proposal: Simple DVT
Proposals

We look forward to the inclusion of the Simple DVT module in the Lido protocol as we continue pushing forward the adoption of distributed validators across the Ethereum staking ecosystem. Over the next few weeks we will share updates on the progression of the next testnet, along with the security undertakings we have been working on as part of this road to Lido's first new module.

16

sam-ng Oct 2023

Hey, that was an insightful read!

Given that v2 might be on the "distant" horizon, this approach serves as a promising foundation towards achieving a larger network of ~5000 node operators in the next three to four years. Lido's ambition to position itself as the leading staking platform aligns (lul) with the strategy of initially rolling out a permissioned DVT module. This move not only ensures robustness but also provides an opportunity to rigorously assess DVT systems within Lido. By doing so, a more confident transition to permissionless modules can be anticipated in the future.

However, I do have some questions and feedback:

1. **Auditing Details**: How have the findings from SSV's initial audit been addressed? Moreover, when can we expect the results from Obol's subsequent audit rounds?
2. **Evaluation Criteria**: On what grounds would third-party providers be assessed?
3. **Cover Fund**: How was the 6,200 stETH amount determined for the cover fund vault contract? Is it dynamic, or will it need periodic revision?
4. **Cost Analysis**: How does the cost of sourcing third-party cover compare to potential losses in its absence? From my understanding, the cover fund makes more sense. Massive slashing events are already very unlikely, and even more so with DVTs. Hence, the tradeoffs associated with third-party providers don't seem worth it.

Note: Screenshot from the Lido DAO governance forum showing a discussion thread on a proposal.

Figure 2.1: Screenshot of Governance Forum

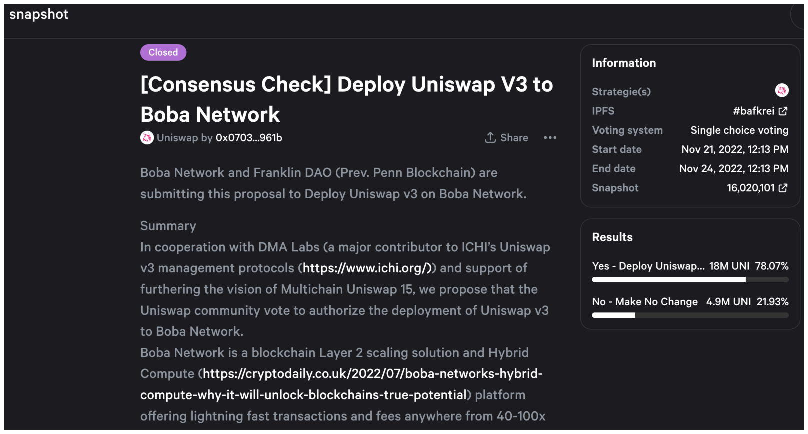

(a) Snapshot Proposal Example

(b) Voters on Snapshot Proposal

Note: Screenshots from the Snapshot governance platform. Panel(a) shows a sample proposal; Panel(b) lists the addresses that cast votes on that proposal.

Figure 2.2: Screenshot of Snapshot Proposal

Figure 2.1 depicts an example of informal discussion on a governance forum. Once a discussion gains significant attention from the community, it may progress to a formal proposal stage on voting platforms such as Snapshot, a widely used tool for governance voting.[8] For most DAOs, the voting period typically lasts between 3 and 7 days, with

[8]The decision-making processes via DAO voting analyzed in this study pertain to off-chain voting, though technically, DAO voting can occur either on- or off-chain. In on-chain voting, votes are recorded on

token holders casting votes in proportion to the size of their holdings. The results are then set to be executed through smart contracts, ensuring that decisions are carried out transparently and immutably (Zhao et al., 2022). Figure 2.2 showcases Snapshot voting. Taken together, the active engagement of token holders in both initial discussions and formal voting highlights the bottom-up nature of community governance, where agendas are identified at the grassroots level and decisions are finalized through collective participation (Ellinger et al., 2024).

2.2.4 Cybersecurity Attacks as Governance Challenges

Cybersecurity attacks are seen as one of the greatest threats to the operation of decentralized platforms (Cong et al., 2024). In particular, vulnerability attacks, where hackers exploit weaknesses or flaws in software, can cause significant damage to DeFi protocols due to the massive scale of assets involved and the reliance on code within a fully transparent virtual environment (Ziolkowski et al., 2020). The critical role of cyberattacks in DeFi DAOs' operations is well-supported by industry statistics and practices. In 2022 alone, $3.8 billion was stolen from cryptocurrency businesses, with approximately 82% ($3.1 billion) lost in DeFi protocols.[9] Given the severe consequences, most DAOs proactively invest substantial resources in preventive measures like audits and bug bounty programs[10] to deter potential breaches (Rabetti, 2024, Cong et al., 2024).

Nevertheless, as complete protection against breaches is unattainable, a DAO's ability to recover swiftly following an incident is of critical importance (Ziolkowski et al., 2020).

the blockchain, incurring transaction (i.e., gas) fees to the voters. To avoid these gas fees, the majority of DAOs, including those in our study, conduct off-chain voting, where token holders can vote on proposals without making blockchain transactions.

[9] See https://www.chainalysis.com/blog/.

[10] Audits are scheduled, structured processes where external professionals conduct a code review, aiming to identify and rectify vulnerabilities before they can be exploited. Regular audits are one of the key signals that the platform meets a certain standard of security and compliance. Bug bounty programs offer financial rewards to white-hat hackers who identify and report potential vulnerabilities in the code.

Effective crisis response requires the presence of well-functioning governance mechanisms at the time of the attack (Gwebu et al., 2018). Moreover, broad stakeholder trust and confidence in the DAO's resilience are essential to prevent mass sell-offs of governance tokens and premature exits (Kaczynski and Kominers, 2021). Below, we illustrate a detailed case of a DAO that experienced a significant breach.

In December 2021, BadgerDAO, governing a DeFi platform named Badger Finance, suffered a $130 million loss due to a front-end exploit.[11] This loss accounted for over 10% of the total value locked in Badger Finance at the time of the attack, exceeding twice the collective funds managed by its DAO. Unsurprisingly, it posed a significant threat to the platform's future. To navigate the crisis, BadgerDAO implemented a comprehensive restitution plan, engaging in extensive discussions within community forums and voting on six Badger Improvement Proposals (BIPs), held just two weeks after the initial exploit. The success of their recovery depended not only on the timely collaboration of 32,000 globally distributed contributors but also on broader stakeholder confidence in the DAO's ability to rebound. This underscores how effective governance and strong community cohesion are essential for a DAO to withstand crises. A core contributor of BadgerDAO later reflected on their recovery process, emphasizing the importance of a well-established and unified community in navigating challenges:

> *"You're talking about people who have never met, from around the world, going through such a dramatic thing as a collective, then mustering up the strength to respond in such a high integrity way."*[12]

To summarize, in this decentralized ecosystem, cyberattacks can serve as a litmus

[11] A front-end exploit occurs when attackers inject malicious code to deceive users into authorizing unauthorized transactions on smart contracts.

[12] https://coindesk.com/tech/2021/12/16/

test for community trust and the effectiveness of internal governance. Given the importance of token ownership structure in shaping a community's composition, examining DAO performance in the aftermath of cyberattacks provides valuable insights into how ownership structure influences community participation and, in turn, overall performance. This perspective underscores how governance structures and stakeholder engagement shape a DAO's ability to weather crises and maintain long-term stability (Murray et al., 2021).

2.3 Hypothesis Development

2.3.1 DAO as a New Form of Organization

Several academic efforts have sought to understand the governance of blockchain-enabled organizations by comparing them to traditional counterparts such as corporations and open-source communities (Yermack, 2017, Murray et al., 2021). Despite some similarities, organizational scholars increasingly agree that DAOs are not a derivative of existing organizations, but a novel form of organizing in digital commons (Ellinger et al., 2024, Lumineau et al., 2021, Beck et al., 2018, Hsieh and Vergne, 2022, Zhao et al., 2022). One of the earliest works on blockchain governance, Yermack (2017), examines the potential implications of blockchain technology for corporate governance. Notably, the token-weighted voting in DAOs resembles the shareholder voting system in publicly-held corporations, where one share typically grants one vote, and shareholders hold varying amounts of shares (Catalini and Boslego, 2019, Tsoukalas and Falk, 2020).

Even so, a critical distinction between DAOs and corporations lies in their decision-making structure. In corporations, final decision-making authority rests with central managers, even though shareholders can cast votes (Edmans and Holderness, 2017), contrasting sharply with DAOs, where token holders directly influence decision-making without central authority. As a result, DAO token holders bear significantly more responsibility and power than shareholders (Tsoukalas and Falk, 2020, Zhao et al., 2022). Moreover, the motivations driving participation in DAOs are more varied than the sole focus on monetary rewards seen in traditional shareholder models (Edmans and Holderness, 2017). DAO members often participate for a range of reasons, including building personal reputation and supporting a shared vision (Ellinger et al., 2024).

Similarly, their dependence on voluntary contributions without central oversight can make DAOs appear as variations of open-source projects (Hsieh and Vergne, 2022, Ziolkowski et al., 2020). Better yet, DAOs possess a key feature that addresses one of the biggest drawbacks of typical open-source communities: offering monetary incentives to contributors, fostering sustained participation (Hsieh and Vergne, 2022). Taken together, the unique combination of collective ownership and voluntary contribution underscores why DAOs are increasingly viewed as a novel phenomenon in governance. This distinct nature largely arises from dual-value tokenization, as discussed in Section 2.2.2.

2.3.2 Token Ownership Structure and the Blockholder Perspective

Although DAOs represent a novel form of organizing, their decision-making processes, particularly token-weighted voting, share structural similarities with traditional corporate governance (Lumineau et al., 2021, Tsoukalas and Falk, 2020). In both contexts, owners vote in proportion to their holdings, and a small group often holds a disproportionately large share of voting power (Edmans and Holderness, 2017, Ellinger et al., 2024). As such, blockholder theory, which examines how large shareholders shape governance outcomes, offers a useful starting point for understanding DAO governance dynamics.

In corporate governance literature, blockholders are widely seen as beneficial. They have the financial incentive and power to monitor management, intervene when necessary, and mitigate agency problems (Shleifer and Vishny, 1986). Conversely, when ownership is highly diffuse, governance often suffers from free-riding: individual shareholders have little reason to participate, leading to free-riding and weak oversight (Holderness, 2009). However, DAOs also differ from corporations in important ways, making it unclear whether the same governance logic applies. On one hand, DAOs are explicitly designed to support decentralized decision-making, with blockchain infrastructure significantly lowering coordination costs. On the other hand, the absence of centralized management means that governance relies entirely on token holder participation. If diffuse ownership leads to free-riding in traditional firms, the problem may be even more severe in DAOs, where collective action is essential for basic functionality (Ellinger et al., 2024).

This tension remains largely unresolved in the literature. Despite the importance of this problem for the sustainability of decentralized platforms, little empirical work has examined whether ownership concentration supports or undermines DAO performance. We address this gap by building on blockholder theory to explore how token ownership structure may shape governance outcomes, especially in the face of disruptions.

2.3.3 The Role of Blockholders: Governance through Voice and Exit

To examine how ownership concentration might influence DAO resilience, we turn to two foundational mechanisms in blockholder theory: voice and exit (Edmans and Holderness, 2017). In the voice (intervention) channel, blockholders can actively engage in decision making, such as voting on proposals, initiating governance discussions, and steering management direction (McCahery et al., 2016). Blockholders are incentivized to engage in governance because their substantial stakes mean the expected financial benefits outweigh the associated costs, such as resource allocation and educational expenses (Shleifer and Vishny, 1986, Edmans and Manso, 2011). Conversely, firms lacking blockholders often encounter a challenge known as the free-rider problem, where numerous small shareholders have limited motivation for governance participation (Shleifer and Vishny, 1986, Holderness, 2009). In summary, the disparity in costs and benefits of governance participation results in only a handful of large blockholders exerting the voice mechanism (Edmans and Manso, 2011).

DAOs reshape the cost-benefit landscape of participation. First, the benefits of governance engagement are more direct. Token holders can influence binding decisions via smart contracts, and any improvements in platform performance may directly raise token value (Ellinger et al., 2024). In corporations, while shareholder voting exists, the final decision-making authority typically remains with professional managers, who may be encouraged, but are not obligated, to follow shareholder recommendations (Edmans and Manso, 2011, Fama and Jensen, 1983). Moreover, smart contracts ensure the automatic execution of these decisions, enhancing transparency and reducing the risk of exploitation by intermediaries such as managers (Edmans and Holderness, 2017). This transparency ensures that the benefits of sound governance are not diluted, giving token holders a more direct and significant influence in DAOs, aligned with community interests.

Second, participating in governance is less costly in DAOs due to reduced coordination costs (Catalini and Boslego, 2019). Online governance forums and off-chain voting platforms like Snapshot reduce coordination barriers, making participation faster and less resource-intensive than traditional shareholder meetings (Ellinger et al., 2024). This structure may also motivate mid-sized holders, not just large ones, potentially mitigating free-riding more effectively than in firms.

In summary, DAO token holders are likely to experience higher expected benefits and lower expected costs associated with governance participation compared to traditional corporate shareholders. As the expected net payoff, i.e., benefits minus costs, rises in DAOs, it becomes plausible that not only blockholders but also smaller token holders may be motivated to participate, potentially mitigating the free-rider problem. As a result, the role of blockholders in DAOs may not be as significant as it is in corporate governance. Nevertheless, it is still important to note that blockholders may retain a structural advantage over smaller holders, unless participation costs are zero.

The second avenue through which blockholders can influence governance is the exit (trading) mechanism (Edmans and Manso, 2011). In traditional firms, blockholders can discipline management by threatening to sell their shares, thereby depressing stock prices and signaling discontent (Edmans and Holderness, 2017). In this sense, the threat of exit is a more precise description of the mechanism at work, rather than the actual act of exiting. The stronger the exit threat, the more likely the manager is to make efforts to reduce the chances of an actual exit (Edmans and Holderness, 2017). In addition to its direct impact on managers, the exit mechanism can also influence other blockholders, particularly in firms with multiple blockholders (Edmans and Manso, 2011). Blockholders can monitor each other, and in some cases, a subset of blockholders may sustain a cooperative agreement in their monitoring activities to prevent a single or a few dominant blockholders from expropriating firm value at the expense of others (Pagano and Röell, 1998). While DAOs operate without central managers, the potential exit of one blockholder can still serve as a critical threat to other holders. A significant amount of exit trading can substantially reduce token prices, acting as a punitive measure for other blockholders. Therefore, blockholders are incentivized to optimize the DAO's performance through their monitoring efforts, whether individually or collaboratively, leveraging their substantial holdings as a means of an exit strategy (Pagano and Röell, 1998).

Taken together, these mechanisms suggest that concentrated ownership in DAOs may serve as a governance resource. Blockholders' incentives, capabilities, and visibility position them to stabilize DAO governance, especially in high-stakes scenarios such as cyberattacks. In the next section, we connect these ideas to the context of organizational resilience.

2.3.4 DAO Vulnerabilities and Their Effects on Performance

Cybersecurity has long been recognized as a major risk to firm operations, due to potential harms such as remediation costs, legal liabilities, and reputational damage from reduced stakeholder trust (Gwebu et al., 2018). However, recent studies suggest that the performance effects of cyberattacks may be less significant or persistent than previously thought. For example, using data from 2005–2018, Richardson et al. (2019) find that most firms show little to no change in revenue, sales growth, or return on assets. Likewise, Foerderer and Schuetz (2022) argue that firms' strategic timing of breach disclosures, bundling them with other negative news, may account for the mixed findings in earlier research.

As detailed in 2.2.4, cyberattacks pose even greater threats to decentralized platforms, which operate entirely in a virtual space where code is law and the root of trust (Ziolkowski et al., 2020, Cong et al., 2024). While the open-source nature of blockchain offers transparency benefits, it also makes them susceptible to exploitation, demanding extra attention to security to deter breaches ((Hsieh and Vergne, 2022)). Moreover, cyberattacks on blockchain networks are inherently transparent, revealing attack details, such as attack types, methods as the attack unfolds (Jung, 2019). As such, this prevents decentralized platforms from controlling the timing of attack disclosures, necessarily triggering immediate market reactions. As a result, DAOs may suffer negative consequences from breaches, as cyberattacks undermine the trust that forms their foundation while the transparency ensures immediate market disclosure. Therefore, we hypothesize that cyberattacks will significantly weaken DAOs' financial performance over time.

Hypothesis 1. *Cyberattacks have a negative impact on short-term and long-term performance of DAOs.*

No matter how extensively firms strive to enhance the security, complete prevention of cyberattacks is impossible since attackers are constantly evolving their tactics and exploiting unforeseen weaknesses in even the most secure systems (Ziolkowski et al., 2020, Kamiya et al., 2021). As such, beyond investing in security measures to prevent breaches, being adequately prepared for the aftermath of such incidents is critical for firms to weather crises and recover swiftly (Gwebu et al., 2018). Maintaining strong internal governance is one key element of such preparedness, serving as a vital intangible resource that strengthens a firm's resilience to financial crises (Ray et al., 2004, Albuquerque et al., 2020).

In the context of DAOs, operating with radical transparency, breaches are visible to the entire community in real-time. This transparency magnifies the repercussions, eroding trust, depressing token prices, and triggering community-wide uncertainty before any recovery strategy can be mobilized. As a result, the ability of a DAO to recover from such disruptions may vary significantly depending on the strength and preparedness of its governance. Strong governance in DAOs is not just about having decision-making structures in place; it involves engaged stakeholders who are ready to collaborate quickly to mitigate the impact (Ellinger et al., 2024, Hsieh and Vergne, 2022). Blockholders, with their larger stakes and more vested interests, may play a stabilizing role here. Their consistent participation provides institutional memory and continuity, which are essential for coordinating collective action and restoring trust (Gwebu et al., 2018).

Moreover, DAOs with strong governance structures, where decision-making is clear, transparent, and trusted, are better positioned to act swiftly in the aftermath of a breach. This can include rapid deployment of recovery measures, clear communication with the community, and coordinated efforts to rebuild confidence in the platform (Gwebu et al., 2018, Kaczynski and Kominers, 2021). Blockholders are more likely to mobilize these resources and lead efforts to restore stability, as they are both financially and reputationally invested in the platform's long-term success (Ellinger et al., 2024).

Taken together, in the case of cyberattacks, the strength of governance becomes paramount. A well-established governance framework, characterized by engaged and informed stakeholders, can mitigate the free-rider problem and mobilize action quickly. In contrast, a DAO with weak governance or dispersed participation may struggle to organize an effective recovery, prolonging the crisis and leading to a slower recovery process. Blockholders, by virtue of their concentrated stakes and vested interests, are well-positioned to initiate and coordinate governance responses when disruptions occur. With token ownership structure serving as a key indicator of governance strength in DAOs, we hypothesize that DAOs with a more concentrated and engaged ownership base will be better positioned to recover after cyberattacks.

Hypothesis 2. *The presence of blockholders is linked to stronger DAO performance; specifically, greater ownership concentration is associated with stronger post-attack recovery.*

2.4 Data and Measurements

2.4.1 Data

Table 2.1: Variable Descriptions

Variable	**Description**
Dependent Variables	*Source*:CoinGeckoAPI
$\text{CAR}_{i,t}$	Cumulative abnormal return: cumulative sum of the differences between the daily token-price return and its expected return for DAO i on day t (see Section2.4.2).
Daily $\text{Return}_{i,t}$	Daily percentage change in DAO i's token price from day $t-1$ to day t.
Ownership Concentration	*Source*:IntoTheBlock
Ownership $\text{Concentration}_{i,t}$	Dynamic token-ownership concentration measured by Theil's T index (Section2.4.2).
Ownership Concentration_{i}	Static token-ownership concentration measured by Theil's T index.
DAO Profile	*Source*:IntoTheBlock
Transaction $\text{Volume}_{i,t}$	Aggregate on-chain transaction volume (USD) for DAO i on day t.
# of Token $\text{Holders}_{i,t}$	Total number of addresses holding DAO i's token on day t.
Active Holder $\text{Ratio}_{i,t}$	Ratio of active addresses to all addresses with a balance for DAO i on day t.
Average Transaction $\text{Fee}_{i,t}$	Mean on-chain transaction fee (USD) for DAO i on day t.
Voting Participation	*Source*:Snapshot GraphQL API
% of $\text{Voters}_{i,p}$	Number of voters in proposal p divided by the number of token holders at the proposal's close.
% of Tokens $\text{Voted}_{i,p}$	Number of tokens cast in proposal p divided by total circulating tokens at the proposal's close.

Note: All variables are measured at the DAO–day level except the voting-participation variables, which are at the DAO–proposal level.

We compiled a dataset by aggregating information from six platforms. To ensure a comprehensive coverage of platforms, we began with the complete list of DeFi platforms whose tokens are listed on CoinGecko [13] and supplemented it with data from CoinMarketCap [14], both reputable crypto data aggregators. We identified 115 DAOs operating between

[13] https://www.coingecko.com/
[14] https://coinmarketcap.com/

June 2018 and September 2022, with market capitalizations ranging from $1.6millionto4.7$ billion as of this writing. Next, we collected daily DAO profile data from IntoTheBlock [15], a source providing DAO-related financial, performance, and ownership distribution information. Daily historical DAO token prices were obtained from the CoinGecko API [16]. Definitions and summary statistics for each variable are provided in Table 3.1 and Table 2.2, respectively. We proceeded to compile a list of cyberattack incidents in DAOs by aggregating incidents from multiple reputable sources including SlowMist, Rekt, and Chainsec[17], some of which offer detailed summaries of each event, including the date of the incident, size of damage, type of incidents, and sources. For those without the incident summary, we conducted a manual search for detailed information and finalized the list. We identified 42 hacking incidents for DAOs during the observation period, which will be detailed in Table 2.3.

Table 2.2: Summary Statistics

Variable	*N*	**Mean**	**S.D.**	**Min.**	**Max.**
Dependent Variables					
$\text{CAR}_{i,t}$	10 343	4.23	35.15	–55.31	534.96
Daily $\text{Return}_{i,t}$	18 773	0.00	0.06	–0.20	0.24
Ownership Concentration					
Ownership $\text{Concentration}_{i,t}$	10 343	6.42	1.15	0.00	8.52
Ownership Concentration_{i}	10 343	6.34	1.10	3.63	8.49
DAO Profile					
Transaction $\text{Volume}_{i,t}$	10 343	40 356 609.9	103 462 582	0	662 759 808
# of Token $\text{Holders}_{i,t}$	10 343	93 977.04	302 454.6	2	1 514 412
Active Holder $\text{Ratio}_{i,t}$	10 343	0.04	0.06	0.00	0.32
Average Transaction $\text{Fee}_{i,t}$	10 343	0.01	0.00	0.00	0.02
Voting Participation					
% of $\text{Voters}_{i,p}$	2 905	2.67	0.06	0.01	66.19
% of Tokens $\text{Voted}_{i,p}$	2 905	3.25	0.05	0.01	95.79

Note: All variables are at the DAO–day level, except voting-participation variables, which are measured at the DAO–proposal level.

[15] https://app.intotheblock.com/

[16] https://www.coingecko.com/en/api

[17] https://hacked.slowmist.io/en/,https://rekt.eth.link/leaderboard/, https://chainsec.io/defi-hacks/

2.4.2 Variables and Measures

Performance and Resilience

Research has established that financial performance serves as a reliable indicator of resilience (Albuquerque et al., 2020). In a DAO's context, resilience is reflected in its ability to recover from financial losses incurred due to a cyberattack within a defined period. The token price of a DAO is regarded as an indicator of its long-term development prospects and its effectiveness in fulfilling monetary and non-monetary goals (Zhao et al., 2022, Appel and Grennan, 2023). Therefore, a drop in daily abnormal returns following a cyberattack is interpreted as a performance loss, while platforms that incur smaller losses and recover more quickly are considered more resilient. This perspective allows for a nuanced understanding of resilience, aligning financial performance with a DAO's ability to withstand and recover from disruptive events, particularly in the evolving landscape of DAOs.

We gauge DAO-level resilience through the lens of daily risk-adjusted abnormal returns, in line with Albuquerque et al. (2020), which explored the impact of the Covid-19 pandemic on stock market crashes and examined whether firms with higher environmental and social ratings experienced superior returns. Specifically, we investigate how the concentration of token ownership moderates the impact of the cybersecurity attacks on performance. Our underlying assumption is that the market responds differently to cyberattacks depending on the degree of ownership distribution within a DAO, thereby influencing the DAO's resilience. To quantify this, we analyze daily risk-adjusted abnormal returns.

The formula for abnormal returns is derived within the framework of the Capital Asset Market Pricing Model (CAPM). In CAPM, the expected return ($Return_{expected}$) of an asset is expressed as a linear function of the risk-free rate (R_f) and the asset's sensitivity to the market movements (β) multiplied by the market risk premium ($M_t - R_f$) where R_f

is the risk-free rate; β is the asset's beta representing its sensitivity to market movements, and M_t is the market return at time t. The abnormal return (AR_{market}) is then defined as the difference between the observed return ($Return_{observed}$) and the expected return ($Return_{expected}$):

$$\text{Return}_{expected} = \text{R}_f + \beta \cdot (\text{M}_t - \text{R}_f) \tag{2.1}$$

$$\text{AR}_{market} = \text{Return}_{observed} - \text{Return}_{expected} \tag{2.2}$$

This formulation enables the evaluation of an asset's performance beyond what is predicted by its beta and market conditions. We estimate this model by regressing our sample DAO's daily logarithmic return on daily logarithmic return of the market, where the estimation window spans from 60 to 30 days before each incident. As to the market index, we utilized the daily logarithmic return of the top 100 governance tokens, ranked by total market valuation, providing the overall market return similar to S&P 500 in the stock market.

Ownership structure

We measure ownership concentration in DAOs using Theil's T index, a widely used measure of economic inequality. It captures disparity through an entropy-based approach, ranging from 0 (perfect equality) to $\ln(N)$ (maximum inequality). Mathematically, it is expressed as:

$$\text{Theil}_{i,t} = \sum_{h=1}^{N} w_h \ln w_h \tag{2.3}$$

where w_h represents the share of tokens held by holder h, computed as the number of tokens held by h divided by the total token supply of DAO i on day t. N is the total number of token holders. A low Theil's T index indicates an equal distribution of tokens, suggesting broad governance power among the token holders. Conversely, a high Theil's T index points to an unequal distribution, indicating that governance power is concentrated among a few token holders. Theil's T index is highly sensitive to changes in the distribution of ownership shares, particularly at different points in the distribution. Unlike the Gini coefficient, which measures overall inequality but is less sensitive to changes at the extremes of the distribution, Theil's T index can capture shifts in inequality more precisely[18]. Figure 2.3 visualizes the distribution of ownership concentration measured by Theil's T index. Given that the average Theil's T index value is about 6, we confirm that the overall DAO's ownership distribution is significantly concentrated.

We construct two types of ownership concentration measures. Ownership Concentration$_{i,t}$, is a dynamic measure that changes daily, thus operating at the DAO-day level. The second, Ownership Concentration$_i$, is a static measure representing the overall extent of ownership concentration, calculated as the aggregated average value of the ownership concentration measure one month before each incident for each DAO. The motivation for constructing the static measure lies in the potential for ownership concentration to change following a cyberattack due to the entrance of new token holders, the exit or changes in the holding size of current holders.

[18]The conceptual differences are documented in Appendix.

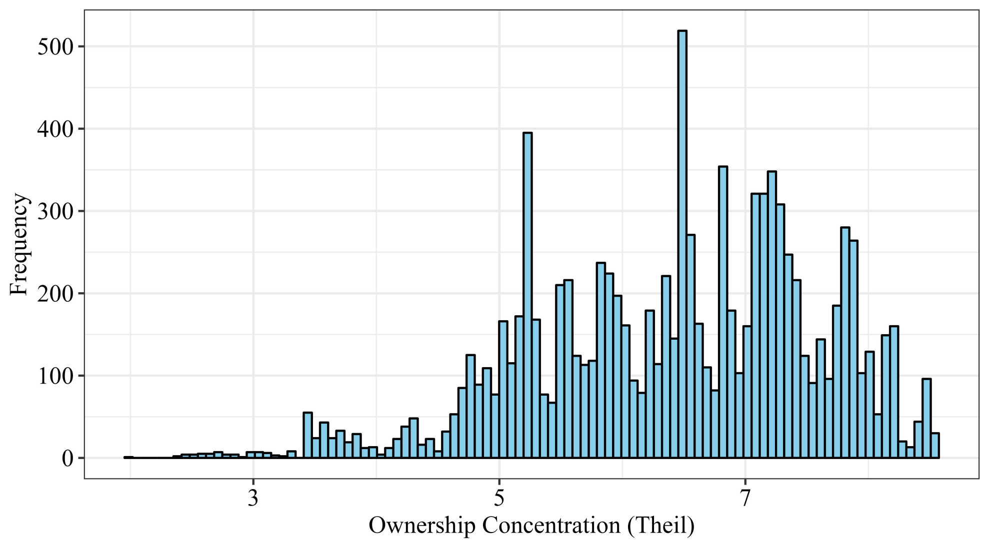

Note: Distribution of ownership concentration (Theil index) across DAOs in the sample.

Figure 2.3: Distribution of Ownership Concentration

2.5 Identification Strategy

To examine the role of DAO ownership structure on DAO resilience, we analyze 42 significant cyberattack incidents on DAOs, treating these incidents as destructive exogenous shocks. We employ a staggered Difference-in-Difference (DiD) approach (Autor, 2003, Angrist and Pischke, 2008, Burtch et al., 2018). Table 2.3 summarizes these DAO cyberattack incidents, including the name of the DAO, its ticker in InToTheBlock, and the incident date. The vast majority of these DAOs operate within the Ethereum ecosystem, with some on the Polygon and Binance Smart Chain networks. The most devastating incident in terms of damage in our list was the Badger DAO attack on December 2, 2021. According to their report [19], BadgerDAO lost $120 million due to a front-end (user interface) attack,

[19] CoinDesk (Decemeber 2, 2021). URL

resulting in a 21% decrease in the BADGER token's value within 24 hours after the attack. Figure 2.4 visually represents the overall change in cumulative abnormal returns of our sample DAOs before and after the incidents. As evident in the graph, DAO tokens' abnormal returns exhibit a notable decline after hack events.

Table 2.3: List of DAO Cyberattacks

DAO Name	**Ticker**	**Incident Date**
Bancor	bnt	2018-07-10
Fusion	fsn	2019-09-28
NULS	nuls	2019-12-23
MakerDAO	mkr	2020-03-12
Loopring	lrc	2020-05-07
Bancor	bnt	2020-06-18
Curve DAO	crv	2020-10-11
Value Liquidity	value	2020-11-14
Pickle Finance	pickle	2020-11-22
Compound	comp	2020-11-26
Rari	rgt	2020-11-30
Sushi	sushi	2021-01-27
yearn.finance	yfi	2021-02-05
Alpha Finance	alpha	2021-02-13
Cream	cream	2021-02-13
ARMOR	armor	2021-02-28
Curve DAO	crv	2021-03-05
EasyFi	easy	2021-04-20
Value Liquidity	value	2021-05-05
Alchemix	alcx	2021-06-16
Popsicle Finance	ice	2021-08-04
Cream	cream	2021-08-31
Poly Network	pnt	2021-09-20
Compound	comp	2021-09-30
Lido DAO	ldo	2021-10-05
Indexed Finance	ndx	2021-10-15
Alpha Finance	alpha	2021-10-23
Cream	cream	2021-10-27
bZx	bzrx	2021-11-05

Continued on next page

DAO Name	Ticker	Incident Date
		Continued from previous page
Curve DAO	crv	2021-11-11
OlympusDAO	ohm	2021-11-23
dYdX	dydx	2021-11-27
Badger DAO	badger	2021-12-02
Gelato Network	gel	2021-12-11
Rari	rgt	2022-01-15
Fei Protocol	fei	2022-04-30
Quickswap	quick	2022-05-14
Ribbon Finance	rbn	2022-06-24
Uniswap	uni	2022-07-11
Curve DAO	crv	2022-08-09
Celer	celr	2022-08-18
Kyber Network Crystal	knc	2022-09-02

Note: Only the incident dates are shown here; loss amounts appear in Table 2.3. Dates are formatted as YYYY-MM-DD for consistency.

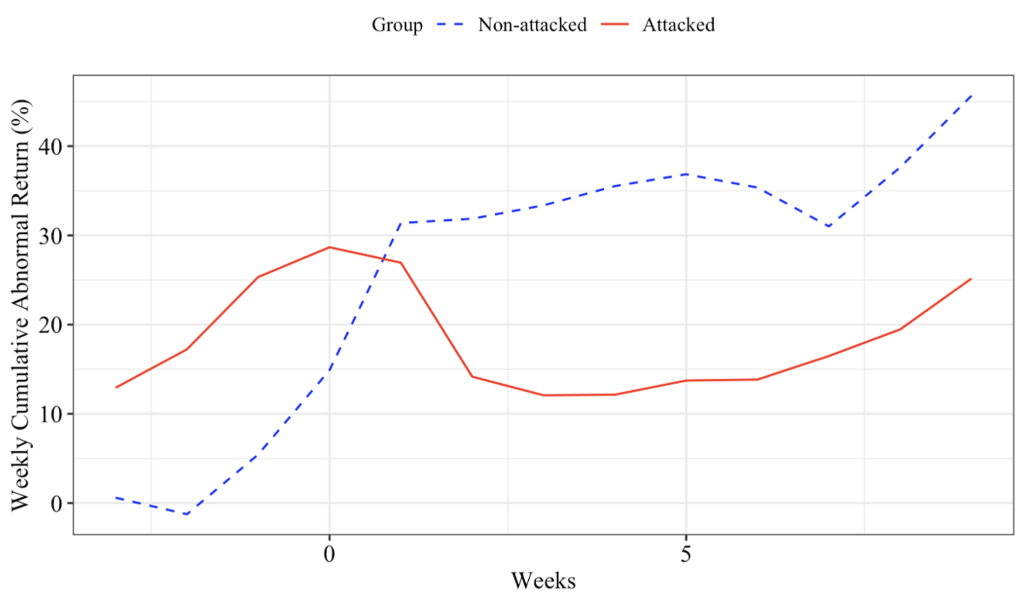

Note: Week0 denotes the week when the cyber-attack occurred. The y-axis shows the weekly cumulative abnormal return (%), computed as the sum of daily abnormal returns within each week.

Figure 2.4: Visualization of Impact of Hack Incidents on DAO Performance

2.5.1 Model Specification

Our primary econometric specification employs a staggered Difference-in-Differences (DiD) approach using DAO-day observations, encompassing the periods 1 month before and 1, 2, and 3 months after each cyberattack on different DAO platforms. Although the two-way fixed effects (TWFE) DiD model is traditionally considered the gold standard for estimation, we turn to the conventional DiD model, as represented by Equation (4), for two principal reasons. First, the impact of the staggered cyberattacks is conspicuously heterogeneous in terms of damages measured in dollar amounts. Recent findings in Econometrics advise against using the TWFE DiD model when treatment effects are dynamic and vary with treatment timing because it can result in biased and inconsistent estimates. These dynamic treatment effects can introduce significant bias when using TWFE models, making the conventional DiD approach more suitable for our analysis (Callaway and Sant'Anna, 2021, Goodman-Bacon, 2021, Athey and Imbens, 2022, Baker et al., 2022). Second, it is methodologically unfeasible to use a DAO-level ownership concentration measure as a moderator within the TWFE DiD model specification. The DAO fixed effects in the TWFE model would absorb all DAO-level time-invariant characteristics, including overall ownership concentration levels. Consequently, using DAO-level ownership concentration as a moderator in the TWFE model would only capture the marginal effect of daily changes in ownership concentration on the dependent variable. Therefore, we adopt the conventional DiD specification, as presented below:

$$y_{i,t} = \beta_0 + \pi \cdot \text{Cyberattack}_i + \rho \cdot \text{Post}_t + \beta_1(\text{Cyberattack}_i \cdot \text{Post}_t) + \theta_{i,t} + \epsilon_{i,t} \tag{2.4}$$

Where the $y_{i,t}$ is the risk-adjusted cumulative abnormal return, capturing token performance changes before and after an attack for DAO i at time t. Cyberattack_i is an indicator that equals to one if a DAO is in treatment group, zero otherwise. Post_t takes the value of one if the period is after the staggered cyberattacks, and zero otherwise. $\text{Cyberattack}_i \cdot \text{Post}_t$ is the DiD estimator, equal to one if DAO i experiences a cyberattack at time t and after and zero if it is a non-attacked DAO. $\theta_{i,t}$ includes vectors of controls for each DAO. $\epsilon_{i,t}$ is the error term, representing unobserved factors influencing $y_{i,t}$ not captured by the other variables in the model.

After examining the overall impact of cybersecurity attacks on the financial performance of DAO tokens, our main analysis focuses on the moderating role of ownership structure on the impact of these attacks. Therefore, we introduce a moderator, $\text{Ownership Concentration}_{i,t}$, which is interacted with the DiD estimator. This allows us to assess how ownership concentration influences the observed impact of cyberattacks on DAO resilience. In Equation (5), the three-way interaction term κ in the updated model is critical for capturing the moderated treatment effects of cyberattacks, providing insights into the differential impacts across DAOs with varying ownership concentration levels.

$$\begin{aligned} y_{i,t} = \beta_0 + \pi \cdot \text{Cyberattack}_i + \rho \cdot \text{Post}_t + \beta_1(\text{Cyberattack}_i \cdot \text{Post}_t) + \gamma \cdot \text{Ownership}_{i,t} \\ + \delta(\text{Cyberattack}_i \cdot \text{Ownership}_{i,t}) + \lambda(\text{Post}_t \cdot \text{Ownership}_{i,t}) \\ + \kappa(\text{Cyberattack}_i \cdot \text{Post}_t \cdot \text{Ownership}_{i,t}) + \theta_{i,t} + \epsilon_{i,t} \end{aligned} \tag{2.5}$$

2.5.2 Propensity Score Matching

To mitigate potential biases stemming from assuming a functional form relationship between our dependent variables and DAO-specific variables, we employ a propensity score matching (PSM) approach. The naive approach of designating the control group as the remaining DAOs in the market may introduce model dependence and biases into our estimates (Bertrand et al., 2004). To address this concern, we reconstruct our DAO-day level panel data by matching attacked DAOs with non-attacked DAOs.

We implement one-to-one nearest neighbor matching, pairing each treated DAO with an untreated DAO that has never experienced any cyberattack. Matching is based on pre-treatment values of Ownership Concentration$_{i,t}$, Transaction Volume$_{i,t}$, # of Token Holders$_{i,t}$, Active Holder Ratio$_{i,t}$, Average Transaction Fee$_{i,t}$. This procedure aims to control for observable DAO-level heterogeneities across treatment and control groups. The covariance balance test demonstrates that the balance between treatment and control group significantly improved.[20]

We conduct the matching without replacement, ensuring each attacked DAO is matched with a non-attacked DAO having the closest propensity score. The caliper for matching is set at 0.1, following Kamiya et al. (2021), which examines the impact of cyberattacks on firm performance. This strategy enhances the robustness of the analysis by creating more comparable treatment and control groups, thereby improving the validity of the causal inferences about the impact of cyberattacks on DAO performance.

[20] The results are omitted due to space limitations but are available upon request.

2.6 Estimation Results

2.6.1 Effect of Cyberattack on DAO Performance

Table 2.4: Difference-in-Difference Estimation Results

	$\text{CAR}_{i,t}$				
	1Month		2Months	3Months	
	OLS	OLS	OLS	OLS	TWFE
	(1)	(2)	(3)	(4)	(5)
DiD Variables					
Cyberattack_i	2.16	–1.42	–4.47	–3.20	
	(2.91)	(2.74)	(4.08)	(5.19)	
Post_t	20.44***	25.56***	47.68***	71.71***	
	(4.08)	(4.01)	(5.06)	(5.96)	
$\text{Cyberattack}_i \times \text{Post}_t$	–28.24***	–30.03***	–39.93***	–52.12***	–31.06***
	(5.37)	(5.14)	(6.17)	(7.26)	(3.16)
DAO Control Variables					
$\text{\#ofTokenHolders}_{i,t}$		–9.30***	–13.34***	–13.99***	
		(0.76)	(0.95)	(1.04)	
$\text{ActiveHolderRatio}_{i,t}$		59.36*	326.69***	679.59***	
		(32.82)	(51.96)	(69.23)	
$\text{AverageTransactionFee}_{i,t}$		–770.38**	–2,810.60***	–910.45*	
		(322.19)	(391.07)	(466.94)	
$\text{TransactionVolume}_{i,t}$		5.02***	7.13***	8.59***	
		(0.29)	(0.37)	(0.42)	
DAO FE	NO	NO	NO	NO	YES
Year–Month FE	NO	NO	NO	NO	YES
Observations	5,286	5,286	7,874	10,343	10,343
Adjusted R^2	0.01	0.10	0.10	0.09	0.87

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors are clustered at the DAO level. Continuous variables are winsorized at the 1% level to mitigate extreme outliers.

To test Hypothesis 1, we first estimate the effects of the cyberattack on the performance of attacked DAOs using Equation (4), where the dependent variable is $\text{CAR}_{i,t}$, the cumulative abnormal return (CAR) at the DAO-day level. Table 2.4 shows the results, confirming that cyberattacks have negative impact on short-term and long-term performance of DAOs. Our baseline model estimations, with and without DAO-level control variables, are presented in Columns (1) and (2), capturing the short-term effects over one month after the attack. The point estimates are statistically significant at the 1% level, and the observed decline in abnormal returns is substantial, measuring -28% and -30%. The effect is economically significant, given the average $\text{CAR}_{i,t}$ of 4.23% for DAOs during our observation period, as shown in Table 2.2. We extend the observation window to 2 and 3 months after each cyberattack in Columns (3) and (4), to capture a sustained long-term impact. Our results reveal that the negative impact of cyberattacks does not attenuate but rather intensifies over time, reaching -39% and -52% after 2 and 3 months, respectively. In Column (5), we also include our DiD estimation results based on the two-way fixed effects (TWFE) model. Although the effect size decreases, the DiD estimator remains statistically significant at the 1% level. The result remains consistent when we use alternate market indices in estimating the CAR, following Appel and Grennan (2023).[21]

This result stands in sharp contrast to recent findings in the cybersecurity literature, which suggest that breaches have a limited impact on firm performance, both in magnitude and duration. In prior studies, the magnitude of the negative effect in CARs ranged from 2.1% to as low as 0.3% within a few days (Cavusoglu et al., 2004, Richardson et al., 2019). Moreover, the gap in CARs between breached and non-breached firms typically subsides within a few weeks, if not sooner, resulting in no significant long-term impact (Richardson et al., 2019). In line with this, Foerderer and Schuetz (2022) explain the limited impact of breaches by noting that firms can strategically time breach announcements, often during periods of high market noise. In their sample, the average (median) delay between breach occurrence and disclosure is 39 (28) days, aligning with typical regulatory requirements that mandate notification within 30 days.

[21] For robustness, we use Ethereum, Bitcoin, and top 5 governance tokens as proxies for market returns, instead of top 100 governance tokens, which generates consistent results (available upon request).

The more severe and persistent losses from cyberattacks in DAOs are understandable, given the structural characteristics of the blockchain ecosystem. Unlike traditional firms, which can manage the timing of breach disclosures, DAOs operate in an environment where disclosure is effectively mandatory, as cyberattacks are instantly visible on-chain (Jung, 2019). This enforced transparency, combined with the heightened emphasis on security in decentralized systems, likely amplifies the shock of a breach. Consequently, market reactions are not only more severe but also persistent, reflecting a deeper erosion of stakeholder trust and elevated sensitivity among token holders to perceived protocol vulnerabilities. Put differently, while traditional firms can mitigate short-term market disruptions through strategic disclosure, DAOs remain highly vulnerable to cyberattacks, which represent not only technical failures but also serious governance challenges.

2.6.2 Moderating Effect of Ownership Concentration on DAO Resilience

To further explore the relationship between DAO ownership structure and resilience (Hypothesis 2), we introduce the ownership concentration variable into the analysis, as denoted by Equation (5). This regression analysis incorporates ownership concentration and its interaction with the shock of cyberattacks. Table 2.5 shows the result. We highlight two key takeaways. First, we find that the observed effect size of the DiD estimator, $\text{Cyberattack}_i \cdot \text{Post}_t$, significantly increases after incorporating the ownership concentration measure, from about -30% (Average Treatment Effect on the Treated, ATT) in Table 2.4 to about -113% in Table 2.5.

Table 2.5: Moderating Effect of Ownership Concentration on Resilience

	$\text{CAR}_{i,t}$		
	1Month	2Months	3Months
	(1)	(2)	(3)
DiD Variables			
Cyberattack_i	78.99***	136.21***	204.12***
	(17.83)	(27.30)	(35.90)
Post_t	118.10***	183.11***	241.65***
	(22.58)	(24.80)	(25.51)
$\text{Cyberattack}_i \times \text{Post}_t$	–113.74***	–133.73***	–137.97***
	(31.54)	(39.35)	(46.75)
$\text{Ownership Concentration}_{i,t}$	10.23***	6.20***	7.23***
	(1.90)	(2.28)	(2.64)
$\text{Cyberattack}_i \times \text{Ownership Concentration}_{i,t}$	–12.41***	–21.91***	–32.24***
	(2.63)	(3.97)	(5.21)
$\text{Post}_t \times \text{Ownership Concentration}_{i,t}$	–14.93***	–21.89***	–27.53***
	(3.45)	(3.72)	(3.82)
$\text{Cyberattack}_i \times \text{Post}_t \times \text{Ownership Concentration}_{i,t}$	13.58***	15.55***	14.66**
	(4.73)	(5.78)	(6.82)
DAO Control Variables			
$\text{\# of Token Holders}_{i,t}$	–8.01***	–11.07***	–10.57***
	(0.73)	(0.89)	(0.99)
$\text{Active Holder Ratio}_{i,t}$	73.08**	354.40***	682.87***
	(31.88)	(50.84)	(67.07)
$\text{Average Transaction Fee}_{i,t}$	–416.76	–2,263***	–170.15
	(331.09)	(397.91)	(472.54)
$\text{Transaction Volume}_{i,t}$	4.95***	7.03***	8.69***
	(0.29)	(0.36)	(0.42)
Observations	5,286	7,874	10,343
Adjusted R^2	0.11	0.11	0.09

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors are clustered at the DAO level. Continuous variables are winsorized at the 1% level to reduce the influence of extreme outliers.

This increase suggests that ownership concentration has a significant moderating effect on the treatment effect of the cyberattack. Second, we find that the level of ownership concentration in a DAO significantly enhances its resilience to cyberattacks. Specifically, the 3-way interaction term, $\text{Cyberattack}_i \cdot \text{Post}_t \cdot \text{Ownership}_{i,t}$, is positive and statistically significant at the 1% level. The effect size is economically noteworthy, revealing that a

1 unit increase in ownership concentration, as measured by Theil's T index, increases cumulative abnormal returns by 13.58% (Column 1). Importantly, this observed effect endures over time, as evidenced in both Column (2) and Column (3).

As a visual illustration, we depict the marginal effect of ownership concentration on DAOs' resilience in Figure 2.5. The X-axis represents ownership concentration, measured by Theil's T index, while the Y-axis indicates the coefficients from the estimation model used in Column (3) of Table 2.5. This graph provides a clear visual representation of the positive marginal effect of ownership concentration on DAOs' resilience. Thus, we find supporting evidence for our hypothesis. That is, the presence of blockholders is associated with increased resilience following cyberattacks.

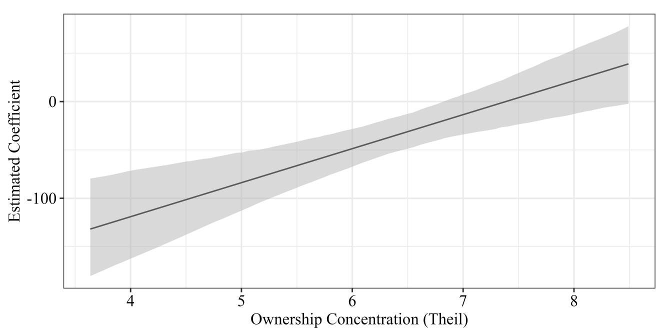

Note: Visualization of the moderating effect of ownership concentration on DAO performance. Shaded bands indicate 95% confidence intervals.

Figure 2.5: Visualization of Moderating Effect of Ownership Concentration

2.6.3 Robustness Analyses

Parallel Trends Assumption Check

To validate our primary difference-in-difference estimation results, which heavily depend on the parallel-trends assumption, we employ a relative time model estimation. This approach involves the use of a lead-and-lags model, allowing for a quasi-experimental design as proposed by Angrist and Pischke (2008). This model enables us to examine whether attacked DAOs and non-attacked DAOs in our sample exhibit similar trends before the staggered treatments. In this quasi-experiment, we construct a weekly-level dependent variable. This choice is motivated by the tendency of abnormal returns to exhibit fluctuations. By aggregating at the weekly level, we aim to smooth out the daily-level fluctuations. Our reference base case is set as the week preceding each cybersecurity event ($t-1$ week). This choice is grounded in the expectation that the impact of new and potentially detrimental information on the market is immediate.

The results from our estimation, presented in Table 2.6, reveal that all pre-treatment time dummies, when interacted with treatment dummies, are statistically insignificant. This suggests that there were no discernible differences in the outcomes of interest between the treatment and control groups during the pre-treatment period. This finding aligns with the model-free evidence presented in Figure 2.4. A noteworthy observation is that the effect size remains relatively consistent and does not exhibit a decrease. This consistency implies a sustained negative impact of cybersecurity events on the resilience of the impacted DAOs.

The insignificance of estimates prior to the treatment is pivotal, serving as a critical baseline for understanding the natural trajectory of observed outcomes. This baseline establishes a foundation for assessing the treatment's effectiveness, providing a clear benchmark against which post-treatment changes can be evaluated. Consequently, any observed alterations in outcomes post-treatment can be more confidently attributed to the treatment itself, rather than being influenced by underlying temporal trends or other confounding factors that might have affected the estimates prior to the intervention.

Table 2.6: Relative Time Model Estimation

Relative Time	$\text{CAR}_{i,t}$
Relative $\text{Time}_{t-4\text{ week}}$	0.001
	(0.26)
Relative $\text{Time}_{t-3\text{ week}}$	0.05
	(0.26)
Relative $\text{Time}_{t-2\text{ week}}$	–0.27
	(0.26)
Relative $\text{Time}_{t-1\text{ week}}$	*Omitted (base)*
Relative $\text{Time}_{t\text{ week}}$	–0.44*
	(0.26)
Relative $\text{Time}_{t+1\text{ week}}$	–0.48*
	(0.26)
Relative $\text{Time}_{t+2\text{ week}}$	–0.51**
	(0.26)
Relative $\text{Time}_{t+3\text{ week}}$	–0.51**
	(0.26)
Relative $\text{Time}_{t+4\text{ week}}$	–0.46*
	(0.26)
Relative $\text{Time}_{t+5\text{ week}}$	–0.39
	(0.26)
Relative $\text{Time}_{t+6\text{ week}}$	–0.42
	(0.26)
Relative $\text{Time}_{t+7\text{ week}}$	–0.43*
	(0.26)
Relative $\text{Time}_{t+8\text{ week}}$	–0.52**
	(0.26)
Relative $\text{Time}_{t+9\text{ week}}$	–0.54**
	(0.26)
Relative $\text{Time}_{t+10\text{ week}}$	–0.54**
	(0.26)
Relative $\text{Time}_{t+11\text{ week}}$	–0.50*
	(0.26)
Relative $\text{Time}_{t+12\text{ week}}$	–0.52**
	(0.26)
DAO FE	YES
Observations	1442
Adjusted R^2	0.84

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the DAO level. Continuous variables are winsorized at the 1% level to mitigate extreme outliers.

Static Measure As an Alternative Ownership Concentration Variable

As explained in Section 2.4.2, we incorporate a static measure of ownership concentration in our DiD regression analysis, as provided in Appendix. We report consistent impact of moderating effect of ownership concentration on DAO resilience after cyberattacks.

Random Implementation Test

We next explore the potential for false significance in the estimates, considering the influence of spurious relationships or serial correlation in the dependent variable. To address this concern, we conduct a random implementation test. We generate pseudo-treatment indicators by randomly assigning the treatment status within the original dataset. The coefficient derived from this pseudo-treatment is stored, and the entire procedure is iterated 1,000 times. This test is to evaluate the potential spuriousness of any significant findings, stemming from autocorrelation in the dependent variable, abnormal returns, in our context (Bertrand et al., 2004).

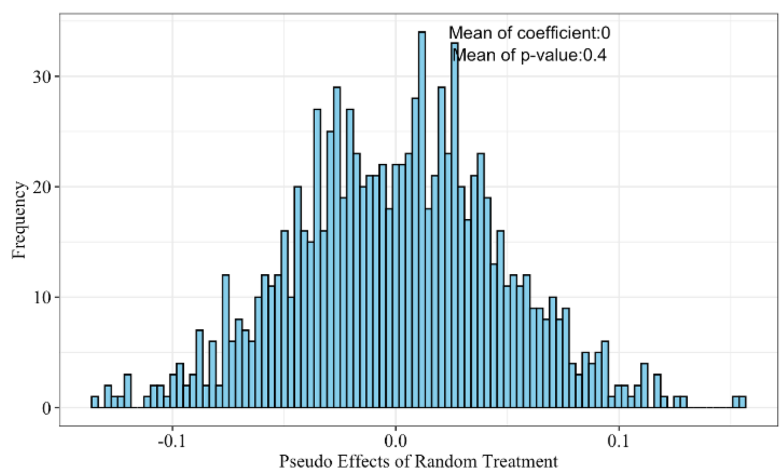

Note: Distribution of pseudo (placebo) treatment effects generated by randomly shuffling the treatment indicator.

Figure 2.6: Distribution of Pseudo Effects of Random Treatment

Figure 2.6 presents the histogram depicting the distribution of pseudo effects associated with randomly assigned treatment coefficients. Notably, we observe that the estimated average linked to the randomly assigned treatment indicator converges to 0, and the corresponding p-value (=0.4) is not statistically significant. This convergence suggests that the primary results are reliably estimated, providing confidence in the robustness of our findings.

2.7 Causal Mediation Analysis

2.7.1 Governance Participation as Mediator

As illustrated in Section 2.2, community's participation in governance is a fundamental element of DAOs' operation. This section explores a potential mechanism behind the relationship between the concentrated token ownership and DAO's financial performance, viewed through governance participation. Mediation analysis is an empirical method to assess whether and how the impact of an independent variable or exposure (in our study: the level of ownership concentration) on a dependent variable or outcome (i.e., performance of DAO) occurs through a mediating variable (i.e., governance participation). To test the presence of mediating role of governance participation, we follow the causal counterfactual framework proposed by Imai et al. (2010), developed from the original mediation model by Baron and Kenny (1986).

More specifically, we decompose the effect of ownership concentration on DAO's performance into direct and mediated components, which are called as the average direct effect (ADE) and the average causal mediating effect (ACME). The average direct effect measures the extent of the performance change that can be attributed to ownership concentration independent of participation rate. In contrast, the average causal mediating effect captures the portion of the performance change that occurs due to ownership concentration's influence on participation. In summary, the total effect of ownership concentration on DAO performance decomposes into ADE and ACME, allowing us to compute how much of the total effect of ownership concentration is actually being mediated through governance participation.

To assess governance participation in DAOs, we turn to *Snapshot*, one of the most widely used governance platforms designed for DAO voting (see Figure 2.2). Among the DAOs included in the DiD analysis reported in Table 2.5, 53 actively utilize Snapshot for off-chain voting, making them subjects of this analysis. Using the GraphQL API[22], we retrieved all historical proposals, their associated voters and tokens voted, resulting in a dataset of 2,905 Snapshot proposals and 2,900,012 votes from 168,949 unique voters.

Table 2.7: Results of Causal Mediation Analysis

	ACME	ADE	Type of Mediation	% Mediated
Mediators				
% of Voters$_{i,p}$	0.006***	0.037***	Complementary	14
% of Tokens Voted$_{i,p}$	0.002***	0.040***	Complementary	5

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcome variable is Daily Return$_{i,t}$. ACME = Average Causal Mediation Effect; ADE = Average Direct Effect; % Mediated = share of the total effect explained by the mediator.

To conduct the mediation analysis, we construct the outcome variable based on the seven-day window following each proposal's end date. Specifically, we use changes in daily token prices (Daily Returni, t) to estimate how existing token holders or potential investors react to the proposal. Our key independent variable, Ownership Concentrationi, t, captures the overall level of token ownership concentration within each DAO. Next, we construct two DAO-proposal level participation variables as our mediators: % of Voters$_{i,p}$ and % of Tokens Voted$_{i,p}$. % of Voters$_{i,p}$ represents the voting participation rate in each proposal, calculated as the number of wallet IDs that cast votes normalized by the total

[22]https://graphql.org/

number of token holders. % of Tokens Voted$_{i,p}$ represents the proportion of tokens used to cast votes in each proposal, calculated as the number of tokens used in voting divided by the total number of tokens. These variables capture the extent of token holder engagement and token utilization in the voting process, providing a measure of overall governance participation levels. For more detailed descriptions of the variables, please refer to Table 3.1.

Table 2.7 shows results of the causal mediation analysis. The estimated ACMEs of governance participation rate are 0.006 for % of Voters$_{i,p}$ and 0.002 for % of Tokens Voted$_{i,p}$, both statistically significant at the 1% level. The corresponding ADEs are 0.037 and 0.040, respectively, also significant at the 1% level. The mediation effects are complementary to the direct effect, as both share the same sign. The point estimates indicate that approximately 14% of the total effect is mediated through % of Votersi, p, and about 5% by % of Tokens Votedi, p. This suggests that active governance participation mediates the relationship between ownership concentration and the performance of a DAO. More detailed explanations and the results from two-step regression for the mediation analysis are provided in Appendix.

2.7.2 Ownership Concentration and Participation: An Inverted-U Relationship

We further examine the potential curvilinear relationship between ownership concentration and the governance participation. As illustrated in Figure 2.7, when DAO ownership concentration exceeds a certain threshold, token holders appear less engaged in the community decision-making process, as reflected in the declining participation rate. Following the seminal work by Aghion et al. (2005) on the nonlinear economic relationship between competition and innovation, we interpret the inverted-U-shaped relationship between ownership concentration and governance participation as follows. At low levels of ownership concentration, DAOs are comprised of widely distributed token holders, with no dominant stakeholders whose opinions carry significant weight. In such a setting, individual token holders are unlikely to perceive a high marginal benefit from voting to justify the associated participation costs, i.e., the time and effort required. Hence, free-riding behavior will emerge as an attractive strategy, leading to lower overall participation rates in DAOs with lower ownership concentration (Ellinger et al., 2024). In other words, if the level of ownership concentration is very low in a given DAO, an increase in ownership concentration should induce more active governance participation.

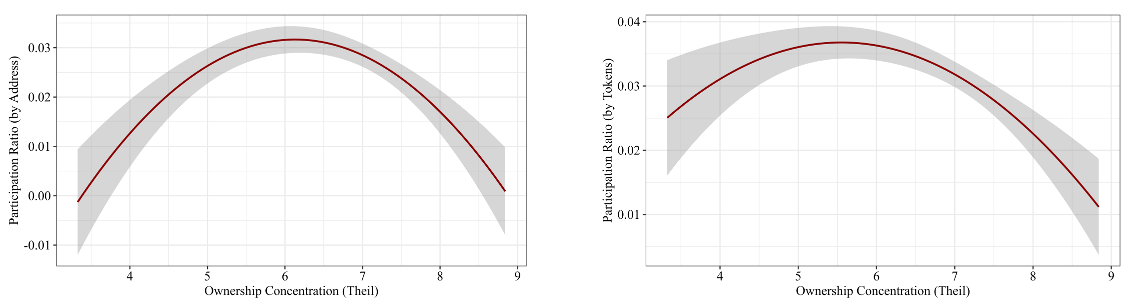

(a) Participation Ratio by Addresses (b) Participation Ratio by Tokens Voted

Note: Curvilinear relationship between ownership concentration and participation on Snapshot. The shaded bands depict 5% confidence intervals.

Figure 2.7: Curvilinear Relationship between Ownership Concentration and Snapshot Participation

On the other hand, when ownership concentration is very high, a small number of large token holders dominate the decision-making process, effectively reducing the influence of smaller holders. In such a setting, smaller token holders may perceive their votes as inconsequential, leading to disengagement from governance activities. As a result, overall participation rates decline at high levels of ownership concentration. This dynamic explains why, beyond a certain threshold, increasing ownership concentration can discourage participation, thereby creating the inverted-U-shaped relationship.

2.8 Discussion

This study examines how token ownership structure influences a DAO's resilience to disruptions, using cyberattacks as a natural setting to observe performance under stress. By analyzing 42 cyberattacks across 115 leading DAOs and employing a staggered difference-in-differences framework, we show that cyberattacks result in significant and prolonged performance losses for DAOs. However, DAOs with more concentrated token ownership, i.e., those with more blockholders, recover more swiftly, experiencing a less severe and more transient drop in cumulative abnormal returns. Furthermore, we find that the resilience effect is partially mediated by governance participation, lending support to the blockholder theory. In addition, we identify an inverted-U-shaped relationship between ownership concentration and governance participation, suggesting that both too little and too much concentration can dampen engagement. This insight underscores the importance of a balanced ownership structure in DAO governance.

In conclusion, by centering cyberattacks as a governance stress test, this study challenges the common presumption that decentralized governance necessarily entails equal distribution of power. Instead, our results suggest that decentralization should be seen as a functional process of governance. In this view, the presence of influential blockholders is not undermining decentralization; it may in fact facilitate it, by sustaining coordination and recovery through effective governance mechanisms such as voice. This insight adds nuance to the ongoing discourse around decentralization, proposing that what matters most is how decisions are made, not how evenly influence is distributed. A DAO may still be meaningfully decentralized if it operates transparently even if some token holders wield greater influence than others.

2.8.1 Theoretical Contributions

This study contributes to the growing literature on DAOs by challenging the prevailing assumption that decentralization is synonymous with equal power distribution (Lumineau et al., 2021, Beck et al., 2018). Instead, our findings suggest a redefinition: decentralization is better understood as a process of transparent, participatory governance rather than as an egalitarian power structure. By showing that DAOs with higher ownership concentration, those with more blockholders, demonstrate stronger resilience after cyberattacks, we emphasize that the functionality of governance mechanisms matters more than the symmetry of token distribution. This reconceptualization aligns with broader organizational theories that prioritize decision-making quality, coordination, and stakeholder engagement as the true indicators of robustness in collective systems (Ellinger et al., 2024, Hsieh and Vergne, 2022, Zhao et al., 2022).

Moreover, we empirically extend blockholder theory to decentralized settings, demonstrating that even in the absence of centralized management, blockholders can play a critical role in facilitating decentralized governance by enhancing participation and collective efforts (Shleifer and Vishny, 1986, Edmans, 2009). This also marks a shift in how decentralization is theorized: central figures or concentrated holders do not contradict decentralization when they operate within transparent, rule-based systems that facilitate collective input and responsiveness.

Finally, our study foregrounds resilience as a governance outcome rather than merely a performance metric, positioning cybersecurity as a critical yet broadly underexplored lens in the study of decentralized organizations (Lumineau et al., 2021, Cong et al., 2024).

By leveraging cyberattacks as exogenous stress tests, we reveal how the structure and dynamics of token ownership shape a DAO's ability to absorb, adapt to, and recover from operational shocks. This framing not only bridges decentralized governance and organizational resilience but also elevates cybersecurity as a foundational concern for assessing the viability and maturity of DAOs. In doing so, our work opens new pathways for future research on crisis management, transparency and trust in decentralized systems (Ellinger et al., 2024).

2.8.2 Managerial Implications

This study offers critical insights for designers, contributors, and researchers involved in DAO ecosystems. Our findings emphasize that effective decentralization is defined less by equal power distribution and more by the ability of collective governance to function reliably, particularly under conditions of stress. Given that DAOs operate without centralized management, resilient governance necessarily arises from the community itself. Recognizing this, ownership concentration can serve as an essential resource rather than a liability. Blockholders, due to their significant stakes and visibility, can play crucial roles in stabilizing governance during crises by sustaining participation and initiating recovery efforts. Nevertheless, our research highlights the risks associated with extremes in both dispersion and concentration of ownership, underscoring the need for a balanced approach. Achieving this balance requires active dialogue among community members, DAO designers, and protocol developers concerning tokenomics and incentive frameworks. Such collaboration is vital to ensuring robust engagement and governance resilience while avoiding problematic extremes.

Additionally, the governance infrastructure significantly influences stakeholder participation. Practical measures to reduce participation friction through off-chain voting tools like Snapshot, token delegation, and asynchronous discussion forums, and targeted community incentives, to name a few, can empower stakeholders to engage more meaningfully. Ultimately, resilient decentralization involves not the removal of influence but embedding influence within transparent, accountable, and responsive processes. Thus, DAO governance must be approached as a dynamic, iterative, and fundamentally social design challenge.

2.8.3 Limitations and Future Research Directions

As one of the earliest efforts to examine how decentralized governance operates in DAOs, particularly in the context of cyberattacks, this study has some limitations that also highlight promising directions for future research. First, as most decentralized platforms are still in the initial stages of organizational development, the governance structures and behavioral dynamics observed in this study may continue to evolve. At the same time, even in their current form, DAOs exhibit diverse voting rules and governance mechanisms across platforms, variations that are not fully captured in our analysis. This can raise a valid question about the boundary conditions of our findings. Nevertheless, it is important to note that our study centers on the general relationship between token ownership, governance participation, and organizational performance, addressing a fundamental issue of ownership and control in organizations without centralized authority. This focus comes from the core assumption that individuals with greater token ownership have stronger incentives to engage in governance, an incentive structure that remains consistent across DAOs. Accordingly, even with variation in specific governance rules, the theoretical logic underlying our findings should hold broad relevance for understanding decentralized governance in DAOs where the decisions are made via token-weighted community voting.

Second, our findings reflect general patterns of shock and recovery observed across

DAOs following cyberattacks, without capturing the specific recovery strategies employed by individual DAOs. In other words, the analysis does not fully account for the possibility that DAOs with similarly active token holders may adopt different post-attack approaches, resulting in divergent recovery trajectories. Future research could explore the heterogeneity of responses to cyberattacks across DAOs, examining how governance structures and community decision-making processes influence their responses to breaches and, consequently, their ability to recover. Beyond the specific focus on governance in our study, the heightened significance of security risks in blockchain ecosystems presents a promising avenue for IS scholars to examine how technological infrastructure shapes organizational culture in decentralized contexts.

Finally, this study takes a high-level approach, focusing on aggregate patterns of ownership, participation, and performance. As a result, it does not directly account for individual-level factors, such as the characteristics of token holders and voting participants, which are likely to play a critical role in shaping the effectiveness of DAO governance. Future research could explore how the composition and experience of the token holder base shape governance quality and organizational outcomes, by leveraging data that traces individual token holders' voting behavior across DAOs, e.g., whether they are active participants in multiple DAOs. Such analyses would allow for a deeper understanding of how individual-level community dynamics influence decentralized governance in DAOs.

3 Tokenized Access: How NFT Market Empowers Minority Artists

3.1 Introduction

Non-Fungible Tokens (NFTs) are blockchain-based digital assets that certify ownership and authenticity of unique items, such as artworks, music, or collectibles, through tamper-proof, decentralized ledgers. Unlike traditional digital files, NFTs carry verifiable provenance and scarcity, enabling digital content to be treated as tradable assets. Over the past few years, the NFT art market has grown rapidly, offering artists direct access to global audiences without relying on intermediaries such as galleries or auction houses.

This disintermediation holds particular promise for underrepresented artists, including women and racial minorities, who have historically faced barriers to entry, recognition, and compensation in the conventional art world. However, whether NFT platforms genuinely level the playing field remains an open question. Markets characterized by high uncertainty and aesthetic subjectivity, such as art, often reproduce inequality through indirect mechanisms. In the absence of standardized benchmarks or institutional endorsements, evaluators may rely on heuristics and observable cues such as artist identity, profile narratives, or perceived credibility (Adams et al. 2021, Bohren et al. 2019, Cui et al. 2020).

This study investigates whether demographic disparities exist in the NFT art market, and if so, how such disparities can be mitigated. Using transaction-level data from SuperRare, a leading curated NFT platform, we analyze approximately 27,000 sales and 90,000 bids from 2,500 artists. We ask three main questions: (1) Are artworks by underrepresented artists, female and non-White individuals, less likely to be sold and sold at lower prices? (2) Does providing artist information attenuate these disparities? (3) If so, which types of information are most effective, e.g., greater detail, third-party-verifiable signals, or contextual cues?

Our findings contribute to ongoing debates on equity in digital creative markets. While decentralization expands access, it does not automatically eliminate outcome disparities. We show that female and non-White artists face significant disadvantages in sale probability and pricing, especially in the case of racial minority artists. However, these disparities are meaningfully reduced when artists provide credible information. Notably, exhibition history and location in advanced economies substantially mitigate racial valuation gaps, while client history and narrative elaboration benefit female artists. These findings underscore the importance of quality signals in promoting fairness in decentralized art platforms, where artists bear greater responsibility for shaping how their work is perceived. In the absence of institutional curation, strategic self-presentation becomes a critical tool for creators to overcome structural disadvantages.

3.2 Theoretical Background and Related Literature

3.2.1 Digitization and Its Limits in the Art Market

Digital technologies have reshaped the structure of traditional markets by reducing search costs, lowering entry barriers, and reallocating control away from entrenched intermediaries (Parker and Van Alstyne, 2005). Across sectors, online platforms have enabled creators and sellers to reach broad audiences directly. For example, digitization has expanded consumer access to product information and reviews, weakening the dominance of incumbent firms and increasing market transparency (Brynjolfsson et al. 2009, Reimers and Waldfogel 2021). These changes have altered the dynamics of competition, allowing smaller players to challenge established ones by overcoming traditional entry barriers and engaging directly with audiences online (Kokkodis et al., 2023).

In content and creative markets, similar forces have enabled individuals to bypass publishers, studios, and other traditional gatekeepers. The rise of platforms such as YouTube and Kickstarter has shown how digital infrastructure allows creators to distribute their work and raise capital independently (Waldfogel, 2017). Empirical work documents how reputation systems and crowdfunding mechanisms facilitate trust, reduce uncertainty, and empower new entrants in domains long governed by centralized institutions (Kokkodis et al. 2023, Agrawal et al. 2014). These shifts have opened participation in cultural production to more diverse voices, positioning digitization as a powerful democratizing force across industries.

Despite these shifts, the fine art market has stood apart. It continues to rely on a concentrated set of intermediaries, such as galleries, curators, and auction houses, for authentication[1], valuation, and market access (Whitaker, 2019). Unlike consumer goods or digital media, artwork is characterized by its non-fungibility[2], aesthetic subjectivity, and limited verifiability (Li et al., 2024). These attributes have historically justified the dominant role of institutions as arbiters of quality and legitimacy. Indeed, the literature in art economics has emphasized the authority of art dealers and curators in establishing value and market access (Velthuis 2005, Alexander 1996). While artists have been able to display their work through online portfolios and retail-like platforms, the authentication of originality, the establishment of value, and access to serious collectors have continued to demand mediation by recognized institutions (Li et al. 2024, Kräussl and Tugnetti 2024). As a result, the art market has remained largely resistant to the democratizing effects observed in other digitally transformed sectors. The following section discusses how blockchain-based NFTs represent a structural break from this institutional dependence.

[1] Authentication is to prove the correct authorship, operated through various methods, such as human expertise, scientific analysis, and certificate of authenticity.

[2] Non-fungibility refers to an asset's uniqueness and indivisibility; unlike fungible goods such as currency, each item is distinct. This property underlies the term "non-fungible token" (NFT).

3.2.2 NFTs as a Bridge from Peripheral to Curated Art Markets

Before blockchain, digital art largely circulated on open Web 2.0 platforms with little connection to the elite fine art world. These venues suffered from extreme oversupply, low-to-mid price points, and no reliable mechanism for scarcity or provenance. Collectors often questioned authenticity or ownership because digital files can be copied infinitely. Payments via banks or processors faced delays, and artists bore printing and shipping costs for limited editions.

Blockchain technology introduces a new infrastructure of trust by encoding ownership and provenance on a shared ledger. NFTs (non-fungible tokens) operationalize this system for digital art, recording authorship, transaction history, and scarcity in permanent on-chain metadata (Howell et al. 2020b, Kräussl and Tugnetti 2024). This durable record solves the duplication problem and enables trustless verification of an artwork's authenticity; legitimacy no longer relies on centralized gatekeepers (Whitaker 2019, Kräussl and Tugnetti 2024).

Curated NFT marketplaces exploit these features to replicate the scarcity and prestige of the traditional art market. Admission is controlled—only a few new creators are approved each month—so supply remains scarce and prices routinely reach several thousand dollars, with top works selling in the five- to six-figure range. Each artwork is minted as a unique on-chain token, making provenance irrefutable and eliminating unauthorized copies. Payments settle instantly in cryptocurrency, bypassing banking friction; smart contracts automatically deliver resale royalties to the artist on every secondary sale (Whitaker 2019). Because the artwork is digital data, artists avoid all shipping, insurance, and customs costs.

The impact on access and agency is profound. Without geographic or institutional gatekeepers, artists can reach global audiences based on the merits of their work alone (Li et al. 2024, Whitaker 2019). Galleries and auction houses traditionally emphasized credentials such as formal education or location; on NFT platforms, creators instead highlight personal narratives, philosophies, and community engagement. Emerging artists from peripheral regions, previously excluded by infrastructure limitations or capital controls, can now participate fully, since blockchain payments bypass national banking barriers and markets maintain transparent records. Nigerian NFT artist Osinachi, for example, attributes his commercial success to blockchain, which enabled him to reach international audiences without leaving his home country.[3] These platform-induced changes imply that blockchain's capacity to broaden access in the art market should be assessed based on measurable changes in price formation, revenue persistence, and market access for creators outside traditional art hubs.

Overall, the curated NFT ecosystem presents a new, blockchain-backed model of the fine art market for digital mediums. Platforms like SuperRare, with their rigorous curation and one-of-one token standards, offer digital artists a pathway into a segment of the market that had no close analog before blockchain. By enforcing scarcity, transparent ownership, and perpetual royalties on-chain, this platform design redistributes informational agency toward creators. These patterns suggest a promising direction for future research on how blockchain can transform structural access and expand opportunities for creators in the art market.

3.2.3 Artwork Valuation and the Female Premium vs. Discount Debate

In traditional art markets, the pricing of artworks reflects not only objective artistic characteristics, such as medium, size, and style, but also subjective factors, including ownership utility and anticipated resale value (Renneboog and Spaenjers 2013, Li et al. 2024). Given

[3] See his interview with MoMA (the Museum of Modern Art) Magazine, "Collaboration and Creativity on Blockchain," Oct. 2023. https://www.moma.org/magazine/articles/953

the inherent uniqueness of artworks, pervasive information asymmetry, and risk of fakes and forgeries, trust built through reputable intermediaries (e.g., galleries and auction houses) has been essential. Hence, provenance, such as pedigree, i.e., the history of prior owners, and certificate of authenticity (COA) significantly influences market valuation (Li et al. 2024, Whitaker 2019).

Relatedly, the distinctive characteristics of the art market also influence valuation disparities linked to socio-cultural factors, such as gender and race of the artists. Early attention to gender disparities was famously initiated by Linda Nochlin's seminal essay, *"Why Have There Been No Great Women Artists?"* published in 1971, which set the stage for empirical explorations into systematic market inequalities across demographic groups. A growing body of research has documented disparities in artists' economic outcomes, particularly with respect to gender. Multiple studies report that artworks by female artists sell at unconditional discounts in global art markets. For instance, Adams et al. (2021) find that artworks by women sell for 42.1% less than those by men, based on a dataset of 1.9 million auction sales across 49 countries. Similarly, LeBlanc and Sheppard (2022) identify a discount of 13% to 19%, after controlling for quality indicators such as medium, dimensions, and auction house prestige. These studies suggest a consistent valuation gap that cannot be fully explained by observable differences in artistic output. Additional evidence from experiments suggests that buyer-side perceptions may contribute to observed discounts. Adams et al. (2021) and Hoffmann and Coate (2022) show that participants in valuation tasks rely on gender stereotypes, tending to associate prestige and higher price with male-associated works.

In contrast, alternative empirical evidence highlights scenarios under which female artists achieve a valuation premium. For instance, Cameron et al. (2019) examine auction results for graduates from the Yale School of Art and report that, conditional on reaching auction markets, works by female artists command higher prices compared to their male counterparts. This finding aligns with the "higher-bar" hypothesis, suggesting that institutional barriers disproportionately restrict women's entry, resulting in a highly selective group whose artworks may be perceived as possessing superior quality. Similarly, Bocart et al. (2022) find that while female artists encounter greater difficulty transitioning from primary to secondary markets, those who succeed secure an average premium of 4.4% relative to male artists. However, their analysis also reveals that this premium diminishes and eventually reverses within the highest market segment, specifically for artworks exceeding $1 million, where male artists continue to dominate sales.

Collectively, these findings indicate that pricing disparities in the art market stem from an interplay between demand-side perceptual biases and supply-side institutional barriers. The simultaneous occurrence of unconditional discounts and context-specific premiums suggests that disparities are structurally embedded and influenced by multiple underlying mechanisms. This complexity underscores the importance of examining whether and how emerging market infrastructures, such as blockchain-based platforms, might mitigate or perpetuate these valuation patterns. Moreover, despite increased scholarly attention to gender-based disparities, empirical research on racial or ethnic valuation disparities remains limited, primarily due to severe historical under-representation resulting in insufficient market presence of nonwhite artists to enable reliable quantitative analysis (Whitaker and Kräussl 2020, Adams et al. 2021).

3.3 Hypothesis Development

In this section, we develop hypotheses on disparities in the curated NFT art market. Building on pricing disparity debates in the art economics literature, we examine whether artists from underrepresented groups, hereafter minority artists (i.e., female and non-white artists), face sales disadvantages in NFT marketplaces. To understand how information provision may shape market outcomes, we draw on signaling theory and research on the role of information in online platforms. Specifically, we investigate whether artist-provided information such as exhibition history, client credentials, and geographic location can help reduce these disparities in the absence of traditional gatekeepers.

3.3.1 Sales Disparities for Minority Artists in the NFT Market

As discussed in Section 3.2.2, the curated NFT art market differs from its traditional counterpart in several ways that may alter the dynamics of identity-based disparities. Building on insights from art economics discussed in Section 3.2.3, we examine three perspectives that may explain how such disparities emerge in the NFT art market: supply-side market entry conditions, and two demand-side mechanisms related to investor profiles and information availability.

From a supply-side perspective, which emphasizes the role of institutional barriers that historically imposed higher entry thresholds for female artists (Cameron et al. 2019, Bocart et al. 2022), the removal of intermediaries in the NFT art market, i.e., galleries and auction houses, may introduce a significant shift. With the gate now open to all artists and only minimal quality-based filtering in place, the structural disadvantage faced by female artists at the point of market entry may be alleviated. As a result, the pricing premium previously observed for minority artists, potentially driven by the perception that only exceptionally qualified minority artists could enter the market, may not be relevant in the NFT art market.

From a demand-side perspective, prior studies suggest that artworks by female artists sell for significantly less than those by male artists, mainly because of the biased curators and collectors, who associate prestige and higher price with male-associated works (LeBlanc and Sheppard 2022, Hoffmann and Coate 2022). Prior research suggests that gender-based pricing disparities in the traditional art market may be partly driven by characteristics of its collector base, namely, affluent male participants who frequently visit galleries and exhibit stronger stereotypical beliefs (Adams et al., 2021). Given that NFT investors tend to be younger, more digitally native, and culturally distinct from conventional collectors, the same gender discount may not manifest to the same extent. While this does not preclude the presence of disparities, it opens the possibility that the NFT market may exhibit different dynamics with respect to artist identity.

Lastly, prior research across digital platforms shows that outcome gaps across demographic groups often emerge when product quality is difficult to evaluate or when assessments are inherently subjective (Bohren et al. 2019, Freeman and Huang 2014). Under such conditions of uncertainty, decision-makers are more likely to rely on stereotypical heuristics, thereby increasing the scope for discrimination. These dynamics can be particularly relevant in art markets, where aesthetic judgments lack standardized benchmarks and are highly susceptible to perception-driven disparities (Velthuis, 2005). In traditional art markets, intermediaries such as galleries and auction houses have played a central role in providing quality signals and background information, which partially offsets this ambiguity (Renneboog and Spaenjers 2013, Li et al. 2024). However, in NFT marketplaces, institutional validation is absent, artist profiles often lack standardized information, and the provision of such information remains optional. As a result, in this low-information environment, art collectors may place greater weight on visible identity cues when forming judgments, which may exacerbate disparities in sales outcomes for minority artists (Doleac and Stein 2013, Laouénan and Rathelot 2022).

All in all, while the structural openness and evolving collector demographics in NFT markets may reduce certain barriers, the limited availability of artist-level information may still leave room for identity-based disparities. We therefore hypothesize:

Hypothesis 1: *In decentralized NFT marketplaces lacking traditional gatekeepers, artworks by minority artists, including female and non-white individuals, are associated with lower likelihoods of sale and lower sale prices relative to artworks by majority-group artists.*

3.3.2 The Benefits of Artist Self-Curation via Information Disclosure in Reducing Minority Discounts

While NFT marketplaces may lack standardized artist-level information, artists have the opportunity to self-curate by voluntarily providing relevant background details, such as exhibition history, client credentials, or geographic location. Prior research suggests that in environments characterized by uncertainty and limited information, such disclosures can meaningfully influence perceptions and outcomes. In particular, credible and task-relevant signals can shift attention away from heuristic-based judgments, including those tied to gender or racial identity (Cui et al. 2020, Laouénan and Rathelot 2022, Agrawal et al. 2016).

This logic is especially relevant in the art market, where aesthetic evaluation is inherently subjective and quality benchmarks are often ambiguous (Velthuis 2005, Freeman and Huang 2014). In the absence of institutional validation, artist-provided information may function as a substitute for traditional gatekeepers, helping buyers infer credibility and artistic merit. Empirical studies across digital labor and rental markets similarly show that detailed information provision reduces demographic disparities in outcomes (Agrawal et al. 2016, Cui et al. 2020). We therefore expect that artist self-curation through information disclosure will improve sales performance in NFT marketplaces.

Hypothesis 2. *In the NFT art market lacking traditional gatekeepers, artist-provided information enhances the likelihood of sale and increases sale prices.*

While information disclosure can improve sales outcomes for artists in general, prior research suggests that such disclosure may be especially valuable for artists from underrepresented groups. In contexts where demographic cues trigger biased expectations or group-based priors, providing credible, performance-relevant information can help mitigate the effects of those priors and shift attention toward individual-level qualifications (Bohren et al. 2019, Cui et al. 2020, Laouénan and Rathelot 2022).

Because minority artists are more likely to face negative assumptions based on identity, the potential gains from clarifying their qualifications through self-curated information may be greater. In digital labor and rental markets, for instance, underrepresented individuals (i.e., workers from developing countries and African American hosts) benefit disproportionately from providing detailed information or verified credentials (Cui et al. 2020, Agrawal et al. 2016). We extend this logic to the NFT art market, where the absence of institutional validation and the salience of identity cues create conditions under which such information may carry particular weight for minority artists.

We further specify three mechanisms through which information provision may attenuate disparities in the NFT art market. Prior research across various digital platforms suggests that increasing the availability of information enhances perceived credibility, improves discoverability, and facilitates inclusion. For example, in the context of crowdfunding, detailed project descriptions and visual materials are positively associated with campaign success, as they help reduce uncertainty and attract backers' attention (Agrawal et al. 2014, Mollick 2014). Drawing on this logic, we posit that providing greater volume of information, regardless of its specific content, may help underrepresented artists overcome baseline disadvantages in visibility or trust, thereby improving sales outcomes. Therefore, we propose:

Hypothesis 2a. *In the NFT art market, providing greater volume of information is* *associated with greater reduction in sales disparities (in terms of likelihoods of sale and prices) for minority artists.*

While increasing the volume of artist information may improve visibility and help mitigate disparities, the credibility and verifiability of that information are equally, if not more, important. A growing body of literature suggests that unverifiable signals, such as self-reported biographical details, often fail to shift perceptions under uncertainty. In contrast, externally validated signals can reduce ambiguity and help decision-makers form more accurate beliefs. For instance, in online marketplaces, reputation systems based on third-party feedback, such as ratings and reviews, have been shown to play a critical role in building trust and reducing transaction frictions (Resnick and Zeckhauser 2002, Pavlou and Gefen 2004). Similarly, in peer-to-peer platforms such as Airbnb, verifiable user attributes like customer reviews have been found to significantly improve outcomes for marginalized users, whereas unverifiable content like personal bios has limited effectiveness in correcting disparities (Cui et al. 2020, Laouénan and Rathelot 2022).

In the context of NFT art markets, where institutional gatekeepers are absent and standards for artistic quality are ambiguous, third-party-verifiable information, such as exhibition history, or client affiliations, can serve as proxy signals of artistic legitimacy. These credible signals may be particularly effective in mitigating disparities by redirecting evaluative attention toward market-relevant indicators.

Hypothesis 2b. *In the NFT art market, providing verifiable artist information (e.g., exhibition, client history) reduces sales disparities (in terms of likelihoods of sale and prices) for minority artists.*

3.3.3 Cultural Capital and the Geography of Artistic Valuation

In traditional art markets, provenance, including exhibition history in cultural capitals, has long served as a proxy for authenticity, expert vetting, and perceived quality, often yielding significant price premiums at auction (Li et al., 2024). However, beyond its role in authentication, recent research in economic geography and the sociology of art highlights the symbolic importance of geographic location itself in shaping market expectations and perceptions of artistic legitimacy. It is often noted that galleries and collectors often interpret location (e.g., a studio in Berlin or New York) as a proxy for quality and seriousness (Velthuis, 2005). Hellmanzik (2010) demonstrates that artworks produced in prominent creative hubs such as Paris and New York systematically command price premiums, up to 43% in certain periods, even after controlling for artist reputation and artwork characteristics. These cluster premiums are attributed to spillover effects, higher-quality peer environments, and greater market visibility. Similarly, Kelly and O'Hagan (2007) find persistent geographic clustering among elite Western visual artists and argue that such clustering reflects both the endogenous nature of artistic innovation and its predictive power for market success. Sociological studies complement this view by framing place as a socially constructed form of cultural capital. For instance, geographic co-location fosters artistic schools not only through collaboration and shared aesthetics, but also by lending a symbolic brand to artists. Being based in a "culturally loaded" city functions as a shorthand for legitimacy and seriousness, especially in opaque or nascent markets (Oberlin and Gieryn, 2015).

In the context of digital art markets like NFTs, where traditional forms of institutional validation (e.g., gallery representation or curated exhibition records) are absent, geographic location may continue to serve as a salient perceptual filter. We extend this logic by showing that even artists from minority groups, who might otherwise face recognition gaps, can obtain valuation premiums if located in major cultural cities. This premium likely stems from implicit associations with global art networks, professional infrastructure, and international standards of artistic practice. In short, location functions as a proxy for market readiness and cultural credibility. Taken together, these insights suggest that the geography of artistic production continues to shape how value is constructed and perceived, even in decentralized, digital art markets like NFTs.

Hypothesis 2c. *In the NFT art market, providing geographic information that signals the artist's presence in culturally prominent location is associated reduces sales disparities for minority artists.*

3.4 Data and Variables

To empirically evaluate the hypotheses, we collect comprehensive data from SuperRare[4] covering the period from April 2018 to December 2023. This dataset includes approximately 27,000 sales and 90,000 offers and bids, encompassing 40,000 NFTs created by 2,500 artists. We then construct variables at the artist, NFT, and transaction levels. Details for the each category of variables are outlined in Table 3.1. Summary statistics for each variable are reported in Table 3.2

Table 3.1: Variable Descriptions

Variable	Description
Dependent Variables	
Minting	A binary indicator variable representing whether an NFT successfully sells on the primary market.
Sales Price ($)	The price at which an NFT is sold, denominated in US dollars.
Artist-level Variables	
Female	A binary indicator variable representing whether an NFT artist identifies as female (1 if yes, 0 otherwise).

[4]https://superrare.com/

Continued on next page

$Non-White$	A binary indicator variable representing whether an NFT artist identifies as non-White (1 if yes, 0 otherwise).
$Artist\,Sales\,(\$)$	The historical sales volume of NFTs by an artist, measured in U.S. dollars (logarithmic).
$Artist\,NFTs\,Minted$	The number of NFTs minted by the artist (logarithmic).
$Sig.Qual.Exhibition$	A binary indicator variable representing whether an NFT artist reports previous exhibition(s) (1 if yes, 0 otherwise). Examples of notable art fairs include the Venice Biennale, Art Basel, Frieze Art Fair, Berlin Biennale, and NADA, among others.
$Sig.Qual.Clien$	A binary indicator variable representing whether an NFT artist mentions previous or current clients (1 if yes, 0 otherwise). Examples of such clients include Warner Bros, Perkins and Will, Paramount Pictures, Apple, LG Display, Adobe, Sony Music, Gibson, and LVMH, among others.
$Sig.Qual.Location$	A binary indicator variable denoting whether an NFT artist is currently based in one of the 42 advanced economies as designated by the International Monetary Fund (IMF). The list of these economies is available at https://www.imf.org/en/Publications/WEO/weo-database/2023/April/groups-and-aggregates.
$Artist\,Word\,Count$	The number of words used in an artist's self-description on Superrare (logarithmic).
$Artist\,URL$	A binary indicator variable representing whether an external URL, such as those from Twitter, Facebook, TikTok, YouTube, Discord, or a personal website, exists for an NFT artist on Superrare (1 if yes, 0 otherwise).
NFT-level Variables	
$NFT\,Size$	The file size of an NFT, measured in kilobytes (logarithmic).
$NFT\,Word\,Count$	The number of words used in the description of an NFT on Superrare (logarithmic).
$NFT\,ResNet\,PCA\,1$	The first principal component dimension extracted from ResNet-101 embeddings for an NFT.
$NFT\,ResNet\,PCA\,2$	The second principal component dimension extracted from ResNet-101 embeddings for an NFT.
$NFT\,ResNet\,PCA\,3$	The third principal component dimension extracted from ResNet-101 embeddings for an NFT.

	Continued on next page
NFT # of Bid/Offer	The cumulative number of bids and offers received for an NFT.
NFT Age	The number of days elapsed after an NFT is listed on the primary market (logarithmic).
NFT Genre	Binary indicator variables representing various genres associated with an NFT, such as Abstract, Illustration, Surrealism, Glitch, Generative, Portrait, Minimalism, Expressionism, Conceptual, Cubism, Realism, Impressionism.
NFT Media	Binary indicator variables representing various media types associated with an NFT, such as Drawing, Painting, Video, Photography, Collage, Sculpture, Installation.
Transaction-level Variables	
Transaction Fee in ETH	The total fees paid to Superrare in Ethereum, including gas fees and transaction fees (logarithmic).
Auction	A binary indicator variable representing whether a primary or secondary sale is conducted via auction (1 if yes, 0 otherwise).

Table 3.2: Summary Statistics

	N	Min	Median	Max	Mean	SD
Dependent Variables						
Minting	35.3K	0.00	1.00	1.00	0.68	0.47
Sales Price ($)	23.9K	0.01	904.35	1.1M	4.8K	18.1K
Artist-level Variables						
Female	35.3K	0.00	0.00	1.00	0.17	0.38
Non-White	35.3K	0.00	0.00	1.00	0.20	0.40
Artist Sales ($)	35.3K	0.00	6.2K	3.6M	50.5K	162.5K
Artist NFTs Minted	35.3K	1.00	18.00	852.00	53.3	93.5
Sig.Qual.Exhibition	35.3K	0.00	0.00	1.00	0.05	0.21
Sig.Qual.Client	35.3K	0.00	0.00	1.00	0.39	0.49
Sig.Qual.Location	35.3K	0.00	0.00	1.00	0.44	0.50
Artist Word Count	35.3K	1.00	27.00	845.00	52.2	73.4
Artist URL	35.3K	0.00	1.00	1.00	0.93	0.25
NFT-level Variables						
NFT Size	35.3K	0.00	12.2M	260.0M	18.4M	17.9M
NFT Word Count	35.3K	1.00	24.00	15.8K	36.6	95.2
NFT ResNet PCA 1	35.3K	-3.09	-0.16	6.37	0.03	1.58
NFT ResNet PCA 2	35.3K	-3.62	-0.02	4.88	0.01	1.16
NFT ResNet PCA 3	35.3K	-3.68	-0.04	3.95	-0.02	1.10

NFT # of Bid/Offer	23.9K	0.00	1.00	39.00	2.53	2.90
NFT Age	23.9K	0.00	7.00	1.8K	57.9	138.0
NFT Genre	35.3K	Omitted for brevity				
NFT Media	35.3K	Omitted for brevity				
Transaction-level Variables						
Transaction Fee in ETH	35.3K	0.00	0.01	0.19	0.01	0.01
Auction	23.9K	0.00	0.00	1.00	0.24	0.43

Continued on next page

Note. Values are rounded: K = thousands, M = millions. Standard deviations are in the final column.

3.4.1 Dependent Variables: Art Valuation Measures

Art valuation is notoriously challenging, even for seasoned art historians (Whitaker and Abrams, 2023). In the field of art economics, art prices often serve as a proxy for how collectors value a piece of art. The study of the NFT art market presents a unique advantage: unlike the traditional art auction market, the transaction history of an artwork can be fully traced. This traceability allows for the construction of three distinct art valuation measures, which are visualized in the flowchart presented in Figure 3.1. The following sections provide a detailed explanation of each of these valuation measures.

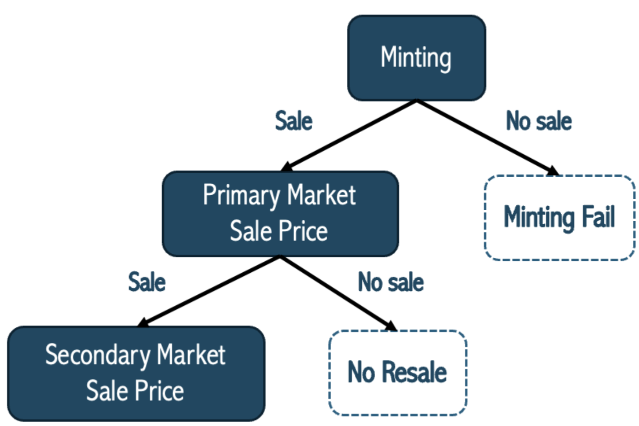

Figure 3.1: Flow Chart of Art Valuation Measures

- **Minting**. To sell NFTs, the first step for an NFT artist is to register ("mint") their NFT to a smart contract[5], which functions like a vending machine for NFTs (Oh et al., 2022). Once this registered ("minted") NFT is initially sold to a collector, the successful transaction is commonly referred to as a "mint-out." This valuation measure is crucial because the potential for an NFT art sale depends on whether it mints out in the first place. In our sample, the average mint-out rate is approximately 68% (about 27,000 out of 40,000 NFTs), which is higher than the 62% observed in the traditional art auction market (Bocart et al., 2022). Unlike the traditional art auction market, where some artworks are those created hundreds of years ago, the average age of NFT art in our sample is less than two years.

- **Primary Market Sale Price**. The primary market refers to the process by which NFTs are initially created on the blockchain ("minted") and sold to collectors. Once an NFT mints out, it acquires a primary sale price, a widely used proxy for art valuation in prior studies on art economics, such as Adams et al. (2021), Lovo and Spaenjers (2018), and Renneboog and Spaenjers (2013). It is important to note that we construct a winsorized sale price measure at both the 1% level. This approach is necessary due to the well-established *superstar effect* in art economics, wherein the top 1% of artists account for a significant portion of the total art auction market share. Given that our primary objective is to estimate general gender and racial gap in the NFT art market, we aim to mitigate the effects of this skewed distribution by adjusting the top and bottom extremes of sales volume.

[5] A Smart Contract is a self-executing contract with the terms of the agreement directly written into code. It automatically enforces and executes the agreed-upon conditions once these conditions are met, eliminating the need for intermediaries. Smart contracts are typically deployed on blockchain platforms, ensuring transparency, security, and immutability.

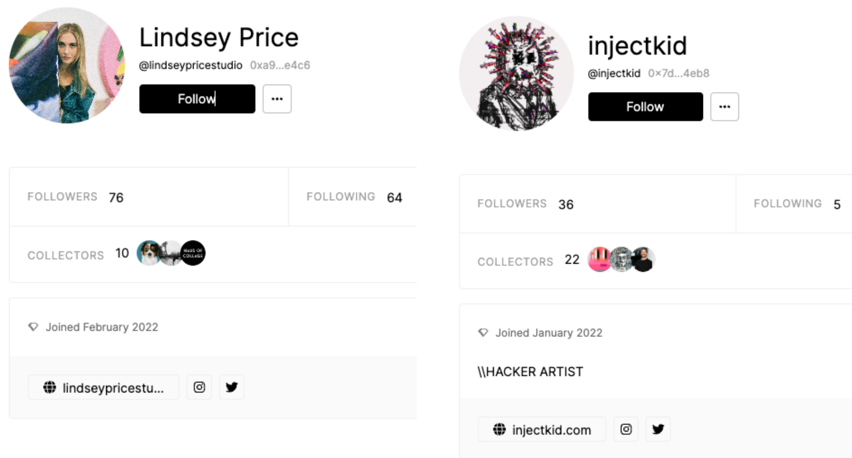

(a) Identity-classified NFT Artist (b) Anonymous NFT Artist

Note: In Panel (a), Lindsey Price, a Los Angeles-based NFT artist, identifies herself as female and was born in Boston, Massachusetts. During the data collection process, we categorize Lindsey Price as a White female. (Link: Lindsey Price's Profile). In Panel (b), the artist pseudonym injectkid remains unidentifiable. injectkid entirely conceals its identity, even across external platforms such as Instagram, X, and other NFT markets. Therefore, we label it as ”anonymous” and omit it from our analysis. (Link: injectkid's Profile).

Figure 3.2: Disclosure of Identity in Superrare

3.4.2 Independent Variables: Gender and Race

Since we estimate how NFT collectors value NFT artworks contingent upon NFT artists' gender and ethnicity background, $Female$ and $non-White$ become our independent variables, serving White male as a baseline group. Classifying the gender and race of each NFT artists entail three steps. First, we manually classify them by searching their profile in Superrare, social networks such as X, Instagram, Meta, YouTube, and their personal websites. Second, we use Namsor software[6] to do name-to-gender and name-to-ethnicity inference, based on their first and last name. For example, Namsor algorithm classifies NFT artist Jeehyun Kim and Juan Mingarro as Asian and Hispanic, with the probablity of

[6]https://namsor.app/api-documentation/. Namsor is a widely used software (See, e.g., Yang et al. (2022)

97% and and 94%, respectively. We drop samples if the most dominant probability of classification is below 80% and if an NFT artist does not reveal their identity. Lastly, if there is a conflicting results between our inference and Namsor's name-based inference, we double check our manual classification, which rarely happens. The criteria and methodology for defining gender and race are also illustrated with examples in Figure 3.2.

3.4.3 Moderating Variables: Signal of Quality Measures

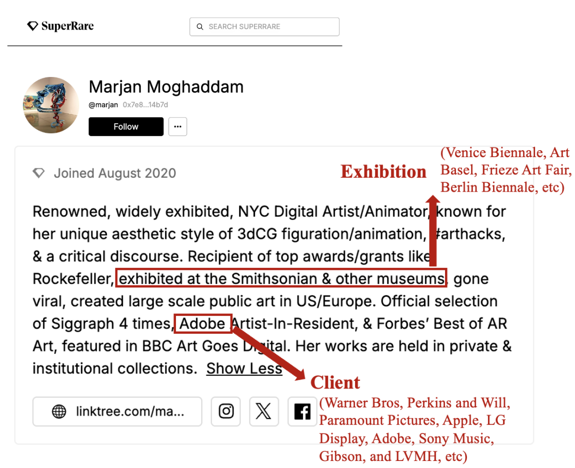

Note. Source: https://superrare.com/marjan.

Figure 3.3: Example of Artist Self-curation in Superrare

In SuperRare, artists are responsible for curating their own profiles. To capture signals of artistic quality, we construct three measures: $Sig.Qual.Exhibition$ indicates whether an artist shares information about their previous participation in art fairs or exhibitions. As shown in Figure 3.3, for instance, Marjan Moghaddam notes that she has exhibited at the Smithsonian and other major museums. $Sig.Qual.Client$ reflects whether the artist discloses any notable past or current clients. In the same example, the artist indicates that she was an in-house artist for Adobe. We record client affiliations across various industries, including entertainment (e.g., Warner Bros) and electronics (e.g., LG), covering hundreds of recognizable brands. $Sig.Qual.Location$ captures whether the artist is based in a culturally prominent location, such as France, the United Kingdom, or the United States. While some artists originate from developing countries, many indicate current residence in these more established cultural centers. Following Li et al. (2024), this variable is constructed using references such as the European Capital of Culture, 39 UN City 2016, City Mayors EU 500, City Mayors World 300, and other global cultural hubs characterized by a high concentration of museums, galleries, and auction houses.

3.4.4 Control Variables

To accurately estimate the impact of NFT artists' gender and ethnicity on art collectors' valuations, it is crucial to control for potential confounding variables. Therefore, we incorporate artist-, NFT-, and transaction-level characteristics within our analysis.

Artist-level Controls

We consider several characteristics that could influence an NFT artist's profile, as displayed on SuperRare. First, we measure each artist's historical sales in U.S. dollars. As explored by Velthuis (2013), the relationship between prices and their symbolic meanings in the contemporary art market suggests that prices can act as indicators of an artist's reputation and market standing. Thus, we use $Artist\,Sales(\$)$ as a proxy for the artist's reputation. Additionally, we control for the extent of the artist's self-description, quantified by the word count in their SuperRare profile, denoted as $Artist\,Word\,Count$. Furthermore, we account for whether the artist reveals additional information through social networks such as Instagram, X, TikTok, or a personal website, captured by $Artist\,URL$. For instance, as illustrated in Figure 3.2, NFT artists Lindsey Price and injectkid share their personal website and social media links via clickable external URLs.

NFT-level Controls

Given that each NFT artwork is inherently unique and ”non-fungible,” we control for various NFT-level characteristics. Even NFTs created by the same artist can exhibit significant differences. As shown in Figure 3.4, NFT collectors on SuperRare can access metadata such as file size ($NFT\,Size$) and pixel dimensions [7], as well as medium. We also control for the word count of the artist's description of the NFT, $NFT\,Word\,Count$, for reasons discussed in the previous section.

Medium	Image (JPEG)
File Size	7.5 MB
Dimensions	2350×2938
Contract Address	0xdf7...3ff0b
Token Standard	ERC-721
Blockchain	Ethereum

Note. Refer to https://superrare.com/artwork/eth/0xdf75cf0a3dad9bf64c87d4902d1ffdbfaaa3ff0b/2 for more information. This is an random example from Superrare.

Figure 3.4: Metadata of an NFT in Superrare

The visual characteristics of an artwork are well-documented predictors of sale price, as highlighted by Aubry et al. (2023). To capture these features, we downloaded each NFT file in PNG format (approximately 500 gigabytes for around 40,000 NFTs). We then used a pre-trained ResNet101 model (He et al., 2016) in PyTorch to extract embedding vectors for the images in our sample. Specifically, we stored the output from layer 4 of the ResNet101 model, resulting in a 2048x7x7 feature volume. After flattening this volume, we

[7]Dimensions were excluded from our regression analysis due to high collinearity with file size.

obtained a vector of size 100,352, which we compressed into three dimensions using principal component analysis (PCA), yielding $NFT\,ResNet\,PCA\,1$, $NFT\,ResNet\,PCA\,2$, and $NFT\,ResNet\,PCA\,3$. Additionally, we control for each NFT's artistic style, $NFT\,Genre$, and medium, $NFT, Media$, acknowledging that certain genres, such as Impressionism, may be more favored by collectors.

Transaction-level Controls

For each NFT transaction, buyers are required to pay transaction fees, including gas fees and commissions to SuperRare, in Ethereum, on top of the purchase price. These fees are likely to influence buyers' willingness to purchase and their willingness to pay, so we include $Transaction\,Fee\,in\,ETH$. Furthermore, there are two primary methods of purchasing NFTs: bidding in an auction or making an offer available 24/7, captured by a dummy variable, *Auction*. These purchasing methods can potentially influence pricing dynamics. Auctions, creating a sense of urgency and competition, may drive prices higher, especially when multiple bidders are interested in the same NFT. In contrast, the offer system provides a more stable pricing environment, allowing buyers time to consider their purchases, potentially leading to lower prices compared to the auction model.

3.5 Empirical Model

3.5.1 Model Specification

To formally assess how NFT collectors' valuation over NFT artworks vary by artists' gender and race (Hypothesis 2), we estimate regression specifications of the following form:

$$\begin{aligned}\text{NFT Valuation Measures}_{i,t} = \alpha + \beta_1 \cdot \text{Female}_j + \beta_2 \cdot \text{non-White}_j + \gamma \cdot \text{Artist Controls}_{j,t} + \\ \theta \cdot \text{NFT Controls}_{k,t} + \lambda \cdot \text{Transaction Controls}_{i,t} + \tau_t + \varepsilon_{i,t}\end{aligned} \tag{3.1}$$

where β_1 and β_2 are the coefficients for the main independent variables Female_j and non-White_j, assessing the impact of the artist's gender and race on the NFT artwork valuation. NFT Valuation $\text{Measures}_{i,t}$ are either Sales $\text{Success}_{i,t}$ or Sales $\text{Price}_{i,t}$, depending on model specification. Accordingly, the regression model also depends on $\text{Minting}_{i,t}$ (logistic regression) and Sales $\text{Price}_{i,t}$ (OLS regression). γ, θ, and λ are vectors of coefficients for the respective control variables. Artist $\text{Controls}_{j,t}$ control for both time-variant and time-invariant factors related to the artist that may influence the artwork's price, ensuring that the estimated effects of gender and race are not confounded by these variables. NFT $\text{Controls}_{k,t}$ control for various attributes of the artwork that can affect the price, such as visual features, genre and media, allowing for a more accurate estimation of the effects of the artist's characteristics. Transaction $\text{Controls}_{i,t}$ control for transaction-specific factors such as fees and auction indicator to isolate their impact on the artwork's price. τ_t are year-month fixed effects to address common time trends. The definition for each variable is described in Table 3.1.

We proceed by estimating the moderating effect of artist information provision (Hypothesis 3) using the following regression model:

$$
\begin{aligned}
\text{NFT Valuation Measures}_{i,t} = \alpha &+ \beta_1 \cdot \text{Female}_j + \beta_2 \cdot \text{non-White}_j + \gamma \cdot \text{Artist Controls}_{j,t} \\
&+ \beta_3 \cdot \text{Female}_j \times \text{Information Provision}_{j,t} \cdot \\
&+ \beta_4 \cdot \text{non-White}_j \times \text{Information Provision}_{j,t} \\
&\theta \cdot \text{NFT Controls}_{k,t} + \lambda \cdot \text{Transaction Controls}_{i,t} + \tau_t + \varepsilon_{i,t}
\end{aligned} \tag{3.2}
$$

Most variables in this model are consistent with those in Equation 3.1, with the addition of the sets of moderating variables Information Provision$_{j,t}$, including quantitative measure, Artist Word Count$_{j,t}$, and qualitative measures, Sig.Qual.Exhibition$_{j,t}$, Sig.Qual.Clients$_{j,t}$, and Sig.Qual.Location$_{j,t}$.

3.6 Results

3.6.1 Gender and Racial Diversity in NFT Market Entries

We begin by examining the composition of the NFT market concerning the gender and race of NFT artists. Our findings indicate that the NFT art market is predominantly dominated by white males in terms of supply volume: about 23% of NFT artists are women, and only about 24% are non-white (Figure 3.5). While these statistics suggest a significant gender and racial imbalance, it is essential to consider them in the context of the traditional art market. From 2000 to 2017, only about 4% of the works sold at auction were by women (Cameron et al., 2019). Similarly, between 2012 and 2022, works by male artists sold for a total of $102 billion, compared to just $6.7 billion (6.5%) for works by women (Bocart et al., 2022). In terms of racial disparity, only 16% of sales in the U.S. art market were attributed to non-white artists during the same period (LeBlanc and Sheppard, 2022). In summary, although the NFT market exhibits a notable gender and racial imbalance, it is less severe than that observed in the traditional art market.

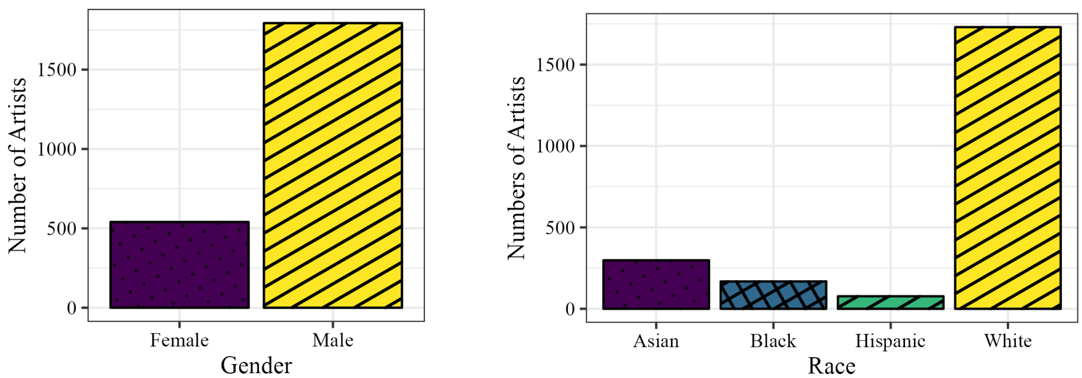

(a) # of Artists by Gender (b) # of Artists by Race

Note: The distribution of additional statistics, including (i) the number of NFTs categorized by gender and race, and (ii) the number of NFT sales categorized by gender and race, exhibits a high level of consistency with the distributions presented in Panel A and Panel B of this figure.

Figure 3.5: # of NFT Artists by Gender and Race

This difference is meaningful because the NFT market is often positioned as a decentralized and democratized alternative to legacy systems that have long been criticized for gatekeeping and exclusion. The relatively higher representation of women and non-white artists in the NFT space suggests a potential shift in access and participation. Understanding the extent and limitations of this shift is essential for assessing whether decentralized technologies genuinely foster inclusion, or whether they replicate existing inequalities in new forms. As such, examining the demographic composition of NFT creators provides a foundation for evaluating how far the promises of equity and decentralization could be realized in practice.

3.6.2 Gender and Racial Disparity: Model-free Evidence

As discussed in the Data section, we measure NFT valuations using the probability of sales (”Minting”) and primary sales price. Below are the model-free descriptive statistics for each measure based on the pooled sample:

Minting. Female and male NFT artists have an average sales probability of 66% and 69%, respectively, in the primary market. White and Asian NFT artists have sales probabilities of 70% on average, while Black and Hispanic artists have probabilities of 61% and 37%, respectively. The sales probability for Hispanic NFT artists, at 37%, is particularly notable as it is approximately half that of White artists.

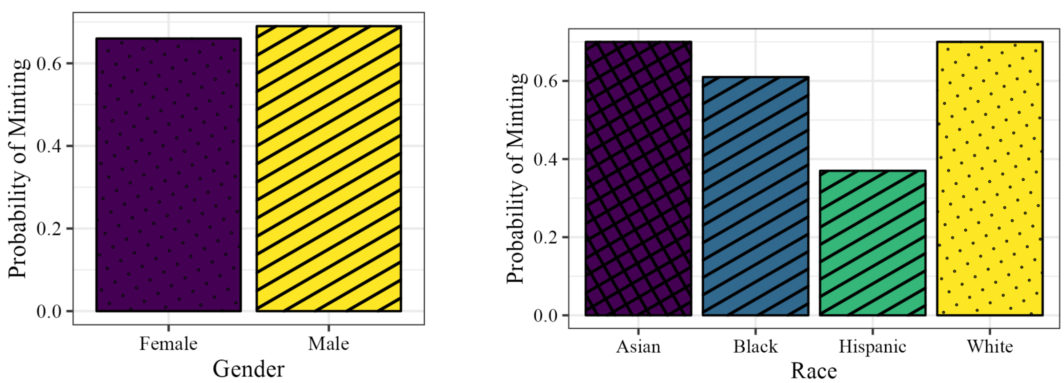

(a) Minting Probability by Gender

(b) Minting Probability by Race

Figure 3.6: Minting Probability by Gender and Race

Primary Sales Price. Female and male NFT artists sell their works for an average of \$3873 and \$4109, respectively, in the primary market. White, Black, Hispanic, and Asian artists sell their NFTs for \$4025, \$2472, \$3001, and \$6208, respectively. Figure 3.8 illustrates the temporal trend of price disparities over time. As shown in Panel a, there is no noticeable gender disparity between female and male artists during our observational period, which covers a full boom-bust cycle. However, significant racial disparities exist: Asian artists experienced a substantial price escalation during booming market periods, while Black NFT artists consistently had the lowest average prices. Overall, racial disparities diminish as the market transitions to bust periods.

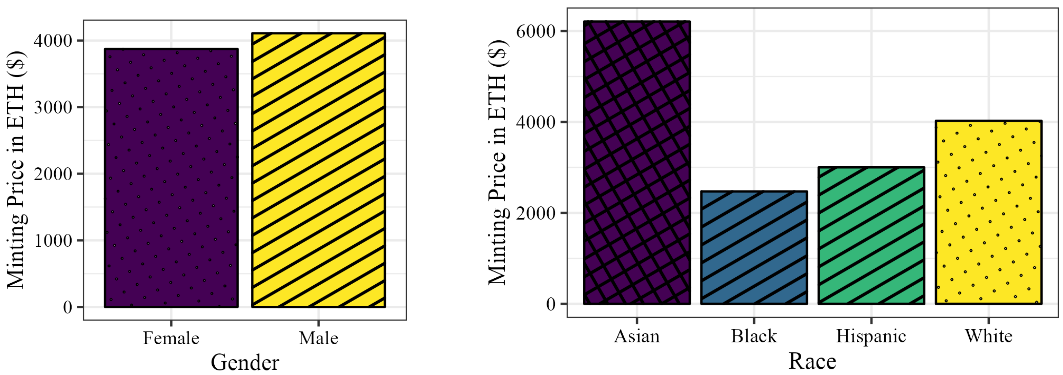

(a) Primary Price by Gender (b) Primary Price by Race

Figure 3.7: Primary Price by Gender and Race

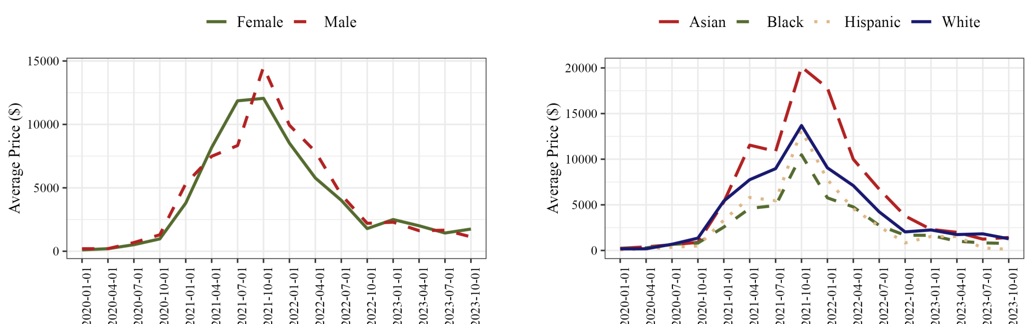

(a) Price Trend by Gender (b) Price Trend by Race

Note. NFT prices ($) are winsorized at 1% to reduce the influence of extreme outliers.

Figure 3.8: NFT Price Trend by Gender and Race

3.6.3 Gender and Racial Disparity: Model-based Evidence

In this section, we present the formal estimation of NFT collectors' valuations for NFT artworks, as shown in Table 3.3 and Table 3.4.

Valuation Measure: Minting Probability

Column (1) of Table 3.3 presents the baseline logistic regression results examining the relationship between artist demographics and the likelihood of an NFT being sold on the primary market. The dependent variable, *Minting*, is a binary indicator equal to 1 if the NFT is sold and 0 otherwise. All regressions include a comprehensive set of artist-, NFT-, and transaction-level control variables, as well as fixed effects for NFT genre, media, and time, to account for potential confounding factors.

The coefficient on the *Female* indicator is -0.010 with a standard error of 0.035, suggesting no statistically significant difference in the probability of sale between male and female artists. The corresponding odds ratio is $\exp(-0.010) \approx 0.990$, indicating a 1% decrease in the odds of sale for female artists, though this effect is not statistically distinguishable from zero.

The *non-White* coefficient is -0.308 (standard error = 0.035), statistically significant at the 1% level. This implies that the odds of sale for non-White artists are approximately 26.5% lower than those for White artists, with an odds ratio of $\exp(-0.308) \approx 0.735$. To interpret this effect in terms of probability, we approximate the marginal effect at the mean (MEM). Given the sample average sale probability of 0.68, the odds at the mean are $\frac{0.68}{1-0.68} = 2.125$. Applying the odds ratio yields adjusted odds of $2.125 \times 0.735 = 1.561$, which translates back to a predicted probability of $\frac{1.561}{1+1.561} \approx 0.610$. This suggests that identifying as a non-White artist is associated with an approximate 6.9 percentage point reduction in the likelihood of sale, holding other factors constant.

Table 3.3: NFT Artist Gender, Race, and Probability of Sales

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.010	−0.130***	−0.015	−0.076*	−0.106**	−0.199***
	(0.035)	(0.045)	(0.035)	(0.044)	(0.045)	(0.053)
non-White	−0.308***	−0.205***	−0.323***	−0.277***	−0.286***	−0.207***
	(0.035)	(0.043)	(0.036)	(0.043)	(0.038)	(0.047)
Artist Word Count	−0.001***	−0.001**	−0.001***	−0.001***	−0.001***	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig.Qual.Exhibition	0.129**	0.140**	0.039	0.130**	0.132**	0.018
	(0.066)	(0.066)	(0.079)	(0.066)	(0.066)	(0.080)
Sig.Qual.Client	0.089***	0.090***	0.090***	0.073**	0.086***	0.070*
	(0.031)	(0.031)	(0.031)	(0.036)	(0.031)	(0.036)
Sig.Qual.Location	−0.024	−0.028	−0.023	−0.025	−0.053	−0.051
	(0.029)	(0.029)	(0.029)	(0.029)	(0.033)	(0.032)
Female × Artist Word Count		0.002***				0.002***
		(0.001)				(0.001)
non-White × Artist Word Count		−0.002***				−0.003***
		(0.000)				(0.001)
Female × Sig.Qual.Exhibition			0.187			0.070
			(0.180)			(0.185)
non-White × Sig.Qual.Exhibition			0.311**			0.559***
			(0.145)			(0.156)
Female × Sig.Qual.Client				0.159**		0.019
				(0.070)		(0.078)
non-White × Sig.Qual.Client				−0.069		0.107
				(0.065)		(0.078)
Female × Sig.Qual.Location					0.236***	0.197***
					(0.070)	(0.071)
non-White × Sig.Qual.Location					−0.122	−0.093
					(0.099)	(0.100)

Continued on next page

Table 3.3 continued

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Sales	0.157***	0.155***	0.157***	0.156***	0.157***	0.155***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Artist NFTs Minted	−0.553***	−0.550***	−0.553***	−0.552***	−0.555***	−0.550***
	(0.013)	(0.014)	(0.013)	(0.013)	(0.013)	(0.014)
External URL	−0.184***	−0.192***	−0.184***	−0.186***	−0.179***	−0.190***
	(0.056)	(0.056)	(0.056)	(0.056)	(0.056)	(0.056)
NFT Size	0.004	0.005	0.004	0.005	0.006	0.005
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
NFT Word Count	−0.003***	−0.003***	−0.003***	−0.003***	−0.002***	−0.003***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Transaction Fee	5.105***	5.201***	5.094***	5.118***	5.088***	5.193***
	(1.764)	(1.763)	(1.765)	(1.765)	(1.765)	(1.765)
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
Year-Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	35 326	35 326	35 326	35 326	35 326	35 326
AIC	36 388.6	36 362.0	36 387.4	36 386.7	36 380.5	36 351.3
BIC	37 261.3	37 251.6	37 277.0	37 276.3	37 270.1	37 291.7
Log. Lik.	−18 091.3	−18 076.0	−18 088.6	−18 088.3	−18 085.2	−18 064.6

* p < 0.1, ** p < 0.05, *** p < 0.01

Note. Robust standard errors are clustered at the artist level. Continuous control variables are winsorized at the 1% level to reduce the influence of extreme outliers.

Columns (2) through (6) in Table 3.3 extend the baseline specification by incorporating interaction terms that examine our Hypothesis 2 that enhanced provision of artist information is associated with reduced gender and racial disparities in art valuations within the NFT market.

Column (2) introduces interactions with *Artist Word Count*, the number of words in the artist's self-description. The main effect of *Artist Word Count* is negative and statistically significant (-0.001, $p < 0.01$), suggesting that longer bios are generally associated with a lower probability of sale, which is quite an unexpected result. However, the interaction terms reveal important subgroup variation. Specifically, the effect is less negative (even positive) for female artists ($Female \times Artist\ Word\ Count = 0.002$, $p < 0.01$), indicating that detailed self-presentations may benefit female artists in overcoming baseline disadvantages. Conversely, the interaction term for non-White artists is negative and significant (-0.002, $p < 0.01$), suggesting that lengthier bios further depress sale probability for non-White artists. This divergence may reflect variation in how audiences interpret self-promotional efforts as a function of the artist's identity.

Column (3) examines the moderating role of providing information on prior exhibitions, captured by *Sig.Qual.Exhibition*. The main effect is positive and significant (0.140, $p < 0.05$), indicating that artists who report a notable exhibition history have a higher likelihood of selling their NFTs. The interaction with *non-White* is positive and significant, suggesting a larger increase in log-odds of sale for non-White artists, helping mitigate baseline racial disadvantages. In contrast, the interaction with *Female* is not statistically significant, implying a more modest or inconsistent effect for female artists.

Column (4) introduces *Sig.Qual.Client*, a dummy indicating whether the artist mentions past or present commercial clients. The main effect is positive and statistically significant (0.090, $p < 0.01$), reflecting that commercial credibility increases the probability of sale.

The interaction with *Female* is also positive and significant, referencing prominent clients is associated with a greater increase in the log-odds of sale for female artists. However, the interaction with *non-White* is negative and insignificant, suggesting that referencing clients does not provide a comparable advantage for non-White artists.

Column (5) includes interactions with *Sig.Qual.Location*, which indicates whether the artist is based in a cultural hub. The main effect of this variable is not statistically significant, implying no average difference in sale probability based solely on location. However, the interaction with *Female* is positive and significant, suggesting that among female artists, being based in a cultural hub is associated with a higher log-odds of sale for female artists. The interaction with *non-White* is not significant.

Column (6) includes all the interaction terms jointly. The main effects of *Artist Word Count* and *Sig.Qual.Client* remain statistically significant, indicating that on average, longer bios and client references have meaningful (though complex) relationships with sale probability. The interaction *non-White* × *Sig.Qual.Exhibition* grows stronger, reinforcing the importance of institutional validation in offsetting racial disparities. The interaction *Female* × *Artist Word Count* remains significant and positive, and *Female* × *Sig.Qual.Location* also remains significant, supporting earlier findings. Other interaction term, *Female* × *Sig.Qual.Client*, is not statistically significant when all moderators are included.

Valuation Measure: Sales Price

Table 3.4 presents results from OLS regressions where the dependent variable is the natural logarithm of the NFT's primary market sales price, denominated in U.S. dollars.

Table 3.4: NFT Artist Gender, Race, and Primary Sales Price

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.001	0.004	−0.002	0.001	−0.002	0.000
	(0.007)	(0.009)	(0.007)	(0.009)	(0.009)	(0.011)
non-White	−0.006	−0.029***	−0.013*	−0.015*	−0.018**	−0.041***
	(0.007)	(0.009)	(0.007)	(0.009)	(0.008)	(0.010)
Artist Word Count	0.000***	0.000**	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
FairExhibition	−0.014	−0.018	−0.055***	−0.014	−0.012	−0.050***
	(0.015)	(0.015)	(0.018)	(0.015)	(0.015)	(0.018)
Client	−0.003	−0.003	−0.002	−0.007	−0.002	−0.002
	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)
Premium Location	−0.008	−0.007	−0.007	−0.007	−0.014**	−0.013**
	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Female × Artist Word Count		0.000				0.000
		(0.000)				(0.000)
non-White × Artist Word Count		0.001***				0.000***
		(0.000)				(0.000)
Female × Sig.Qual.Exhibition			0.082**			0.086**
			(0.039)			(0.039)
non-White × Sig.Qual.Exhibition			0.128***			0.094***
			(0.036)			(0.039)
Female × Sig.Qual.Client				−0.002		0.003
				(0.014)		(0.016)
non-White × Sig.Qual.Client				0.029**		−0.001
				(0.015)		(0.016)
Female × Sig.Qual.Location					0.002	0.002
					(0.014)	(0.014)
non-White × Sig.Qual.Location					0.074***	0.076***
					(0.018)	(0.018)

Continued on next page

Table 3.4 continued

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Sales	0.323***	0.324***	0.323***	0.323***	0.323***	0.323***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Artist NFTs Minted	−0.330***	−0.330***	−0.330***	−0.330***	−0.330***	−0.330***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
External URL	0.014	0.015	0.014	0.014	0.012	0.013
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
NFT Size	0.012***	0.011***	0.011***	0.012***	0.011***	0.011***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
NFT Word Count	0.001***	0.001***	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
NFT BidOffer Num.	0.027***	0.027***	0.027***	0.027***	0.027***	0.027***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
NFT Age	−0.012***	−0.012***	−0.012***	−0.012***	−0.012***	−0.013***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Transaction Fee	2.598***	2.582***	2.598***	2.591***	2.593***	2.575***
	(0.348)	(0.348)	(0.348)	(0.349)	(0.348)	(0.348)
Auction	−0.084***	−0.083***	−0.084***	−0.084***	−0.083***	−0.083***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
Year-Month FE	YES	YES	YES	YES	YES	YES
Num.Obs.	23 919	23 919	23 919	23 919	23 919	23 919
R2	0.668	0.668	0.668	0.668	0.668	0.669
R2 Adj.	0.667	0.667	0.667	0.667	0.667	0.667

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Importantly, this outcome is observed only for NFTs that were successfully sold on the primary market. Therefore, this analysis is conditional on minting—i.e., it captures valuation dynamics among NFTs that cleared the initial barrier of market acceptance. As such, the coefficients reflect price effects given that the NFT was minted, rather than providing insight into overall market inclusion.

In the baseline model (Column 1), neither gender nor race has a statistically significant association with sale price. The coefficient on *Female* is close to zero (-0.001), and the coefficient on *non-White* is slightly negative (-0.006), but not significant. These estimates suggest that, on average, conditional on sale, NFTs by female or non-White artists are priced similarly to those by White male artists. However, as additional artist-level moderators are introduced in Columns (2) through (6), significant effects emerge, particularly for race.

In Column (2), the main effect of *non-White* becomes statistically significant and more negative (-0.029, $p < 0.01$). This corresponds to a reduction of approximately 2.9% in sale price for non-White artists, holding other factors constant. However, the interaction term between *non-White* and *Artist Word Count* is positive and significant (0.001, $p < 0.01$), indicating that narrative elaboration (e.g., longer artist bios) helps counteract the pricing penalty. This suggests that detailed self-presentation may serve as a strategic tool to increase perceived value, particularly for racially marginalized artists. The interaction with *Female* remains insignificant.

In Column (3), we examine signaling how prior participation in exhibitions moderates demographic pricing effects. The interaction between *non-White* and *Sig.Qual.Exhibition* is positive and highly significant (0.128, $p < 0.01$), corresponding to an estimated 13.6% price increase for non-White artists who have exhibited at art fairs. For female artists, the interaction is also positive and significant (0.082, $p < 0.05$), which translates to an approximate 8.5% price premium. These results suggest that institutional validation has a disproportionately positive effect on the valuation of underrepresented groups, likely by enhancing buyer confidence in artistic quality or signaling elite market recognition.

In Column (4), we introduce interactions with commercial client affiliation. The interaction between *non-White* and *Client* is positive and marginally significant (0.029, $p < 0.05$), implying a 3% price increase for non-White artists who report reputable clients. However, this interaction becomes insignificant in the full model in Column (6), suggesting that its effect may be confounded or moderated by other artist-level signals. No significant interaction is found between *Female* and *Client*, and the main effect of client affiliation remains close to zero throughout.

Column (5) introduces geographic context through the *Sig.Qual.Location* variable. The main effect of being based in a culturally prominent location is negative and significant for the baseline group (-0.014, $p < 0.05$), suggesting a 1.4% price reduction for White male artists in advanced economies. However, the interaction with *non-White* is strongly positive and significant (0.074, $p < 0.01$), corresponding to an estimated 7.7% price premium for non-White artists based in such regions. This finding implies that non-White artists may signal enhanced visibility, infrastructure, or perceived legitimacy when they are located in a cultural hub. Again, the interaction with *Female* is not significant.

Finally, Column (6) includes all interaction terms jointly. The coefficient for *non-White* remains large and highly significant (-0.041, $p < 0.01$), indicating an average 4% lower price for non-White artists, even after accounting for other characteristics. Several interaction terms remain robust in this full specification: the moderating effects of *Artist Word Count* and *Sig.Qual.Exhibition* for non-White artists, and the effect of *Sig.Qual.Exhibition* for female artists, all remain significant and of similar magnitude. The interaction between *non-White* and *Sig.Qual.Location* also remains strong, indicating that location-based benefits for racially marginalized artists persist even when multiple forms of validation are included.

Taken together, the results from the minting and sales price models provide consistent evidence supporting Hypothesis 1: racial disparities exist in both minting and price valuation, while gender disparities are more nuanced and context-dependent. Non-White artists face a significant disadvantage in the likelihood of sale (minting) and, conditional on sale, receive lower prices than their White counterparts. In contrast, female artists do not face a statistically significant penalty in sale probability or price once other factors are controlled.

The results also support Hypothesis 2: several forms of information provision help reducing these disparities. For non-White artists, institutional exhibition and being based in advanced economies appear especially effective in mitigating racial penalties in both minting and pricing. For female artists, the interaction effects are less consistent overall, but certain forms of signaling, particularly advanced economy location and client affiliation, are associated with improved likelihood of sale. These findings suggest that while disparities persist, credible signals of artistic quality and market legitimacy, especially those rooted in institutional validation, can meaningfully offset structural disadvantages in decentralized art markets.

3.7 Robustness Check

To assess the robustness of our principal results, we conduct a series of supplementary analyses. Table 3.5 enumerates each potential concern, the methodological strategies employed to address it, and whether the resulting evidence corroborates or challenges our main conclusions.

Table 3.5: Summary of Robustness Checks

Concern	Test	Results	Table
Rarity rather than quality signal	Two-sample t-tests for variables (exhibition, client, location) across gender and race groups	Supported	Table 3.6
Mispricing by minority artists	Regression based only on NFT sales made via offers, not bids	Supported	Table 5.1, 5.2
Average effect driven by outliers	Regressions run excluding superstar artists	Supported	Table 5.3, 5.4
Selection bias	Balancing artist demographics using a matching approach	Supported	Table 5.5, 5.6

3.7.1 Alternative Mechanism: Rarity Rather than Quality Signal

In Hypothesis 2 we posit that minority NFT artists can attenuate valuation discounts by furnishing verifiable, credible curatorial signals of quality. A plausible rival explanation, however, is that the observed premium reflects rarity rather than perceived quality. Prior work shows that scarcity exerts a first-order influence on NFT liquidity and prices (Mekacher et al. 2022, Kaur Nagpal and Renneboog 2023, Hofstetter et al. 2023). Consequently, if the moderators in our model—$ArtistWordCount$, $Sig.Qual.Exhibition$, $Sig.Qual.Client$, and $Sig.Qual.Location$—differ systematically across demographic groups, collectors might simply reward the unusualness of those traits. For instance, if non-White artists seldom disclose exhibition histories or notable clients, the few who do could command a premium consistent with the "higher-bar" argument.

Table 3.6: Two-Sample t — Test Results for Artist-Information Variables

	Gender			Race		
Variable	Male	Female	p-value	White	Non-White	p-value
$ArtistWordCount$	42.79	39.16	0.16	39.99	40.06	0.98
$Sig.Qual.Exhibition$	5%	6%	0.61	5%	6%	0.94
$Sig.Qual.Client$	39%	39%	0.80	39%	38%	0.70
$Sig.Qual.Location$	40%	43%	0.22	50%	15%	0.00

Note: Means are reported for each subgroup. p-values are from two-sample t-tests of equal means. Variable definitions appear in Table 3.1.

Table 3.6 reports two-sample t-tests comparing these variables across subgroups. None of the curatorial signals differ significantly, except for $Sig.Qual.Location$: White artists are more likely to be based in advanced economies, which is not surprising. Because the other indicators show no systematic rarity, we conclude that our mitigation effects are unlikely to be driven by scarcity alone; rather, they reflect the informational value of credible self-curation.

3.7.2 Alternative Mechanism: Mispricing by Minority Artists

We interpret our principal results through a demand-side lens on gender- and race-based disparities. Building on the canonical distinction between taste-based and statistical discrimination, we first ask whether such disadvantages arise and then test whether providing richer, verifiable artist information attenuates them.

A potential threat to this interpretation is supply-side mispricing: when creators set starting prices in auctions, those prices can serve as anchors that shape subsequent bids, thereby biasing our estimates. To rule out this concern, we re-estimate the model after excluding auction transactions, in which artists have an option to determine the initial price. Auctions represent roughly 24% of the sample, so the revised analysis relies on the remaining 76% of offer-based sales, where price discovery is driven exclusively by buyer demand.

Table 5.1 presents the minting-probability estimates, and Table 5.2 reports the primary-sale price regressions. Across specifications, the coefficients closely track—and in several cases exceed—those in the pooled sample. For instance, in Column (3) of Table 5.1 the interaction term $Non{-}White \times Sig.Qual.Exhibition$ is 0.199 points larger than in Table 3.3. Likewise, in the primary-sale model (Column (3) of Table 5.2) all moderator coefficients increase except for $Female \times Sig.Qual.Exhibition$, which becomes statistically insignificant.

3.7.3 Sensitivity Check: Average Effect Driven by Outliers

Prior research in art economics documents a pronounced "superstar effect" in traditional art markets. Cameron et al. (2019) and Adams et al. (2021) show that a handful of top-ranked female artists account for the lion's share of auction revenues attributable to women. In the valuation domain, Bocart et al. (2022) report that, for transactions above $1 million, works by male artists command an 18.4% premium over those by female artists.

Because our objective is to estimate the average pricing impact of gender and race in the NFT art market, we remove such extreme "superstar" observations. Specifically, we sort artists by cumulative sales volume and exclude the top decile (10%) before conducting the analysis.

Table 5.3 replicates the minting regressions after removing the top 10 decile of artists by cumulative sales. The core results survive virtually unchanged. A sizable, negative $Non - White$ coefficient persists in every specification, while the main $Female$ coefficient continues to register only once extensive moderators are introduced. Crucially, the interaction terms that proxy for verifiable quality or effort, $ig.Qual.Exhibition$, $ig.Qual.Client$, and *ig.Qual.Location*, retain their moderating power. In several cases the mitigating effects are even larger in magnitude, indicating that information signals remain salient once the distorting influence of superstar listings is stripped from the sample.

Table 5.4 repeats the primary price regressions after dropping the top 10 percent of artists by cumulative sales. The core patterns are intact. The interaction terms capturing verifiable information signals, particularly *ig.Qual.Exhibition* and *ig.Qual.Client*, continue to attenuate that penalty. In some cases the moderating coefficients are even larger once superstar listings are removed, underscoring the consistent role of curatorial signals in narrowing racial disparities.

Overall, none of these shifts overturn our principal inference: racial disparities in minting probability are robust, and credible curatorial signals systematically narrow them. Removing the upper tail of the artist distribution therefore reinforces, rather than weakens, the interpretation that the observed gaps reflect broad market dynamics rather than the pricing of a small set of outlier careers.

3.7.4 Addressing Observable Selection Bias via Matching

A remaining concern is that the gender (or race) coefficient in our pooled regressions may be driven by observable sample composition effects rather than genuine price discrimination. Even with exhaustive artwork-, artist-, and transaction-level controls, the covariate distributions can still differ systematically across demographic groups. When this occurs, the ordinary least-squares (OLS) estimator must extrapolate into regions of the covariate space where the opposite group is thinly represented, making the estimate highly model-dependent. To attenuate this selection on observables, we apply propensity–score matching (PSM) following Rosenbaum and Rubin (1983).

Table 3.7: Summary Statistics Before Matching

Variables	**White** ($N = 1725$)	**Non-White** ($N = 541$)	**P-value**
$Female$	0.21 (0.41)	0.30 (0.46)	<0.001
$Artist\,Sales$ (\$)	20363.38 (48160.17)	16948.90 (47689.57)	0.149
$Artist\,NFTs\,Minted$	8.89 (17.79)	7.06 (19.48)	0.041
$Artist\,Word\,Count$	39.99 (51.27)	40.06 (52.57)	0.979
$Sig.Qual.Exhibition$	0.05 (0.21)	0.06 (0.22)	0.939
$Sig.Qual.Client$	0.39 (0.42)	0.38 (0.42)	0.704
$Sig.Qual.Location$	0.84 (0.37)	0.45 (0.50)	<0.001
$Artist\,URL$	0.95 (0.21)	0.94 (0.24)	0.126
$NFT\,Size$	23127159 (16275388)	20509684 (14309394)	0.001
$NFT\,Word\,Count$	47.85 (37.20)	52.33 (43.04)	0.019

Note. Standard errors are in parentheses. The p-value indicates the statistical significance of two-sample t-test results.

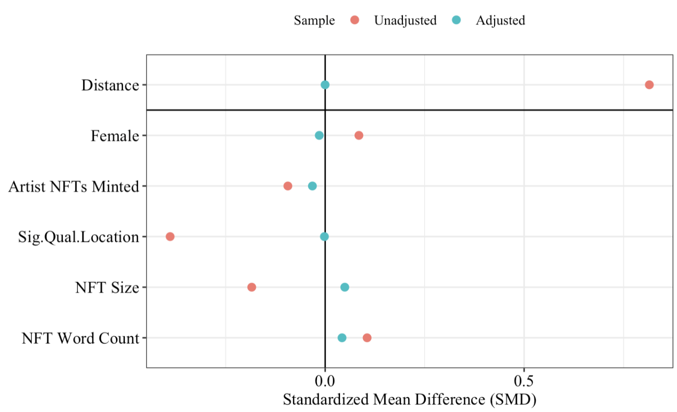

Figure 3.9: Covariate Balance Test

Let $p(X) = \Pr(\text{Non-White} = 1 \mid X)$ denote the probability that an NFT is created by a Non-White artist, conditional on the vector of observables X. We first compare group characteristics via t-tests (Table 3.7) and identify five covariates with p-values below 0.05: $Female$, $Artist\,NFTs\,Minted$, $Sig.Qual.Location$, $NFT\,Size$, and $NFT\,Word\,Count$. These variables enter a logit model used to estimate $p(X)$. Each Non-White artist is then paired with up to *two* nearest-neighbor White artists within a caliper of 0.1 on the propensity score. Figure 3.9 shows that, after matching, the standardized mean difference for every covariate falls below 0.05, confirming adequate balance and ensuring that the analysis is restricted to a region of common support.

Across all six model specifications as shown in Table 5.5, the propensity-score–matched (PSM) estimates are largely consistent with our main findings. Most notably, the baseline non-White gap in minting probability remains large and negative: the coefficient on non-White is –0.19 to –0.26 (p ¡ 0.01) in Columns (1)–(5) of the matched sample, mirroring the –0.20 to –0.32 range reported in the original regressions. Interaction terms that capture information provision continue to mitigate this disadvantage, and in several cases the magnitudes are even larger after matching. For example, the uplift from $Non-White \times Sig.Qual.Exhibition$ rises from 0.31–0.56 in the pooled sample to 0.50–0.86 in the matched sample, while $Female \times Sig.Qual.Exhibition$ turns statistically significant only after matching.

The PSM exercise does, however, reveal slightly weaker evidence of a direct gender gap. Whereas the pooled regressions show a consistently negative and often significant coefficient on $Female$ once full controls are included, the matched estimates hover near zero and lose significance. Overall, the core pattern of racial disadvantage and the moderating power of credible quality signals survives the like-for-like comparison, but the pure gender effect proves more sensitive to sample composition.

As shown in Table 5.6, the like-for-like analysis based on propensity-score matching confirms the core pricing patterns observed in the pooled regressions. In both panels, artworks by non-White artists command systematically lower primary prices, and this discount survives every specification. The magnitude remains economically meaningful (roughly –3% to –4% in the pooled sample; –5% to –9% after matching) and the coefficient is always significant once a full set of controls is included. Moreover, the information variables continue to offset this racial penalty: the interactions $Non-White \times ArtistWordCount$, $Non-White \times Sig.Qual.Exhibition$, $Non-White \times Sig.Qual.Client$, and $Non-White \times Sig.Qual.Location$ stay positive and highly significant, with slightly larger point estimates in the matched sample.

Gender effects are broadly consistent yet exhibit one noteworthy change. In the pooled regressions the main *Female* coefficient is small and never reaches significance, whereas the matched estimates reveal a clear gap of –8% to –10%. Thus, once female- and male-created NFTs are compared on observably similar listings, a gender discount emerges. Importantly, the moderating role of curatorial signals remains: $Female \times Sig.Qual.Exhibition$ and $Female \times Sig.Qual.Location$ attain significance in several matched specifications and have magnitudes comparable to the pooled results. Overall, the PSM exercise corroborates the racial pricing disparity and the protective value of verifiable quality signals, while uncovering a gender penalty that the pooled model understated.

3.8 Discussion

This study asked whether the promise of Web 3—to flatten hierarchies and widen access—actually translates into more equitable outcomes for artists who have long been marginalized in the traditional art world. Analyzing about 27,000 primary market sales in SuperRare, we find a nuanced picture: female and non-White creators participate at higher rates than they do at auction, yet White men still command the lion's share of market activity. Conditional on participation, non-White artists face a sizable penalty in both probability of sale and price, while the female discount is smaller and, in prices, often statistically insignificant. Crucially, self-curated quality signals, such as prior exhibitions, prominent clients, detailed bios, and being based in advanced economies, offset a meaningful share of these penalties.

Our evidence shows that decentralization alone does not entirely remove minority sales disadvantage; rather, it shifts the arena in which reputational signals are produced and consumed. In a ledger-based market where curation costs are pushed onto creators, verifiable markers of expertise become a scarce form of capital. Viewed through signaling theory, NFTs function less like frictionless commodities and more like credence goods whose value still hinges on social proof, only now the artist, not the gallerist, must supply it. Future work could explore platform-level interventions (e.g., algorithmic surfacing of under-signaled talent or third-party badge systems) and how these interact with downturns, when liquidity evaporates fastest.

3.9 Contribution and Implication

This study is the first to bring large-scale, transaction-level evidence from a blockchain art market to the literature on gender gaps in art valuation (Adams et al. (2021), Bocart et al. (2022)). By analyzing NFT trades—the first serious attempt to run a curatorial art market without human intermediaries—we show how digitization can shrink, but not erase, inequities for underrepresented creators. Methodologically, we marry econometric modeling with computer-vision embeddings, answering recent calls for richer measurement in art-economics research (Lee et al. 2024). Theoretically, we reframe the "higher-bar" debate: when gallery gatekeepers disappear, demand-side bias persists yet can be attenuated by self-curated signals.

For NFT platforms, the results surface an actionable lever: build low-cost interfaces that nudge artists to disclose credible résumé elements—structured fields for exhibitions, on-chain certificates of past clients, AI-assisted profile check-lists, and algorithmic surfacing of lightly signaled talent. Such design choices can enlarge inventory (e.g., artists and NFTs) diversity, deepen liquidity, and give collectors a disciplined alternative to heuristic cues like race or gender. In short, blockchain's promise to bridge the gap hinges on purchase decisions that make credibility signals salient at the point of purchase.

Minority artists can make the most of the system by treating profile curation as strategic branding. Detailed bios, verifiable links to prior shows, collaborations with reputable galleries or DAOs (Decentralized Autonomous Organizations), and on-chain proofs of critical recognition materially raise both sale odds and prices. Intermediaries, whether traditional galleries experimenting with NFTs or diversity, focused investment funds—can use these verifiable signals to allocate capital in ways that advance equity and showcase commitment to an inclusive Web 3 ecosystem.

4 From Dialogue to Decision: How Forum Discourse Quality Influences on Voting Participation in DAOs

4.1 Introduction

Decentralized Autonomous Organizations (DAOs) have evolved from small governance experiments into an organizational form that stewards a multi-billion-dollar market. Real-time analytics from DeepDAO record more than US$14 billion in liquid treasury assets under DAO control, a figure that places these communities alongside mid-cap public companies in financial scale[1]. Flagship protocols such as *Uniswap*, *Aave* and *Lido* rely on DAO votes to ratify software upgrades, set risk parameters and distribute incentive budgets, meaning that effective collective decision-making is now critical to the broader decentralized-finance ecosystem.

Unlike shareholder assemblies or board meetings, DAO governance is carried out entirely online. Any token holder may submit a proposal, and ballots are cast asynchronously, with voting power proportional to the tokens each wallet controls. Because participants are globally dispersed and typically pseudonymous, high-quality digital deliberation is essential: the community must inform itself, surface objections and build consensus without ever convening face-to-face.

Two platforms operationalize this workflow. Discussion unfolds on *Discourse*, an open source forum adopted by most large DAOs, while final ballots are collected on *Snapshot*, an off-chain voting service used by roughly 96 % of active DAOs according to industry coverage[2]. Discourse provides long-form debate, versioned edits and rich hyperlinking; Snapshot offers gas-less voting with flexible rules. Together the pair has become the canonical "deliberate-then-decide" stack for decentralized governance.

[1] Real-time treasury figures from `https://deepdao.io` (accessed 23 July 2025).
[2] See, e.g., `https://theblock.co/` industry report "DAO Tooling Landscape," February 2025.

Despite the centrality of this pipeline, previous academic work on DAO governance, such as Appel and Grennan (2023), has treated deliberation and turnout as separate subjects. In consequence, we still lack empirical evidence on how the texture of forum conversations shapes subsequent voting behavior. Addressing this gap, this study draws on the literature on social-media engagement to conduct an exploratory analysis associating discursive characteristics of community discourse to community engagement, i.e. voting-participation rates on Snapshot.

To this end, we compile the largest combined forum-and-voting dataset to date. Using the Boardroom API we harvest 248,256 replies from 15,244 Discourse posts covering 71 DAOs. We match these threads to 11,489 Snapshot proposals comprising 4.2 million individual ballots by 293,180 unique voter addresses. Natural-language processing yields metrics of semantic convergence, sentiment, argument role, cognitive presence and network centrality for each thread, which are then merged with two participation outcomes: the proportion of wallets that vote and the proportion of tokens cast.

The regressions reveal several robust patterns. First, discussions that converge semantically and revolve around reason, giving premises attract significantly higher turnout from both small and large holders, suggesting that clear, consensus-oriented dialogue lowers informational barriers to participation. Second, early-stage exploration rather than problem statements or solution declarations encourages wallets to vote, whereas highly centralized reply networks and fragmented topic mixes deter token-weighted engagement. Third, facilitative moderation, summaries and guiding prompts, has a modest positive effect, while directive instruction depresses participation among major token holders.

These findings imply that the quality and shape of forum discourse materially influence DAO democracy. Communities that design prompts and moderation guidelines to foster semantic alignment, deliberative reasoning and exploratory exchange can expect broader and deeper electoral engagement, thereby enhancing the legitimacy and resilience of decentralized governance.

4.2 Empirical Context: Discourse and Snapshot in DAO Governance

Decentralized Autonomous Organizations (DAOs) frequently utilize platforms such as Discourse and Snapshot to facilitate community governance and decision-making processes. Understanding the operational dynamics of these platforms is crucial for interpreting the empirical findings presented in this study.

Discourse functions as the primary forum where DAO members engage in extensive discussions and debates, enabling the exchange of ideas and the refinement of proposals before they proceed to formal voting. The platform resembles traditional online discussion forums, structured into threads and topics. Discussions commonly revolve around project funding, strategic partnerships, governance rules, protocol amendments, and community-driven initiatives. Members initiate discussions by creating threads, which subsequently invite community engagement through comments and interactive dialogue. These discussions are critical for achieving community consensus and shaping proposals to reflect diverse perspectives. Figure 4.1 presents an illustrative example of a typical Discourse discussion thread, highlighting the nature of active participation and diverse viewpoints among community members.

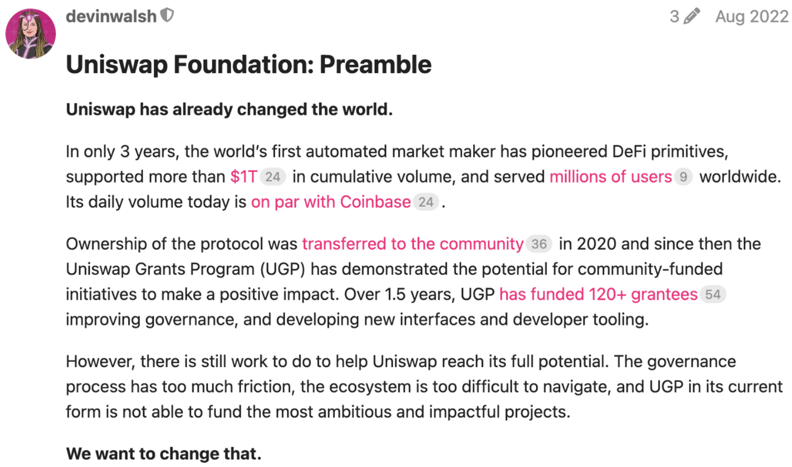

devinwalsh 3 Aug 2022

Uniswap Foundation: Preamble

Uniswap has already changed the world.

In only 3 years, the world's first automated market maker has pioneered DeFi primitives, supported more than $1T 24 in cumulative volume, and served millions of users 9 worldwide. Its daily volume today is on par with Coinbase 24 .

Ownership of the protocol was transferred to the community 36 in 2020 and since then the Uniswap Grants Program (UGP) has demonstrated the potential for community-funded initiatives to make a positive impact. Over 1.5 years, UGP has funded 120+ grantees 54 improving governance, and developing new interfaces and developer tooling.

However, there is still work to do to help Uniswap reach its full potential. The governance process has too much friction, the ecosystem is too difficult to navigate, and UGP in its current form is not able to fund the most ambitious and impactful projects.

We want to change that.

(a) Discourse Post

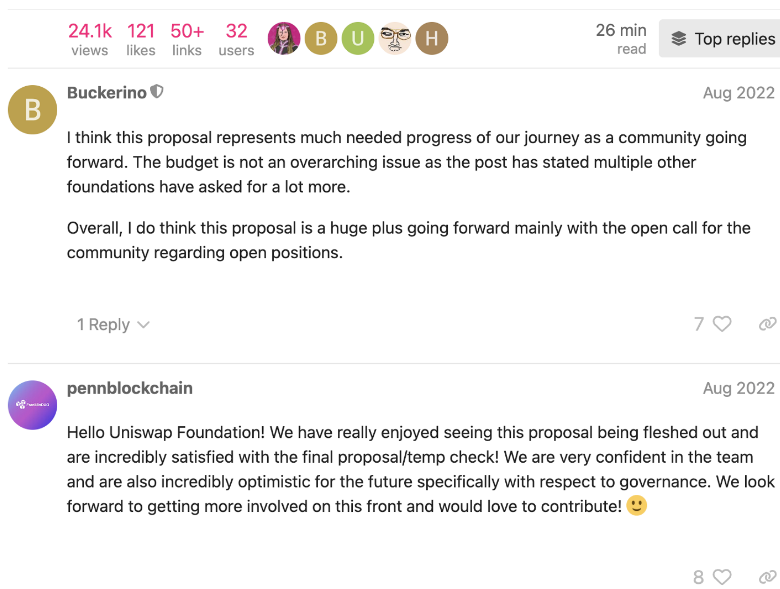

24.1k views 121 likes 50+ links 32 users 26 min read Top replies

Buckerino Aug 2022

I think this proposal represents much needed progress of our journey as a community going forward. The budget is not an overarching issue as the post has stated multiple other foundations have asked for a lot more.

Overall, I do think this proposal is a huge plus going forward mainly with the open call for the community regarding open positions.

1 Reply 7

pennblockchain Aug 2022

Hello Uniswap Foundation! We have really enjoyed seeing this proposal being fleshed out and are incredibly satisfied with the final proposal/temp check! We are very confident in the team and are also incredibly optimistic for the future specifically with respect to governance. We look forward to getting more involved on this front and would love to contribute! 🙂

8

(b) Replies Examples

Note: Source: https://gov.uniswap.org/t/temperature-check-create-the-uniswap-foundation/17358.

Figure 4.1: Screenshot of a Discourse Discussion Thread

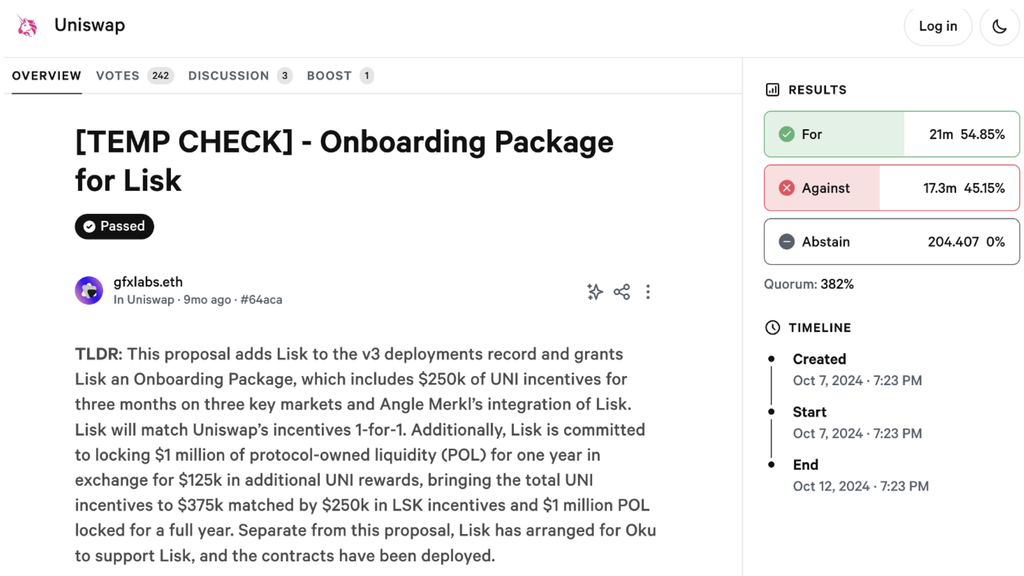

Note: Source: https://snapshot.box//s:uniswapgovernance.eth/proposal/0x64ac9

Figure 4.2: Screenshot of a Snapshot Voting Interface

Snapshot complements Discourse by serving as the designated platform for decentralized and gasless voting. Proposals that achieve maturity and consensus during Discourse discussions transition to Snapshot for formal voting by the community. The procedure involves the structured submission of proposals, detailing objectives, rationale, expected benefits, potential concerns, and clearly defined voting options. Each proposal specifies a fixed duration for voting, typically spanning from several days to one week. During this period, eligible DAO token holders cast votes proportionate to their token holdings by cryptographically signing messages through compatible wallets, such as MetaMask, without incurring transaction fees. The voting outcomes are automatically tallied upon conclusion of the voting window, adhering to predefined governance criteria such as quorum requirements and majority rules. Snapshot proposals typically address critical decisions involving treasury expenditures, protocol upgrades, community representative elections, and strategic organizational directions. Figure 4.2 demonstrates a typical Snapshot voting interface, showcasing an ongoing proposal alongside real-time voting results.

In summary, the combined functionalities of Discourse and Snapshot effectively support DAO governance by fostering transparent deliberative processes and structured decision-making. This dual-platform approach enhances the inclusivity, responsiveness, and efficacy of decentralized organizational governance.

4.3 Literature Review

Coordination in Open-source Communities

While we draw inspiration from open-source coordination theory, our context departs from it by embedding direct economic incentives and token-based governance that introduce strategic behaviors absent in volunteer software communities.

Open-source coordination theory examines how geographically dispersed, volunteer-based software communities self-organize and govern their development processes without a central authority. Early work by Crowston and Howison (2005) highlights how modular software architectures and shared norms allow contributors to coordinate asynchronously, while Linus's Law, "given enough eyeballs, all bugs are shallow," emphasizes the power of collective review and peer production (Raymond, 2001). Subsequent studies have identified key mechanisms such as meritocratic governance, modular task decomposition, transparent communication channels (e.g., mailing lists, issue trackers), and informal reputation systems that enable successful, large-scale collaboration (Weber and Hanssen, 2012, Ducheneaut et al., 2005).

The principles of open-source coordination provide a natural theoretical springboard for understanding DAO governance. Both contexts rely on distributed contributors coordinating without hierarchical oversight, leveraging digital platforms for discussion, decision-making, and artifact production (Reyna et al., 2018). However, DAOs extend these models by embedding economic incentives and on-chain voting mechanisms that directly tie governance outcomes to tokenized stakes. Building on the lessons from OSS communities about communication-based coordination (Crowston et al., 2006), this paper investigates how the quality and structure of off-chain discourse (in Discourse forums) affect on-chain participation (in Snapshot votes).

Coordination theory suggests that effective collaboration emerges from aligning participants' understandings of tasks, roles, and goals (Malone and Crowston, 1994). In OSS, shared repositories and code review workflows serve as coordination artifacts. In DAOs, Discourse threads, proposal templates, and Snapshot interfaces play analogous roles. We draw on this cross-domain insight to hypothesize that semantic, relational structure and emotional, informational content in off-chain dialogue will shape shared mental models, thereby influencing on-chain voting participation.

Recent Research on DAO Governance

Recent advances in DAO governance research have focused intensively on vote mechanics, economic incentives, and market impacts, ranging from novel bond voting schemes that lock tokens to deter Sybil attacks (Mohan et al., 2024) to large scale audits of protocol design failures (Feichtinger et al., 2023), and from formal economic models of token issuance and coalition dynamics (Laturnus, 2023, **?**) to empirical links between forum engagement and token price movements (Appel and Grennan, 2023) or the mitigating role of reputational forums in volatile markets (Han et al., 2023). However, these studies treat governance votes as isolated events and largely ignore the deliberative context—the off chain discourse in which arguments, evidence, and social cues shape participant understanding and, ultimately, voting behavior.

Mohan et al. (2024) introduce a novel “bond voting” mechanism that requires participants to lock tokens for a minimum time commitment to cast governance votes. They theorize that by tying voting rights to non-withdrawable token bonds, DAOs can mitigate Sybil attacks and reduce low-quality or frivolous participation. Through a combination of formal modelling and simulations calibrated to real-world token distributions, they demonstrate that bond voting can substantially increase the effective cost of attack while preserving proportional influence for long-term holders. However, their analysis abstracts away from the deliberative context, treating votes as isolated events rather than the culmination of structured discussion.

Feichtinger et al. (2023) offer one of the first large scale empirical critiques of on-chain governance in DAOs. Analyzing proposal lifecycles across multiple blockchains, they document systematic shortcomings: decision delays, ambiguous quorum rules, and post-vote reversals. Their study reveals that many protocols under-utilize governance tools or implement them inconsistently, leading to voter confusion and low engagement. While rigorous in its on-chain data gathering, this work does not examine the off-chain conversation that typically precedes these votes, leaving open the question of how forum discourse might alleviate or exacerbate the problems they identify.

Laturnus (2023) provides a broad economic framework for DAOs, situating them between classic firm theory and platform markets. He develops a model in which token issuance, voting rules, and contributor rewards interact to determine both participation equilibria and project survival thresholds. Laturnus’s calibrated simulations suggest that DAOs with low delegation costs and moderate token supply growth achieve the highest long-run viability. Although highly influential for DAO design, his paper treats deliberation efficiency as exogenous and does not integrate any measure of communication quality or community dialogue.

Sockin and Xiong (2023) analyze how tokenization itself underpins decentralization, showing that the availability of divisible, transferable voting rights transforms stakeholder incentives and coalition dynamics. Using theoretical bargaining models, they demonstrate that token flexibility can both empower fringe participants and create vote-splitting risks. Han et al. (2023) build on this by empirically examining governance outcomes across dozens of DAOs, finding that higher token volatility tends to depress participation and that reputation tracking via off-chain forums can mitigate these effects. Together, these studies demonstrate the importance of economic and reputational levers yet fall short of analyzing the content of the reputational signals, namely, the discourse in which arguments and social cues are exchanged.

Appel and Grennan (2023) link governance activity to digital asset prices, showing that elevated forum engagement often precedes positive price returns for governance tokens. Their event-study approach documents a roughly 2 percent abnormal return on proposal days with above-median discussion volume. While this finding suggests that discourse intensity matters for market outcomes, the authors focus on aggregate counts and do not disaggregate by the linguistic or structural qualities of the conversation itself.

Literature on Social Media Engagement

In this section, we review the literature on social media engagement to establish the theoretical foundation for our study. Research on social media engagement often distinguishes structural features from the content of messages in social network. Structure is the recognizable patterns formed by nodes and connections within the network, whereas content refers to resources accessible within a network, such as informational and emotional elements (Kane et al., 2014).

Researchers increasingly use network metrics (e.g., degree centrality, betweenness, cohesiveness) to quantify the **structural dimension** of social media interactions (Kane et al., 2014). Relational structure captures how users are positioned in relation to one another; for instance, outdegree reflects how actively a user engages with others, while centrality metrics identify influential nodes within the network (Wasserman and Faust, 1994, Ransbotham et al., 2012). In the context of online discourse, influential users who exhibit higher centrality in online communities, i.e., who frequently posts receiving many responses, play a key role in sustaining engagement and information dissemination (Zhang et al., 2007, Zhao et al., 2014).

Beyond these interactional structures, researchers have also examined semantic structure within conversations. This involves measuring how discussions evolve in terms of shared meaning and topical focus. For example, Kaplan (2008) demonstrates that in strategy-making under uncertainty, actors engage in framing contests to align divergent cognitive frames, with successful convergence shaping organizational decisions and power dynamics. Similarly, Miranda et al. (2022) analyze semantic convergence in blockchain discourse by examining how diverse frames introduced by different discursive fields gradually align into a coherent collective understanding. Such convergence reflects the extent to which participants build upon each other's contributions, forming a deliberative structure at the conceptual level. In this study, we incorporate both types of structure, relational (network-based) and semantic (text-based), to capture the layered architecture of community discourse.

The **content dimension** focuses on what is being communicated. A useful distinction is between informational content (factual, topic-driven messages) and emotional content (posts conveying sentiment or feelings). Informational content includes news, data, instructions, or knowledge sharing. Schanke et al. (2023) find that on corporate Facebook pages, posts emphasizing unexpected content types for the context. For instance, an emotional post on a typically "functional" product page garnered higher user engagement. Their findings reinforce that content type (informational vs. emotional) critically shapes user response, and that incongruity or mix of the two can attract attention. Cheng et al. (2024) investigate how firms' responses to initial user comments on social media with varying degree of comment sentiment, controversialness, response uniqueness, and timeliness, influence subsequent user engagement. They find that responding promptly and uniquely, especially to negative and controversial comments, increases user comment volume but can also amplify negativity in subsequent discussions.

By contrast, emotional content emphasizes feelings, opinions, or social support. Emotions on social media are pervasive and can powerfully influence reactions (Kane et al., 2014). Numerous studies have applied sentiment analysis in social network research. For instance, Lee and Lee (2019) demonstrates that the emotional tone and brand personality embedded in social media content play a key role in driving user engagement, with sentiment-based elements like humor and emotion consistently linked to positive reactions. Similarly, Schanke et al. (2023) highlight how sentiment, along with stylistic choices, influences user engagement outcomes, showing that emotionally incongruous and informally framed content tends to resonate more with users, particularly within the informal norms of platforms like Facebook. Overall, the content dimension encompasses message substance and tone; whether a post primarily delivers factual news or evokes emotion. Both aspects can matter: informational posts convey knowledge, while emotional posts connect to innate social needs, such as expressions of support and empathy. The structure and content aspect of discussions in Discourse is summarized in Table 4.1.

Structure		Content	
Relational	**Semantic**	**Emotional**	**Informational**
Participants'	Discussion Convergence	Affective Tone	Argumentative moves
Network	Topical Diversity	Social Tone	Cognitive moves
			Instructive moves

Table 4.1: Structural-Content Framework for Discourse Analysis

4.4 Hypotheses Development

Relational Structure and User Engagement

Users who occupy central positions in the social network, particularly those with high out-degree centrality, are more likely to influence broader community behavior by disseminating information and shaping attention flows. In deliberative governance contexts, high out-degree actors (i.e., users who frequently reply to or quote others) may function as informal information brokers, increasing the visibility and reach of a discussion thread. We theorize that threads involving users with higher centrality, especially maximum or average out-degree, are associated with greater voter turnout, as the information is more likely to diffuse through the network. Therefore, we propose:

H1. *Threads with higher network centrality, measured by maximum and mean out-degree of participants, will exhibit higher subsequent voting participation on the associated proposal.*

Semantic Structure and User Engagement

Semantic convergence refers to the extent to which participants in a discussion thread coalesce around shared language and framing over time. A steeper convergence slope and higher convergence similarity signal stronger alignment in terminology and conceptual focus. Prior research in collective intelligence and open-source collaboration suggests that such convergence facilitates mutual understanding, coordination, and follow-through. In DAO governance, semantic convergence likely signals a maturing consensus or shared interpretation of a proposal, lowering cognitive barriers for less-engaged members. Thus, we expect that threads with higher semantic convergence will correspond to higher voter turnout. Therefore:

H2. Threads that exhibit greater semantic convergence, measured by convergence slope and convergence similarity, will be associated with higher subsequent voting participation.

Emotional Contents and User Engagement

Emotional appeals can be a double-edged sword in governance discourse. On one hand, passion and moral urgency in forum posts may rally members by resonating with their values or fears, thereby mobilizing them to vote. Prior studies of political social networks show that moral-emotional language increases content sharing and can spur collective action. On the other hand, excessive affect or ”affective banter” might undermine deliberative credibility. Given that our context is deliberative and complex, participants might discount overly emotional posts as biased or noisy, particularly large token holders who are wary of hype. Preliminary evidence in DAO forums suggests that threads filled with casual or emotion-heavy chatter see lower turnout, indicating that serious voters may disengage when discourse becomes too frivolous or heated.

H3. Threads with higher emotional or social tone will exhibit lower voting participation than threads with a more neutral or rational tone. In other words, as the proportion of emotional content in a discussion increases, voter turnout on the associated proposal will decrease.

Informational Contents and User Engagement

We posit that the provision of rich information and reasoned arguments in forum discourse lowers barriers to participation and encourages voting. Informational content, such as factual evidence, analytical reasoning, and clear explanations, equips community members with the knowledge needed to make an informed decision. In marketing, informational (functional) appeals are found to be crucial for driving follow-through behaviors like purchases. Analogously, in DAOs, a thread high in informative quality should help members understand what is at stake and how a proposal aligns with their interests or the community's goals. Literature on collective intelligence and open-source projects emphasizes that shared understanding (semantic convergence) and premise-driven reasoning facilitate coordinated action. We expect the same for token-holder voting: when discourse focuses on facts, logic, and clarity, essentially treating members as decision-makers who need accurate information, more members (including large stakeholders) will be comfortable participating in the vote.

H4. Threads with a higher level of informational content (e.g., evidence, logical arguments, and clear factual details) will have higher voting participation than threads low in informational content. Greater presence of reasoned, informative discourse in a discussion is associated with an increase in voter turnout on its proposal.

4.5 Data and Variables

Dependent Variable: Official Proposal Voting Outcome Data

To quantify on-chain governance activity, we collected voting records from Snapshot (`https://snapshot.box`). Snapshot assigns each DAO a dedicated "space" (ticker); by manually cross-referencing the discourse list we located active spaces for 71 of the 76 DAOs in our sample. Using Snapshot's public GraphQL endpoint (`https://graphql.org/`) we downloaded complete proposal logs, including the proposal's title, free-text description, voting schema (single-choice, weighted, quadratic, etc.), opening and closing block times, and the proposer's wallet address. The final dataset contains 11 489 proposals, 4,267,818 ballots, and 293,180 unique voter addresses.

From these raw ledgers we derive two complementary governance-participation measures: *Participation (Users)*, the count of distinct wallets casting at least one vote, and *Participation (Tokens)*, the total voting power those wallets deploy. Formal definitions of all dependent and independent variables appear in Table 4.2.

Table 4.2: Variable Definitions

Category	Name	Definition
Dependent Variables		
—	*Participation (Users)*	Proportion of unique token-holders who voted on a proposal relative to the total number of token-holders.
—	*Participation (Tokens)*	Proportion of tokens cast in the proposal relative to the total circulating supply of tokens.
Independent Variables: Relational Structure		
Network centrality (Freeman, 2002)	*Max Out-degree*	Largest out-degree of any author in the thread (number of outgoing "replied-to-someone-else" edges).
	Mean Out-degree	Average out-degree across all participating authors; indicates overall propensity to address previous posters.
Independent Variables: Semantic Structure		
Semantic convergence (Chowdhury, 2010)	*Converge Slope*	Regression slope of consecutive-reply similarity; positive = convergence, negative = divergence.

Continued on next page

Category	**Name**	**Definition**
	Converge Similarity	Mean cosine similarity between each reply and its immediate predecessor.
Topical diversity (Shannon, 1948)	*Entropy Net*	Shannon entropy of category distribution within a thread.
Independent Variables: Emotional Content		
Affective sentiment (Hutto and Gilbert, 2014)	*Polarity*	VADER compound sentiment score for the thread's aggregated text.
Social presence (Rourke, 2001)	*Prob.Soc. Affective*	Mean probability of Social Presence – Affective Expression (emotion, self-disclosure, humour).
	Prob.Soc. Cohesion	Mean probability of Social – Group Cohesion (vocatives, inclusive pronouns, salutations).
Independent Variables: Informational Content		
Argument-role mix (Reimers and Gurevych, 2019)	*Prob. of Claim*	Mean soft-max probability a reply is a *Claim.*
	Prob. of Premise	Mean probability a reply supplies a *Premise* supporting/attacking a claim.
	Prob. of Evidence	Mean probability a reply contains factual *Evidence.*
	Prob. of Counterargument	Mean probability a reply is a *Counterargument.*
Cognitive presence (Garrison et al., 1999)	*Prob.Cog. Trigger*	Similarity-based probability a reply signals *Cognitive – Triggering.*
	Prob.Cog. Exploration	Probability replies are *Cognitive – Exploration.*
	Prob.Cog. Integration	Probability replies are *Cognitive – Integration.*
	Prob.Cog. Resolution	Probability replies are *Cognitive – Resolution.*

Continued on next page

Category	**Name**	**Definition**
Instructive presence (Anderson et al., 2001)	*Prob.Teach. Design*	Probability of Instructive *Design & Organization* moves.
	Prob.Teach. Facilitation	Probability of Instructive *Facilitation.*
	Prob.Teach. Direct	Probability of Instructive *Direct Instruction.*
Control Variables		
—	*Topic Prob.*	Posterior probability that the thread's aggregated reply text belongs to latent topic j as estimated by BERTopic. (Grootendorst, 2022)
—	*Number Reads*	Total number of reads for each post (Discourse).
—	*Network Size*	Total number of token-holders for the DAO on day t (IntoTheBlock).

Table 4.3: Summary Statistics

	N	**Mean**	**S.D.**	**Median**	**Min**	**Max**
Dependent Variables						
Participation (Users)	24 436	0.07	0.16	0.05	0	11.50
Participation (Tokens)	24 447	0.02	0.82	0.01	0	127.66
Independent Variables						
Max Out–degree	24 447	2.01	2.18	1	0	33
Mean Out–degree	24 447	0.91	0.57	1	0	3.62
Converge Slope	16 969	0.00	0.06	0.00	−0.30	0.48
Converge Similarity	20 943	0.38	0.16	0.37	−0.09	1
Polarity	24 447	0.94	0.24	1	−1	1
Entropy Net	24 447	1.37	0.94	1.39	0	4.46
Prob. of Claim	24 447	0.26	0.01	0.26	0.23	0.29
Prob. of Premise	24 447	0.25	0.01	0.25	0.22	0.27
Prob. of Evidence	24 447	0.25	0.00	0.25	0.24	0.29
Prob. of Counterargument	24 447	0.24	0.01	0.24	0.22	0.29

						Continued on next page
Prob. Cog. Trigger	24 447	0.05	0.04	0.05	−0.14	0.41
Prob. Cog. Exploration	24 447	0.13	0.05	0.13	0.09	0.35
Prob. Cog. Integration	24 447	0.07	0.03	0.07	−0.06	0.23
Prob. Cog. Resolution	24 447	0.05	0.04	0.05	−0.14	0.40
Prob. Soc. Affective	24 447	0.09	0.04	0.10	−0.09	0.34
Prob. Soc. Cohesion	24 447	0.12	0.04	0.12	0.06	0.32
Prob. Teach. Design	24 447	0.19	0.07	0.19	−0.09	0.44
Prob. Teach. Facilitation	24 447	0.19	0.06	0.19	−0.07	0.45
Prob. Teach. Direct	24 447	0.09	0.04	0.09	−0.13	0.30
Controls						
Topic Prob.			*Omitted for brevity*			
Number Reads	24 447	421	1 185	77	0	31K
Network Size	24 447	423K	458K	133K	47	1280K

Note. Values are rounded: K = thousands.

Independent Variables: Community Discourse Data

To investigate how community discourse shapes participation in DAO governance, we compiled a comprehensive corpus from Discourse[3], the forum platform most broadly adopted across the DAO ecosystem. Discussion threads were obtained via the Boardroom API[4], a data source recently employed by Appel and Grennan (2023). The API currently indexes 76 DAOs, ranging from flagship protocols such as Uniswap, Aave to mid-tier projects like Ampleforth and Frax Finance (see Appendix for the complete roster).

For each thread, we extracted post-level metadata, title, view count, like count, author identifier, and timestamp, yielding 15,244 distinct posts. We then captured every associated reply, recording its text, author identifier, parent-post identifier, view count, like count, and timestamp, for a total of 248,256 replies.

All reply texts were pre-processed uniformly: HTML tags were removed; characters

[3] `https://www.discourse.org/`

[4] `https://docs.boardroom.io/docs/api`

were Unicode-normalized (NFKD) and lower-cased; URLs and non-alphabetic tokens were stripped; English stop words were discarded; and the remaining tokens were lemmatized with WordNet. The resulting lemmas were concatenated to form the analytical corpus used in subsequent machine-learning and statistical modeling.

4.6 Empirical Model

To quantify the extent to which linguistic and interactional properties of forum discourse influence on subsequent near-term governance participation we estimate a fixed-effects model. Standard errors are clustered at the DAO level to account for within-DAO serial correlation.

$$Y_{i,t} \;=\; \beta X_{i,t} \;+\; \gamma C_{i,t} \;+\; \eta_{d(i)} \;+\; \delta_{m(t)} \;+\; \varepsilon_{i,t} \tag{4.1}$$

where $Y_{i,t}$ is one of two turnout measures for thread i in month t: (a) governance-token volume cast *(Participation (Tokens))* or (b) unique voting wallets *(Participation (Users))*. $X_{i,t}$ is a single focal discourse feature (sentiment, centrality, cognitive-presence probability, etc.). $C_{i,t}$ contains baseline controls—the cumulative number of forum reads *(Number Reads)* and, for the token model, the total addresses with a positive governance-token balance *(Network Size)*—together with the full vector of `BERTopic` posterior probabilities (*topic_1_prob* ... *topic_19_prob*). $\eta_{d(i)}$ are DAO fixed effects that absorb all time-invariant heterogeneity across DAOs. Year-month time fixed effects (δ_m) remove common seasonality and network-wide shocks.

Participation outcomes are restricted to proposals opened within seven days of the thread's creation. Let t_i^{disc} be the timestamp of thread i and t_i^{vote} the opening time of proposal p. A proposal is included if

$$0 \;\le\; \Delta_{i,p} \;=\; t_i^{\text{vote}} - t_i^{\text{disc}} \;\le\; 7 \text{ days}. \tag{4.2}$$

4.7 Results

4.7.1 Exploratory Analysis of Discourse Features

In this section, we report and interpret the results of eight sets of ordinary least squares regressions, each isolating a single discourse feature as the key explanatory variable. The two dependent variables are the logarithm of the fraction of unique token-holders who cast a vote ("Holders") and the logarithm of the share of tokens cast ("Tokens"). All models include controls for total read counts, total token balance, the full set of latent topic-probability scores, DAO fixed effects, and year-month fixed effects (these controls are suppressed in the tables for clarity). Because we use log-transformed participation rates, a coefficient of 0.01 corresponds approximately to a 1 percent increase in turnout.

Table 4.4: Relational Structure as Network Centrality

	Model 1		Model 2	
	Holders	Tokens	Holders	Tokens
Max Out – degree	0.000	-0.009^{***}		
	(0.000)	(0.001)		
Mean Out – degree			-0.001	-0.009^{***}
			(0.001)	(0.001)
Number Reads	0.000^{**}	0.001^{***}	0.000^{***}	-0.001^{*}
	(0.000)	(0.000)	(0.000)	(0.000)
Network Size		-0.008^{***}		-0.007^{***}
		(0.000)		(0.000)
Post Topic Controls	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES
Num. Obs.	24 439	24 428	24 439	24 428
R^2	0.301	0.325	0.301	0.323

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

The first set of models (Table 4.4) examines network centrality in the discussion thread, operationalized by each user's out-degree. We find that neither the maximum out-degree nor the mean out-degree has a statistically significant association with the number of distinct voters. However, both measures bear a negative coefficient of approximately -0.009 in the token-share regressions, significant at the 1 percent level. This suggests that when a discussion is dominated by a highly active few, either a single super-replier or a tight cluster of prolific respondents, token-weighted engagement declines by nearly 1 percent for each unit increase in concentration. Such centralization may signal to token-holders that the conversation is captured by insiders, reducing confidence that their own perspectives will be heard and thereby diminishing their willingness to commit tokens.

Table 4.5: Semantic Structure as Convergence

	Model 1: Converge Slope		Model 2: Converge Similarity	
	Holders	Tokens	Holders	Tokens
Converge Slope	0.006	0.016^{**}		
	(0.004)	(0.008)		
Converge Similarity			0.003^{*}	0.014^{***}
			(0.001)	(0.003)
Number Reads	0.001^{***}	-0.001^{***}	0.000^{***}	-0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Network Size	-0.009^{***}	-0.009^{***}	-0.009^{***}	-0.008^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Post Topic Controls	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES
Num. Obs.	16 969	16 959	20 943	20 932
R^2	0.319	0.344	0.305	0.330

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Semantic convergence, how much replies grow closer in meaning over time, appears to mobilize participants (Table 4.5). The convergence *slope* carries a positive coefficient of 0.016 in the *Tokens* regression ($p < 0.05$) and a smaller, positive but non-significant effect for *Holders*. Meanwhile, the average pairwise similarity between adjacent replies predicts both more voters (+0.003, $p < 0.10$) and more tokens (+0.014, $p < 0.01$). These results imply that as participants increasingly echo one another's language and arguments, uncertainty around the proposal's implications diminishes and both light-stake and heavy-stake holders gain confidence to vote. In practical terms, guiding a thread toward conceptual alignment, through summarizing posts or reflective prompts, may lower the cognitive barriers to participation.

By contrast, the breadth of topics discussed, measured by Shannon entropy, has little effect on voter head-count but a negative effect on token turnout of roughly -0.009 ($p < 0.01$) per 0.1 increase in entropy (Table 4.6). This pattern suggests that highly diversified threads, which span many semantic categories, can scatter attention and leave token-rich members uncertain which line of reasoning to support. While the sheer number of voters does not change, their propensity to allocate significant token weight falls when the conversation lacks a clear thematic focus.

Table 4.6: Semantic Structure as Topical Diversity

	Holders	Tokens
Entropy Net	0.000	-0.009^{***}
	(0.001)	(0.001)
Number Reads	0.000	0.000
	(0.000)	(0.000)
Network Size		-0.007^{***}
		(0.000)
Post Topic Controls	YES	YES
DAO FE	YES	YES
Year–Month FE	YES	YES
Num. Obs.	24 439	24 428
R^2	0.301	0.323

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

We next consider the role of affective sentiment (Table 4.7). The VADER polarity score, capturing net positive versus negative tone, yields small coefficients near zero in both regressions and fails to reach statistical significance. Thus, an upbeat or downbeat emotional climate alone does not appear sufficient to motivate or deter voting. This finding aligns with theories of deliberative engagement that emphasize substance over style: participants respond to the logic and structure of discussion rather than its valence.

Table 4.7: Emotional Content as Affective Sentiment

	Holders	Tokens
Polarity	0.001	0.000
	(0.001)	(0.002)
Number Reads	0.000^{***}	-0.002^{***}
	(0.000)	(0.000)
Network Size		-0.007^{***}
		(0.000)
Post Topic Controls	YES	YES
DAO FE	YES	YES
Year–Month FE	YES	YES
Num. Obs.	24 439	24 428
R^2	0.301	0.322

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.8: Emotional Content as Social Presence

	Model 1: Soc. Affective		Model 2: Soc. Cohesion	
	Holders	Tokens	Holders	Tokens
Prob.Soc. Affective	-0.009^{**}	-0.052^{***}		
	(0.004)	(0.008)		
Prob.Soc. Cohesion			0.006	0.008
			(0.004)	(0.008)
Number Reads	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Network Size		-0.007^{***}		-0.007^{***}
		(0.000)		(0.000)
Post Topic Controls	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES
Num. Obs.	24 439	24 428	24 439	24 428
R^2	0.301	0.323	0.301	0.322

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.9: Informational Content as Arguments

	Claim		**Premise**		**Evidence**		**Counterarg.**	
	Holders	Tokens	Holders	Tokens	Holders	Tokens	Holders	Tokens
Prob. of Claim	-0.169^{***}	-0.316^{***}						
	(0.026)	(0.053)						
Prob. of Premise			0.195^{***}	0.474^{***}				
			(0.026)	(0.054)				
Prob. of Evidence					-0.214^{***}	-0.235^{***}		
					(0.032)	(0.066)		
Prob. of Counterargument							0.132^{***}	-0.003
							(0.028)	(0.058)
Number Reads	0.000^{**}	-0.002^{***}	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}	0.000^{**}	-0.002^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Network Size		-0.007^{***}		-0.007^{***}		-0.007^{***}		-0.007^{***}
		(0.000)		(0.000)		(0.000)		(0.000)
Post Topic Controls	YES	YES	YES	YES	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES	YES	YES
Num. Obs.	24 439	24 428	24 439	24 428	24 439	24 428	24 439	24 428
R^2	0.302	0.323	0.303	0.324	0.302	0.322	0.302	0.322

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Post topic controls, DAO FE, and Year–Month FE included where indicated.

Social-presence cues (Table 4.8) yield mixed effects: affective expressions such as humor and self-disclosure reduce both voters (by about 0.9 percent, $p < 0.05$) and token share (by about 5.2 percent, $p < 0.01$), while cohesion signals (inclusive pronouns, salutations) have no discernible effect. Although building rapport can strengthen social bonds, excessive light-heartedness may distract from decision tasks, particularly for token-heavy holders focused on financial stakes.

Argumentative structure proves highly consequential (Table 4.9). Threads rich in premises, statements offering reasons, see approximately 19.5 percent more voters and 47.4 percent more tokens (both $p < 0.01$), whereas threads dominated by claims (assertions) experience declines of 16.9 percent and 31.6 percent, respectively. Evidence-focused posts likewise reduce turnout by roughly 21–24 percent. Counter-argument moves boost the number of voters by about 13 percent ($p < 0.01$) but have no significant effect on token share. Taken together, these results indicate that deliberative reasoning and dialectical exchange, eliciting *why* and *how* rather than *what*—are the most powerful levers for mobilizing both participation margins.

Cognitive-presence phases (Table 4.10) reveal that exploration posts, those sharing and building on ideas, raise the number of voters by 1.6 percent ($p < 0.01$) but do not significantly affect token proportion. In contrast, triggering posts (problem statements) and resolution posts (solution declarations) are associated with reductions in both outcomes, on the order of 2–5 percent (all $p < 0.01$). Integration moves (weaving ideas together) also lower voter count by roughly 4.5 percent. These findings suggest that keeping a thread in an active inquiry mode, rather than opening or closing it too abruptly, sustains engagement, perhaps by maintaining an open invitation to contribute.

Table 4.10: Informational Content as Cognitive Presence

	Trigger		**Exploration**		**Integration**		**Resolution**	
	Holders	Tokens	Holders	Tokens	Holders	Tokens	Holders	Tokens
Prob. Cog. Trigger	-0.023^{***}	-0.043^{***}						
	(0.004)	(0.008)						
Prob. Cog. Exploration			0.016^{***}	0.003				
			(0.003)	(0.007)				
Prob. Cog. Integration					-0.045^{***}	0.017		
					(0.005)	(0.010)		
Prob. Cog. Resolution							-0.025^{***}	-0.056^{***}
							(0.004)	(0.009)
Number Reads	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Network Size		-0.007^{***}		-0.007^{***}		-0.007^{***}		-0.007^{***}
		(0.000)		(0.000)		(0.000)		(0.000)
Post Topic Controls	YES	YES	YES	YES	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES	YES	YES
Num. Obs.	24 439	24 428	24 439	24 428	24 439	24 428	24 439	24 428
R^2	0.302	0.322	0.302	0.322	0.303	0.322	0.302	0.323

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Post topic controls, DAO FE, and Year–Month FE included where indicated.

Table 4.11: Informational Content as Instructive Presence

	Design		Facilitation		Direct	
	Holders	Tokens	Holders	Tokens	Holders	Tokens
Prob.Teach. Design	-0.002	0.001				
	(0.003)	(0.005)				
Prob.Teach. Facilitation			0.006^{**}	0.010^{*}		
			(0.003)	(0.006)		
Prob.Teach. Direct					-0.003	-0.021^{***}
					(0.004)	(0.008)
Number Reads	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}	0.000^{***}	-0.002^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Network Size		-0.007^{***}		-0.007^{***}		-0.007^{***}
		(0.000)		(0.000)		(0.000)
Post Topic Controls	YES	YES	YES	YES	YES	YES
DAO FE	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	24 439	24 428	24 439	24 428	24 439	24 428
R^2	0.301	0.322	0.301	0.322	0.301	0.322

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Finally, instructive-presence moves (Table 4.11) highlight that facilitation, guiding prompts, summaries, and clarifying questions, yields modest but significant increases in both voter count (0.6 percent, $p < 0.05$) and token turnout (1.0 percent, $p < 0.10$). Direct instruction, however, has no significant impact on the number of voters and reduces tokens by approximately 2.1 percent ($p < 0.01$). Design-and-organization posts show no effect. These patterns imply that scaffolding discussion through open-ended guidance is more effective than didactic, top-down directives in encouraging both broad and deep participation.

Overall, these exploratory analysis has shown how conversations unfold on governance forums powerfully shapes off-chain voting. Threads that gradually converge on shared language and argument structures mobilize both light- and heavy-stake participants. Breadth without focus scatters attention, positive or negative mood by itself does little, and social-presence cues can backfire when they distract from decision making. In contrast, reason-giving, dialectical counter-moves, open-ended exploration, and facilitative prompts consistently raise participation. Taken together, the evidence suggests that clarity, deliberative depth, and inclusive moderation, not mere volume or sentiment, are the cornerstones of effective DAO discourse.

Building on these insights, the following section presents a multivariate model that integrates the full set of discourse attributes. Specifically, we include variables identified as more relevant in the previous analyses. These variables are selected to jointly represent four key conceptual constructs, capturing both the structure and content dimensions of discourse, as summarized in Table 4.1. This specification will allow us to assess the relative associations of each mechanism while accounting for their joint effects.

4.7.2 Final Model with Selected Features

To conceptually represent the four key constructs in a single model, we include the following variables: a network centrality measure for Relational Structure, semantic convergence and topical diversity measures for Semantic Structure, polarity for Emotional Content, and finally, argumentation and cognitive presence for Informational Content, alongside control variables. While Argumentation and Cognitive Presence are each composed of multiple sub-measures (please see 4.2), such as claim and premise for Argumentation, and exploration and resolution for Cognitive Presence, we include only one representative measure from each category in each regression model. For instance, Table 4.12 presents results based on the use of claim and exploration as representative indicators.

Table 4.12: Multivariate Associations of Discourse Attributes with Voter Turnout

	Holders	Tokens
Network	0.016^{***}	0.058^{***}
	(0.001)	(0.001)
Converge Similarity	0.004^{**}	0.010^{***}
	(0.002)	(0.008)
Topical Diversity	0.008^{*}	0.011^{***}
	(0.004)	(0.002)
Polarity (Emotional Valence)	0.007^{*}	0.011^{**}
	(0.004)	(0.005)
Argumentation	0.193^{***}	0.158^{**}
	(0.009)	(0.065)
Cognitive Presence	0.010^{**}	0.029^{***}
	(0.004)	(0.001)
Number Reads	0.000	0.000
	(0.000)	(0.000)
Network Size		-0.009^{***}
		(0.000)
Post Topic Controls	YES	YES
DAO FE	YES	YES
Year–Month FE	YES	YES
Num. Obs.	16,969	16,959
R^2	0.328	0.351

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.12 presents the results of the multivariate analysis incorporating one representative variable from each of the four conceptual constructs. Across both specifications (voter turnout by holders and by tokens), we observe that all four dimensions are significantly associated with turnout, albeit with varying effect sizes. In particular, Relational Structure (e.g., network centrality) shows strong positive associations with turnout in both models, suggesting that posts written by more centrally positioned contributors in the discussion network are more likely to generate engagement. Both Semantic Structure variables, converge similarity and topical diversity, also show significant positive associations, indicating that voters are more responsive to discourse that is both thematically cohesive and topically rich. For Emotional Content, higher polarity (more emotionally positive valence) is significantly associated with greater participation, though the effect size is modest. Finally, within Informational Content, both argumentation and cognitive presence contribute positively to turnout, with argumentation showing particularly strong associations, especially in the holders model.

The robustness of these associations across two different turnout measures highlights the multifaceted nature of persuasive or mobilizing discourse. Notably, even when including various controls and fixed effects, the key constructs remain statistically significant, underscoring their explanatory relevance.

Please note that we conduct additional analyses using different combinations of these sub-measures and find that the results remain consistent. Although some individual measures, such as counterarguments, exhibit opposite signs (e.g., a negative coefficient compared to the positive coefficient for claim), the coefficients for other constructs, such as Relational Structure and Emotional Content, remain positively significant across specifications, with varying magnitudes. This consistency supports the robustness of our findings. While the contrasting signs among certain sub-measures (e.g., claim vs. counterargument) could offer valuable insights, we do not elaborate on these results here due to space constraints. Similarly, for constructs represented by more than one variable, such as Network Structure, measured by both maximum out-degree and mean out-degree, we run separate models for each measure and observe consistent patterns.

4.8 Conclusion

In this paper, we set out to explore which elements of community discourse on Discourse forums are systematically associated with voting participation on Snapshot. To do so, we built and analyzed the largest dataset of its kind, linking 248,256 pre-vote replies from 15,244 forum posts across seventy-one DAOs to 11,489 Snapshot proposals and over

4.2 million ballots. We extracted a rich array of thread-level metrics, semantic convergence, sentiment polarity, argument-role probabilities, cognitive-presence phases, social- and instructive-presence cues, and reply-network centrality.

Our results reveal a consistent pattern: discourse that converges in language, centers on premise-driven reasoning, sustains exploratory idea exchange, and employs facilitative moderation is associated with significantly higher turnout from both small and large token holders. In contrast, threads characterized by centralized reply to structures, highly fragmented topical coverage, affective-heavy social banter, or directive instruction correspond to lower participation, particularly in terms of token-weighted share. Sentiment polarity and design-and-organization cues showed little systematic effect, underscoring that it is the substance and structure of argumentation and interaction, not mere tone or volume, that mobilize DAO electorates.

By linking qualitative features of online deliberation to democratic engagement in tokenized governance, our study extends coordination theory from open-source software communities into the emerging landscape of DAOs. Practically, our findings offer actionable guidance for community managers and platform designers: craft prompts and moderation policies that foster consensus building, reason giving, and open exploration rather than authoritative directives or unfocused debate. Future work might leverage experimental interventions or causal inference methods to validate these associations, but the present exploratory analysis lays a foundational empirical roadmap for enhancing transparency, inclusivity, and resilience in decentralized decision-making.

4.9 Contribution and Implication

Our findings contribute to three core theoretical streams. First, they extend the private–collective innovation model of open-source communities (von Hippel and von Krogh,

2003) by demonstrating that discourse practices, specifically semantic convergence and premise-driven reasoning, are not only crucial for code coordination but also for mobilizing token-based participation in DAOs. Whereas prior work has emphasized structural factors such as contribution strata (Crowston and Howison, 2005) and economic incentives (Laturnus, 2023), our results show that the quality and shape of online dialogue materially influence democratic engagement. In doing so, we bridge open-source coordination theory and tokenized governance, highlighting that shared language and reasoned debate function as low-cost signals of group alignment that reduce informational frictions and encourage turnout.

Second, by applying the *Community of Inquiry* framework (Garrison et al., 1999) to a governance context, we demonstrate that cognitive-presence dimensions, i.e., exploration, integration, resolution, are highly associated with voter turnout. Our evidence that exploration posts boost voter headcounts while triggering and resolution moves suppress both wallet- and token-based turnout suggests a non-linear relationship between inquiry phases and participation. This insight reveals that sustaining an open deliberation phase is essential not only for learning outcomes but also for collective decision-making in decentralized organizations, thereby broadening the framework's relevance to digital public spheres.

Third, we contribute to argumentation theory by quantifying the level of argumentation occurring in a discussion relate to persuasion and action in a high-stakes setting (Toulmin 1958). The strong positive effects of premises and counter-arguments on turnout, and the negative effects of assertive claims, underscore that how an argument is structured matters more than its mere presence. These findings align with classical rhetoric's emphasis on logos over mere statement and provide empirical support for the notion that reason-giving moves foster participation more effectively than assertions or data-heavy appeals.

Together, our study offers a theoretical synthesis: effective DAO governance depends on discourse dynamics that foster shared understanding, scaffold inquiry, and structure arguments in ways that signal collective voter turnout. By integrating open-source coordination, CoI, and social media engagement literature into the context of tokenized democracy, we lay the groundwork for future research on the causal mechanisms through which digital deliberation drives electoral engagement in decentralized systems.

4.10 Limitation and Future Study

While our study leverages an unprecedentedly large corpus of forum posts and voting records, several limitations warrant caution. First, our estimation could be inherently correlational: by estimating each discourse feature one at a time in fixed-effects regressions, we cannot definitively establish causal direction. Unobserved confounders, such as off-thread coordination via private chats or external governance signals, may jointly influence both discussion patterns and voting turnout. Second, our natural-language measures, though grounded in established frameworks (Toulmin, 1958, Garrison et al., 1999), rely on supervised classifiers and semantic-similarity scores that may misclassify nuanced or emerging rhetorical styles. Measurement error in these text features could attenuate true effects or create spurious associations.

Future research should pursue causal identification strategies to validate and refine our exploratory findings. Randomized field experiments, such as A/B tests of discussion-prompt wording or moderation styles, would help isolate mechanisms and rule out reverse causality. Structural models of discourse evolution and turnout could illuminate dynamic feedback loops between conversation phases and voting deadlines. Additionally, qualitative studies or surveys could probe token-holders' perceptions of discourse quality, complementing our machine-learning metrics with first-hand accounts of deliberative experience.

Finally, extending the analysis across a broader range of DAO contexts and platforms would test the robustness of our conclusions. Comparative work could investigate whether cultural or linguistic differences alter which discourse features matter most, and whether emerging off-chain governance tools reshape the relationship between deliberation and democratic engagement. By addressing these limitations, future studies can deepen understanding of how digital dialogue drives participation in decentralized systems.

5 Concluding Remarks

This dissertation has explored how blockchain and Web3 technologies are reshaping governance and market interactions by removing traditional intermediaries. By enabling direct, peer-to-peer interactions, these digital innovations alter the underlying structures of both organizational governance and market practices, offering unprecedented opportunities for inclusion and collective empowerment. Yet, the removal of intermediaries does not automatically guarantee equitable or efficient outcomes; rather, it introduces new challenges, requiring careful consideration of how decentralized communities and markets function.

The first essay investigated the resilience of decentralized autonomous organizations (DAOs) in the face of adverse shocks, emphasizing the role of token ownership concentration. It demonstrated that DAOs with concentrated token ownership among committed blockholders exhibit greater resilience following cybersecurity attacks. Empirical analyses revealed a significant positive effect of ownership concentration on post-shock recovery, highlighting an important synergy between decentralization and effective stakeholder engagement. This finding proposes a nuanced insight that while decentralized structures aim for broad participation, strategic consolidation of blockholders' influence can be beneficial during crises.

The second essay examined decentralized NFT art marketplaces, focusing on how removing intermediaries affects historically underrepresented artists in a curated art market. Despite the lowered entry barrier, minority discounts for female and non-White artists exist, albeit to a lesser extent. It also suggests a potential pathway to mitigate the gap: strategic self-curation, such as providing credible signals of artistic quality through verifiable credentials and cultural advantages, significantly reduced the minority discount. These findings suggest that eliminating traditional gatekeepers may expand market access, but also place greater responsibility on artists to strategically represent themselves, showing the continued importance of curatorial strategies in decentralized settings.

The third essay analyzed how the structure and content of DAO governance discussions shape democratic participation. By employing a structural-content framework built on the literature on social media engagement, the study found that discourse characterized by semantic alignment, reasoned argumentation, and having influential participants significantly boosted voter turnout. These insights bridge existing research on open-source communities and social media engagement with the emerging literature on DAO governance, suggesting that understanding decentralized governance demands close attention to the structure and content of community discourse.

On the whole, the essays in this dissertation contribute a multi-dimensional understanding of disintermediation enabled by blockchain technology. Disintermediation brings new challenges that go beyond eliminating intermediaries, such as evolving creator roles and the need for robust community discourse. While blockchain-based transformations offer significant innovations in organizational design and market access, sustainable growth will depend on well-crafted governance mechanisms, a redefinition of artistic agency, and intentional strategies for fostering inclusive and meaningful participation.

In closing, the insights gained through these studies have broader implications for both practitioners and researchers exploring digital transformation. For DAO developers, NFT artists, and broader community members, recognizing the nuanced dynamics uncovered in this research can inform practical interventions and policy designs that support the sustainable growth of decentralized ecosystems. For researchers, the findings open new avenues for investigating the social and economic implications of decentralization in the area of governance and digital marketplace. Ultimately, this dissertation offers insights into how to leverage the opportunities of blockchain-driven disintermediation while proactively addressing the challenges that arise from organizational and market transformations.

Bibliography

Adams, R. B., Kräussl, R., Navone, M., and Verwijmeren, P. (2021). Gendered Prices. *The Review of Financial Studies*, 34(8):3789–3839.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and innovation: An inverted-u relationship. *The Quarterly Journal of Economics*, 120(2):701–728.

Agrawal, A., Catalini, C., and Goldfarb, A. (2014). Some simple economics of crowdfunding. *Innovation Policy and the Economy*, 14(1):63–97.

Agrawal, A., Lacetera, N., and Lyons, E. (2016). Does standardized information in online markets disproportionately benefit job applicants from less developed countries? *Journal of International Economics*, 103:1–12.

Albuquerque, R., Koskinen, Y., Yang, S., and Zhang, C. (2020). Resiliency of environmental and social stocks: An analysis of the exogenous covid-19 market crash. *The Review of Corporate Finance Studies*, 9(3):593–621.

Alexander, V. D. (1996). Pictures at an exhibition: Conflicting pressures in museums and the display of art. *American Journal of Sociology*, 101(4):797–839.

Anderson, T., Liam, R., Garrison, D. R., and Archer, W. (2001). Assessing instructive presence in a computer conferencing context. *Journal name*.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

Appel, I. and Grennan, J. (2023). Decentralized governance and digital asset prices. *Working Paper*.

Athey, S. and Imbens, G. W. (2022). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*, 226(1):62–79.

Aubry, M., Kräussl, R., Manso, G., and Spaenjers, C. (2023). Biased auctioneers. *The Journal of Finance*, 78(2):795–833.

Autor, D. H. (2003). Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics*, 21(1):1–42.

Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395.

Baron, R. M. and Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6):1173.

Beck, R., Müller-Bloch, C., and King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10):1020–1034.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275.

Bocart, F. Y. R. P., Gertsberg, M., and Pownall, R. A. J. (2022). An empirical analysis of price differences for male and female artists in the global art market. *Journal of Cultural Economics*, 46(3):543–565.

Bohren, J., Imas, A., and Rosenberg, M. (2019). The dynamics of discrimination: Theory and evidence. *American Economic Review*, 109(10):3395–3436.

Brynjolfsson, E., Hu, Y., and Rahman, M. S. (2009). Battle of the retail channels: How product selection and geography drive cross-channel competition. *Management Science*, 55(11):1755–1765.

Burtch, G., Carnahan, S., and Greenwood, B. N. (2018). Can you gig it? an empirical examination of the gig economy and entrepreneurial activity. *Management Science*, 64(12):5497–5520.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cameron, L., Goetzmann, W. N., and Nozari, M. (2019). Art and gender: market bias or selection bias? *Journal of Cultural Economics*, 43(2):279–307.

Catalini, C. and Boslego, J. (2019). Blockchain technology and organization science: Decentralization theatre or novel organizational form? *MIT Working Paper*.

Cavusoglu, H., Mishra, B., and Raghunathan, S. (2004). The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. *International Journal of Electronic Commerce*, 9(1):70–104.

Cheng, X., Bala, H., and Yang, M. (2024). Engaging users on social media business pages: The roles of user comments and firm responses. *MIS Quarterly*, 48(2).

Chowdhury, G. (2010). *Introduction to Modern Information Retrieval*. Facet Publishing, London, UK.

Cong, L. W., Harvey, C. R., Rabetti, D., and Wu, Z.-Y. (2024). An anatomy of crypto-enabled cybercrimes. *Management Science*.

Crowston, K. and Howison, J. (2005). The social structure of free and open-source software development. *First Monday*, 10(2).

Crowston, K., Wei, K., Howison, J., and Wiggins, A. (2006). Free/libre open-source software development: What we know and what we do not know. *ACM Computing Surveys*, 44(2):7:1–7:35.

Cui, R., Li, J., and Zhang, D. (2020). Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on airbnb. *Management Science*, 66(3):1071–1094.

Doleac, J. and Stein, L. (2013). The visible hand: Race and online market outcomes. *The Economic Journal*, 123(572):F469–F492.

Ducheneaut, N. et al. (2005). Socialization in an open source software community: A socio-technical analysis. *Computer Supported Cooperative Work*, 14(4):323–368.

Edmans, A. (2009). Blockholder trading, market efficiency, and managerial myopia. *The Journal of Finance*, 64(6):2481–2513.

Edmans, A. and Holderness, C. G. (2017). *Blockholders: A Survey of Theory and Evidence*, volume 1. The Handbook of the Economics of Corporate Governance.

Edmans, A. and Manso, G. (2011). Governance through trading and intervention: A theory of multiple blockholders. *The Review of Financial Studies*, 24(7):2395–2428.

Ellinger, E. W., Gregory, R. W., Mini, T., Widjaja, T., and Henfridsson, O. (2024). Skin in the game: The transformational potential of decentralized autonomous organizations. *MIS Quarterly*.

Fama, E. F. and Jensen, M. C. (1983). Separation of ownership and control. *The Journal of Law and Economics*, 26(2):301–325.

Feichtinger, R., Fritsch, R., Vonlanthen, Y., and Wattenhofer, R. (2023). The hidden shortcomings of (d)aos—an empirical study of on-chain governance. In *International Conference on Financial Cryptography and Data Security*, pages 165–185. Springer Nature.

Foerderer, J. and Schuetz, S. W. (2022). Data breach announcements and stock market reactions: a matter of timing? *Management Science*, 68(10):7298–7322.

Freeman, L. (2002). Centrality in social networks: Conceptual clarification. In Scott, J., editor, *Social Network: Critical Concepts in Sociology*.

Freeman, R. B. and Huang, W. (2014). Collaboration: Strength in diversity. *Nature*, 513(7518):305–305.

Gan, J. R., Tsoukalas, G., and Netessine, S. (2023). Decentralized platforms: Governance, tokenomics, and ico design. *Management Science*, 69(11):6667–6683.

Garrison, D. R., Anderson, T., and Archer, W. (1999). Critical inquiry in a text-based environment: Computer conferencing in higher education. *The Internet and Higher Education*, 2(2–3):87–105.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint*.

Gwebu, K. L., Wang, J., and Wang, L. (2018). The role of corporate reputation and crisis response strategies in data breach management. *Journal of Management Information Systems*, 35(2):683–714.

Han, J., Lee, J., and Li, T. (2023). Dao governance. SSRN working paper.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778. IEEE.

Hellmanzik, C. (2010). Location matters: Estimating cluster premiums for prominent modern artists. *European Economic Review*, 54(2):199–218.

Hoffmann, R. and Coate, B. (2022). Fame, What's your name? quasi and statistical gender discrimination in an art valuation experiment c. *Journal of Economic Behavior & Organization*, 202:184–197.

Hofstetter, R., Fritze, M. P., and Lamberton, C. (2023). Beyond Scarcity: A Social Value-Based Lens for NFT Pricing. *SSRN Electronic Journal*.

Holderness, C. G. (2009). The myth of diffuse ownership in the united states. *The Review of Financial Studies*, 22(4):1377–1408.

Howell, S. T., Niessner, M., and Yermack, D. (2020a). Initial coin offerings: Financing growth with cryptocurrency token sales. *The Review of Financial Studies*, 33(9):3925–3974.

Howell, S. T., Niessner, M., and Yermack, D. (2020b). Initial coin offerings: Financing growth with cryptocurrency token sales. *The Review of Financial Studies*, 33(9):3925–3974.

Hsieh, Y.-Y. and Vergne, J.-P. (2022). The future of the web? the coordination and early-stage growth of decentralized platforms. *Strategic Management Journal*, 44(3):829–857.

Hutto, C. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8, pages 216–225.

Imai, K., Keele, L., and Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, 15(4):309.

Jung, T. (2019). How transparency through blockchain helps the cybersecurity community. IBM Think.

Kaczynski, S. and Kominers, S. D. (2021). How nfts create value. *Harvard Business Review*, 10.

Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., and Stulz, R. M. (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics*, 139(3):719–749.

Kane, G., Alavi, M., Labianca, G., and Borgatti, S. (2014). What's different about social media networks? a framework and research agenda. *MIS Quarterly*, 38(1):275–304.

Kaplan, S. (2008). Framing contests: Strategy making under uncertainty. *Organization Science*, 19(5):729–752.

Kaur Nagpal, G. and Renneboog, L. (2023). Valuing Digital Art: On Non-fungible Tokens, Blockchain Hypes, and the Creation of Scarcity. *SSRN Electronic Journal*.

Kelly, E. and O'Hagan, J. (2007). Geographic clustering of economic activity: The case of prominent western visual artists. *Journal of Cultural Economics*, 31:109–128.

Kokkodis, M., Adamopoulos, P., and Ransbotham, S. (2023). Reputation spillover from agencies on online platforms: Evidence from the entertainment industry. *MIS Quarterly*. Forthcoming.

Kräussl, R. and Tugnetti, A. (2024). Non-fungible tokens (nfts): A review of pricing determinants, applications and opportunities. *Journal of Economic Surveys*, 38(2):555–574.

Laouénan, M. and Rathelot, R. (2022). Can information reduce ethnic discrimination? evidence from airbnb. *American Economic Journal: Applied Economics*, 14(1):107–132.

Laturnus, V. (2023). The economics of decentralized autonomous organizations. SSRN working paper.

LeBlanc, A. and Sheppard, S. (2022). Women artists: gender, ethnicity, origin and contemporary prices. *Journal of Cultural Economics*, 46(3):439–481.

Lee, J. W. and Lee, S. H. (2019). User participation and valuation in digital art platforms: the case of Saatchi Art. *European Journal of Marketing*, 53(6):1125–1151.

Lee, K., Park, J., Goree, S., Crandall, D., and Ahn, Y.-Y. (2024). Social signals predict contemporary art prices better than visual features, particularly in emerging markets. *Scientific Reports*, 14(1):11615.

Li, Y., Ma, X., and Renneboog, L. (2024). In art we trust. *Management Science*, 70(1):98–127.

Lovo, S. and Spaenjers, C. (2018). A Model of Trading in the Art Market. *American Economic Review*, 108(3):744–774.

Lumineau, F., Wang, W., and Schilke, O. (2021). Blockchain governance—a new way of organizing collaborations? *Organization Science*, 32(2):500–521.

Malone, T. and Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87–119.

McCahery, J., Sautner, Z., and Starks, L. (2016). Behind the scenes: The corporate governance preferences of institutional investors. *Journal of Finance*, 71(6):2905–2932.

Mekacher, A., Bracci, A., Nadini, M., Martino, M., Alessandretti, L., Aiello, L. M., and Baronchelli, A. (2022). Heterogeneous rarity patterns drive price dynamics in NFT collections. arXiv:2204.10243 [physics, q-fin].

Miranda, S., Wang, D., and Tian, C. (2022). Discursive fields and the diversity-coherence paradox: An ecological perspective on the blockchain community discourse. *MIS Quarterly*, 46(3).

Mohan, V., Khezr, P., and Berg, C. (2024). Voting with time commitment for decentralized governance: Bond voting as a sybil-resistant mechanism. *Management Science*, 70(12):8709–8734.

Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing*, 29(1):1–16.

Murray, A., Kuban, S., Josefy, M., and Anderson, J. (2021). Contracting in the smart era: The implications of blockchain and decentralized autonomous organizations for contracting and corporate governance. *Academy of Management Perspectives*, 35(4):622–641.

Oberlin, K. C. and Gieryn, T. F. (2015). Place and culture-making: Geographic clumping in the emergence of artistic schools. *Poetics*, 50:20–43.

Oh, S., Rosen, S., and Zhang, A. L. (2022). Investor Experience Matters: Evidence from Generative Art Collections on the Blockchain. *SSRN Electronic Journal*.

Pagano, M. and Röell, A. (1998). The choice of stock ownership structure: Agency costs, monitoring, and the decision to go public. *Quarterly Journal of Economics*, 113(1):187–225.

Parker, G. G. and Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science*, 51(10):1494–1504.

Pavlou, P. A. and Gefen, D. (2004). Building effective online marketplaces with institution-based trust. *Information Systems Research*, 15(1):37–59.

Rabetti, D. (2024). Auditing decentralized finance (defi) protocols. Working Paper.

Ransbotham, S., Kane, G., and Lurie, N. (2012). Network characteristics and the value of collaborative user-generated content. *Marketing Science*, 31(3):387–405.

Ray, G., Barney, J. B., and Muhanna, W. A. (2004). Capabilities, business processes, and competitive advantage: Choosing the dependent variable in empirical tests of the resource-based view. *Strategic Management Journal*, 25(1):23–37.

Raymond, E. (2001). *The Cathedral and the Bazaar*. O'Reilly Media, Sebastopol, CA.

Reimers, I. and Waldfogel, J. (2021). Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review*, 111(6):1944–1971.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint*.

Renneboog, L. and Spaenjers, C. (2013). Buying Beauty: On Prices and Returns in the Art Market. *Management Science*, 59(1):36–53.

Resnick, P. and Zeckhauser, R. (2002). Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In *The Economics of the Internet and E-commerce*, pages 127–157. Emerald Group Publishing Limited.

Reyna, A., Martin, C., Chen, J., Soler, E., and Díaz, M. (2018). On blockchain and its integration with iot: Challenges and opportunities for distributed ledger technologies. *Future Generation Computer Systems*, 88:173–190.

Richardson, V. J., Smith, R. E., and Watson, M. W. (2019). Much ado about nothing: The (lack of) economic impact of data privacy breaches. *Journal of Information Systems*, 33(3):227–265.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects.

Rourke, L. (2001). Assessing social presence in asynchronous text-based computer conferencing. *Journal of Distance Education*, pages 1–17.

Schanke, S., Ray, G., Adomavicius, G., and Wagle, M. (2023). Content and style of firm-generated posts on social media: A study of user engagement on hedonic and utilitarian product pages on facebook. *Journal of the Association for Information Systems*, 25(2):341–361.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.

Shleifer, A. and Vishny, R. W. (1986). Large shareholders and corporate control. *Journal of Political Economy*, 94(3):461–488.

Sockin, M. and Xiong, W. (2023). Decentralization through tokenization. *Journal of Finance*, 78(1):247–299.

Tirole, J. (2021). Digital dystopia. *American Economic Review*, 111(6):2007–2048.

Toulmin, S. (1958). *The Uses of Argument*. Cambridge University Press, Cambridge, UK.

Tsoukalas, G. and Falk, B. H. (2020). Token-weighted crowdsourcing. *Management Science*, 66(9):3843–3859.

Velthuis, O. (2005). *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art*. Princeton University Press.

Velthuis, O. (2013). GLOBALIZATION OF MARKETS FOR CONTEMPORARY ART: Why local ties remain dominant in Amsterdam and Berlin. *European Societies*, 15(2):290–308.

von Hippel, E. and von Krogh, G. (2003). Open source software and the "private-collective" innovation model: Issues for organization science. *Organization Science*, 14(2):209–223.

Waldfogel, J. (2017). How digitization has created a golden age of music, movies, books, and television. *Journal of Economic Perspectives*, 31(3):195–214.

Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK.

Weber, S. and Hanssen, B. (2012). The commercial use of open source software: Open source as subversive practice. *MIS Quarterly*, 36(1):1–13.

Whitaker, A. (2019). Art and Blockchain: A Primer, History, and Taxonomy of Blockchain Use Cases in the Arts. *Artivate: A Journal of Entrepreneurship in the Arts*, 8(2).

Whitaker, A. and Abrams, N. B. (2023). The story of nfts: Artists, technology, and democracy. *Rizzoli Electa*.

Whitaker, A. and Kräussl, R. (2020). Fractional Equity, Blockchain, and the Future of Creative Work. *Management Science*, 66(10):4594–4611.

Yang, Y., Tian, T. Y., Woodruff, T. K., Jones, B. F., and Uzzi, B. (2022). Gender-diverse teams produce more novel and higher-impact scientific ideas. *Proceedings of the National Academy of Sciences*, 119(36):e2200841119.

Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance*, 21(1):7–31.

Zhang, J., Ackerman, M., and Adamic, L. (2007). Expertise networks in online communities: Structure and algorithms. In *Proceedings of the 16th International Conference on World Wide Web*, pages 221–230.

Zhao, K., Yen, J., Greer, G., Qiu, B., Mitra, P., and Portier, K. (2014). Finding influential users of online health communities: A new metric based on sentiment influence. *Journal of the American Medical Informatics Association*, 21(e2):e212–e218.

Zhao, X., Ai, P., Lai, F., Luo, X. R., and Benitez, J. (2022). Task management in decentralized autonomous organization. *Journal of Operations Management*, 68(6-7):649–674.

Ziolkowski, R., Miscione, G., and Schwabe, G. (2020). Decision problems in blockchain governance: Old wine in new bottles or walking in someone else's shoes? *Journal of Management Information Systems*, 37(2):316–348.

Appendices

Table 5.1: NFT Artist Gender, Race, and Probability of Sales

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.016	−0.122***	0.013	−0.067	−0.091**	−0.201***
	(0.036)	(0.047)	(0.036)	(0.045)	(0.046)	(0.055)
non-White	−0.341***	−0.215***	−0.364***	−0.297***	−0.343***	−0.244***
	(0.036)	(0.043)	(0.036)	(0.044)	(0.038)	(0.048)
Artist Word Count	−0.001***	−0.001***	−0.001***	−0.001***	−0.001***	−0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig. Qual. Exhibition	0.132**	0.143**	0.005	0.132**	0.135**	−0.029
	(0.068)	(0.068)	(0.082)	(0.068)	(0.068)	(0.082)
Sig. Qual. Client	0.112***	0.114***	0.113***	0.096***	0.109***	0.090**
	(0.032)	(0.032)	(0.032)	(0.037)	(0.032)	(0.037)
Premium Location	−0.041	−0.047	−0.039	−0.042	−0.087***	−0.085**
	(0.030)	(0.030)	(0.030)	(0.030)	(0.033)	(0.033)
Female × Artist Word Count		0.002***				0.002***
		(0.001)				(0.001)
non-White × Artist Word Count		−0.002***				−0.003***
		(0.000)				(0.001)
Female × Sig. Qual. Exhibition			0.207			0.073
			(0.196)			(0.199)
non-White × Sig. Qual. Exhibition			0.510***			0.844***
			(0.160)			(0.175)
Female × Sig. Qual. Client				0.200***		0.042
				(0.073)		(0.082)
non-White × Sig. Qual. Client				−0.097		0.124
				(0.066)		(0.081)
Female × Premium Location					0.259***	0.205***
					(0.073)	(0.074)
non-White × Premium Location					0.046	0.090
					(0.106)	(0.108)

Continued on next page

Table 5.1 continued

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	34 492	34 492	34 492	34 492	34 492	34 492
AIC	35 101.1	35 064.9	35 094.3	35 095.9	35 092.0	35 040.6
BIC	35 971.3	35 952.0	35 981.4	35 983.0	35 979.1	35 978.4
Log. Lik.	−17 447	−17 427	−17 442	−17 442	−17 440	−17 409

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at the artist level. Continuous control variables are winsorized at the 1% level to reduce the influence of extreme outliers.

Table 5.2: NFT Artist Gender, Race, and Primary Sales Price

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.011	−0.009	−0.011	−0.017	−0.009	−0.012
	(0.009)	(0.012)	(0.009)	(0.011)	(0.013)	(0.014)
non-White	0.001	−0.029***	−0.010	−0.017	−0.019*	−0.054***
	(0.010)	(0.011)	(0.010)	(0.011)	(0.011)	(0.013)
Artist Word Count	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig. Qual. Exhibition	0.022	0.016	−0.037	0.023	0.025	−0.030
	(0.023)	(0.023)	(0.026)	(0.023)	(0.023)	(0.026)
Sig. Qual. Client	0.004	0.004	0.005	−0.007	0.005	−0.002
	(0.008)	(0.008)	(0.008)	(0.009)	(0.008)	(0.009)
Sig. Qual. Location	−0.024***	−0.023***	−0.023***	−0.023***	−0.032***	−0.030***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)
Female × Artist Word Count		0.000				0.000
		(0.000)				(0.000)
non-White × Artist Word Count		0.001***				0.000**
		(0.000)				(0.000)
Female × Sig. Qual. Exhibition			0.066			0.061
			(0.075)			(0.074)
non-White × Sig. Qual. Exhibition			0.193***			0.165***
			(0.052)			(0.055)
Female × Sig. Qual. Client				0.020		0.028
				(0.019)		(0.023)
non-White × Sig. Qual. Client				0.059***		0.024
				(0.020)		(0.022)
Female × Sig. Qual. Location					−0.007	−0.011
					(0.018)	(0.018)
non-White × Sig. Qual. Location					0.107***	0.116***
					(0.025)	(0.025)

Continued on next page

Table 5.2 continued

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	18 242	18 242	18 242	18 242	18 242	18 242
R^2	0.638	0.638	0.638	0.638	0.638	0.639
Adj. R^2	0.636	0.636	0.636	0.636	0.636	0.637

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.3: NFT Artist Gender, Race, and Probability of Sales (No-Star Sample)

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.007	−0.149***	0.004	−0.076*	−0.106**	−0.224***
	(0.036)	(0.047)	(0.037)	(0.046)	(0.047)	(0.055)
non-White	−0.358***	−0.233***	−0.382***	−0.316***	−0.354***	−0.252***
	(0.037)	(0.044)	(0.037)	(0.045)	(0.039)	(0.049)
Artist Word Count	−0.001***	−0.001***	−0.001***	−0.001***	−0.001***	−0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig. Qual. Exhibition	0.138**	0.153**	0.008	0.138**	0.143**	−0.021
	(0.069)	(0.069)	(0.083)	(0.069)	(0.069)	(0.084)
Sig. Qual. Client	0.138***	0.140***	0.140***	0.122***	0.136***	0.121***
	(0.032)	(0.032)	(0.032)	(0.038)	(0.032)	(0.038)
Sig. Qual. Location	−0.062**	−0.070**	−0.060**	−0.064**	−0.108***	−0.106***
	(0.031)	(0.031)	(0.031)	(0.031)	(0.034)	(0.034)
Female × Artist Word Count		0.003***				0.002***
		(0.001)				(0.001)
non-White × Artist Word Count		−0.002***				−0.004***
		(0.000)				(0.001)
Female × Sig. Qual. Exhibition			0.228			0.073
			(0.198)			(0.202)
non-White × Sig. Qual. Exhibition			0.508***			0.855***
			(0.160)			(0.177)
Female × Sig. Qual. Client				0.199***		0.014
				(0.074)		(0.083)
non-White × Sig. Qual. Client				−0.097		0.130
				(0.068)		(0.083)
Female × Sig. Qual. Location					0.275***	0.215***
					(0.074)	(0.075)
non-White × Sig. Qual. Location					0.000	0.020
					(0.114)	(0.116)

Continued on next page

Table **??** *continued*

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	32 903	32 903	32 903	32 903	32 903	32 903
AIC	33 648.5	33 609.2	33 641.8	33 643.7	33 638.8	33 585.5
BIC	34 513.9	34 491.4	34 523.9	34 525.9	34 520.9	34 518.0
Log. Lik.	−16 721	−16 699	−16 715	−16 716	−16 714	−16 681

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.4: NFT Artist Gender, Race, and Primary Sales Price

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.011	0.014	0.011	0.005	0.017	0.014
	(0.008)	(0.010)	(0.008)	(0.009)	(0.011)	(0.012)
non-White	0.000	−0.022**	−0.010	−0.008	−0.005	−0.027**
	(0.008)	(0.009)	(0.008)	(0.009)	(0.009)	(0.011)
Artist Word Count	0.000***	0.000	0.000***	0.000***	0.000***	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig. Qual. Exhibition	0.053***	0.049***	0.004	0.054***	0.053***	0.008
	(0.019)	(0.019)	(0.021)	(0.019)	(0.019)	(0.021)
Sig. Qual. Client	−0.003	−0.003	−0.002	−0.009	−0.002	−0.005
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Sig. Qual. Location	−0.012**	−0.011*	−0.011*	−0.012**	−0.012*	−0.011*
	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Female × Artist Word Count		0.000				0.000
		(0.000)				(0.000)
non-White × Artist Word Count		0.001***				0.000**
		(0.000)				(0.000)
Female × Sig. Qual. Exhibition			0.065			0.064
			(0.052)			(0.051)
non-White × Sig. Qual. Exhibition			0.172***			0.144***
			(0.046)			(0.049)
Female × Sig. Qual. Client				0.017		0.022
				(0.016)		(0.019)
non-White × Sig. Qual. Client				0.027*		0.001
				(0.016)		(0.017)
Female × Sig. Qual. Location					−0.015	−0.017
					(0.015)	(0.015)
non-White × Sig. Qual. Location					0.031	0.040**
					(0.019)	(0.019)

Continued on next page

Table 5.4 continued

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	22 656	22 656	22 656	22 656	22 656	22 656
R^2	0.590	0.590	0.590	0.590	0.590	0.590
Adj. R^2	0.588	0.588	0.588	0.588	0.588	0.588

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.5: NFT Artist Gender, Race, and Probability of Sales (PSM)

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.064	−0.001	0.050	−0.024	−0.044	−0.107
	(0.046)	(0.059)	(0.047)	(0.057)	(0.056)	(0.067)
non-White	−0.230***	−0.069	−0.257***	−0.166***	−0.193***	−0.056
	(0.043)	(0.052)	(0.044)	(0.052)	(0.047)	(0.059)
Artist Word Count	0.000	0.001***	0.000	0.000	0.000	0.002***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)
Sig. Qual. Exhibition	0.012	0.022	−0.232**	0.017	0.018	−0.359***
	(0.088)	(0.089)	(0.118)	(0.088)	(0.088)	(0.122)
Sig. Qual. Client	0.030	0.042	0.029	0.035	0.023	−0.016
	(0.046)	(0.046)	(0.046)	(0.061)	(0.046)	(0.062)
Sig. Qual. Location	0.010	−0.003	0.016	0.005	−0.041	−0.044
	(0.047)	(0.047)	(0.047)	(0.047)	(0.060)	(0.060)
Female × Artist Word Count		0.001				−0.001
		(0.001)				(0.001)
non-White × Artist Word Count		−0.003***				−0.004***
		(0.001)				(0.001)
Female × Sig. Qual. Exhibition			0.492**			0.495**
			(0.244)			(0.249)
non-White × Sig. Qual. Exhibition			0.501***			0.859***
			(0.170)			(0.182)
Female × Sig. Qual. Client				0.219**		0.136
				(0.092)		(0.103)
non-White × Sig. Qual. Client				−0.157*		0.073
				(0.083)		(0.094)
Female × Sig. Qual. Location					0.331***	0.287***
					(0.099)	(0.100)
non-White × Sig. Qual. Location					−0.180	−0.141
					(0.112)	(0.113)

Continued on next page

Table 5.5 continued

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	16 537	16 537	16 537	16 537	16 537	16 537
AIC	17 415.3	17 387.2	17 407.3	17 409.6	17 405.2	17 360.5
BIC	18 209.8	18 197.1	18 217.2	18 219.5	18 215.1	18 216.7
Log. Lik.	−8604	−8588	−8598	−8599	−8597	−8569

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.6: NFT Artist Gender, Race, and Primary Sales Price (PSM)

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.079***	−0.088***	−0.080***	−0.084***	−0.091***	−0.104***
	(0.010)	(0.013)	(0.010)	(0.012)	(0.012)	(0.015)
non-White	−0.031***	−0.080***	−0.045***	−0.045***	−0.045***	−0.092***
	(0.010)	(0.012)	(0.010)	(0.012)	(0.012)	(0.014)
Artist Word Count	0.000**	0.000***	0.000	0.000*	0.000**	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sig. Qual. Exhibition	−0.015	−0.012	−0.130***	−0.014	−0.012	−0.088***
	(0.025)	(0.025)	(0.032)	(0.025)	(0.025)	(0.031)
Sig. Qual. Client	0.019*	0.015	0.019*	0.001	0.017	0.013
	(0.011)	(0.011)	(0.011)	(0.015)	(0.011)	(0.015)
Sig. Qual. Location	0.010	0.012	0.014	0.010	−0.016	−0.013
	(0.011)	(0.011)	(0.011)	(0.011)	(0.015)	(0.015)
Female × Artist Word Count		0.000**				0.000
		(0.000)				(0.000)
non-White × Artist Word Count		0.001***				0.001***
		(0.000)				(0.000)
Female × Sig. Qual. Exhibition			0.088*			0.063
			(0.052)			(0.051)
non-White × Sig. Qual. Exhibition			0.241***			0.162***
			(0.048)			(0.051)
Female × Sig. Qual. Client				0.018		0.023
				(0.021)		(0.023)
non-White × Sig. Qual. Client				0.042**		−0.013
				(0.021)		(0.022)
Female × Sig. Qual. Location					0.039*	0.037
					(0.023)	(0.023)
non-White × Sig. Qual. Location					0.066***	0.074***
					(0.025)	(0.025)

Continued on next page

Table 5.6 continued

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	10 259	10 259	10 259	10 259	10 259	10 259
R2	0.661	0.663	0.662	0.661	0.661	0.663
R2 Adj.	0.657	0.659	0.659	0.658	0.658	0.660

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.7: NFT Artist Gender, Race, and Probability of Sales (LIWC Set 1)

	Dependent Variable: Minting						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	−0.084**	−0.237**	−0.022	−0.052	−0.120*	−0.120***	−0.065
	(0.040)	(0.112)	(0.038)	(0.042)	(0.066)	(0.046)	(0.045)
non-White	−0.303***	−0.357***	−0.265***	−0.310***	−0.213***	−0.341***	−0.293***
	(0.040)	(0.102)	(0.038)	(0.045)	(0.065)	(0.045)	(0.045)
Power	−0.039***						
	(0.004)						
Female × Power	0.040***						
	(0.011)						
non-White × Power	0.035***						
	(0.008)						
Confidence		−0.003***					
		(0.001)					
Female × Confidence		0.004**					
		(0.002)					
non-White × Confidence		0.002					
		(0.001)					
Money			0.016***				
			(0.006)				
Female × Money			0.089**				
			(0.042)				
non-White × Money			0.044***				
			(0.015)				
Social Words				−0.009***			
				(0.002)			
Female × Social Words				0.012**			
				(0.005)			
non-White × Social Words				0.014***			
				(0.004)			

Continued on next page

Table 5.7 continued

	Dependent Variable: Minting						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Emotional Tone					−0.002***		
					(0.000)		
Female × Emotional Tone					0.002**		
					(0.001)		
non-White × Emotional Tone					−0.001		
					(0.001)		
Affection						−0.008***	
						(0.002)	
Female × Affection						0.018***	
						(0.005)	
non-White × Affection						0.015***	
						(0.004)	
Positive Emotion							−0.009***
							(0.002)
Female × Positive Emotion							0.012**
							(0.006)
non-White × Positive Emotion							0.010*
							(0.005)
Artist Controls	YES	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES	YES
Num. Obs.	35 326	35 326	35 326	35 326	35 326	35 326	35 326
AIC	36 264.4	36 363.0	36 338.4	36 362.0	36 342.1	36 335.5	36 359.9
BIC	37 162.5	37 261.1	37 236.5	37 260.1	37 240.1	37 233.6	37 258.0

Continued on next page

Table 5.7 continued

	Dependent Variable: Minting						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log. Lik.	−18 026	−18 075	−18 063	−18 075	−18 065	−18 061	−18 073

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.8: NFT Artist Gender, Race, and Probability of Sales (LIWC Set 2)

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.078**	0.000	−0.008	−0.050	0.083**	−0.002
	(0.039)	(0.035)	(0.035)	(0.037)	(0.037)	(0.035)
non-White	−0.282***	−0.252***	−0.256***	−0.265***	−0.153***	−0.216***
	(0.038)	(0.037)	(0.037)	(0.037)	(0.038)	(0.037)
Negative Emotion	−0.007**					
	(0.003)					
Female × Negative Emotion	0.089***					
	(0.021)					
non-White × Negative Emotion	0.047***					
	(0.011)					
Anxious		0.002				
		(0.010)				
Female × Anxious		−0.017				
		(0.045)				
non-White × Anxious		0.148***				
		(0.051)				
Anger			0.000			
			(0.005)			
Female × Anger			0.039*			
			(0.020)			
non-White × Anger			0.029**			
			(0.012)			
Sad				−0.028**		
				(0.011)		
Female × Sad				0.130***		
				(0.040)		
non-White × Sad				0.149***		
				(0.036)		

Continued on next page

Table 5.8 continued

	Dependent Variable: Minting					
	(1)	(2)	(3)	(4)	(5)	(6)
Death					0.058***	
					(0.008)	
Female × Death					−0.039***	
					(0.007)	
non-White × Death					−0.062***	
					(0.008)	
Swear						0.009
						(0.006)
Female × Swear						0.001
						(0.164)
non-White × Swear						−0.275***
						(0.066)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	35 326	35 326	35 326	35 326	35 326	35 326
AIC	36 316.9	36 369.2	36 365.9	36 319.1	36 232.9	36 356.1
BIC	37 215.0	37 267.2	37 264.0	37 217.2	37 131.0	37 254.2
Log. Lik.	−18 052	−18 078	−18 076	−18 053	−18 010	−18 072

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.9: NFT Artist Gender, Race, and Primary Sales Price (LIWC Set 2)

	Dependent Variable: Primary Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	−0.025***	−0.006	−0.005	−0.007	−0.019	−0.030***	−0.035***
	(0.008)	(0.019)	(0.008)	(0.008)	(0.013)	(0.009)	(0.009)
non-White	−0.025***	−0.082***	−0.025***	−0.042***	−0.027**	−0.015*	−0.017*
	(0.008)	(0.021)	(0.008)	(0.009)	(0.013)	(0.009)	(0.009)
Power	0.004***						
	(0.001)						
Female × Power	0.014***						
	(0.002)						
non-White × Power	−0.001						
	(0.001)						
Confidence		−0.001**					
		(0.000)					
Female × Confidence		0.000					
		(0.000)					
non-White × Confidence		0.001***					
		(0.000)					
Money			−0.002***				
			(0.001)				
Female × Money			0.003				
			(0.008)				
non-White × Money			0.000				
			(0.002)				
Social Words				−0.001***			
				(0.000)			
Female × Social Words				0.001			
				(0.001)			
non-White × Social Words				0.003***			
				(0.001)			
Emotional Tone					0.000		
					(0.000)		

Table 5.9 continued

	Dependent Variable: Primary Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female × Emotional Tone					0.000		
					(0.000)		
non-White × Emotional Tone					0.000		
					(0.000)		
Affection						−0.002***	
						(0.000)	
Female × Affection						0.004***	
						(0.001)	
non-White × Affection						−0.001**	
						(0.001)	
Positive Emotion							−0.003***
							(0.000)
Female × Positive Emotion							0.007***
							(0.001)
non-White × Positive Emotion							−0.002*
							(0.001)
Artist Controls	YES	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES	YES
Num. Obs.	23 919	23 919	23 919	23 919	23 919	23 919	23 919
R2	0.670	0.669	0.669	0.669	0.669	0.670	0.670
R2 Adj.	0.669	0.667	0.667	0.667	0.667	0.668	0.668

* p < 0.1, ** p < 0.05, *** p < 0.01

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.

Table 5.10: NFT Artist Gender, Race, and Primary Sales Price (LIWC Set 1)

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	−0.007	−0.003	−0.003	−0.003	−0.014*	−0.002
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
non-White	−0.020**	−0.026***	−0.022***	−0.019**	−0.033***	−0.025***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Negative Emotion	−0.002***					
	(0.000)					
Female × Negative Emotion	0.004**					
	(0.002)					
non-White × Negative Emotion	−0.003***					
	(0.001)					
Anxious		−0.002*				
		(0.001)				
Female × Anxious		0.011				
		(0.010)				
non-White × Anxious		0.006				
		(0.011)				
Anger			−0.002			
			(0.001)			
Female × Anger			−0.002			
			(0.002)			
non-White × Anger			−0.003*			
			(0.001)			
Sad				0.009***		
				(0.003)		
Female × Sad				−0.002		
				(0.003)		
non-White × Sad				−0.024***		
				(0.004)		

Continued on next page

Table 5.10 continued

	Dependent Variable: Primary Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Death					−0.007***	
					(0.001)	
Female × Death					0.007***	
					(0.002)	
non-White × Death					0.007***	
					(0.002)	
Swear						−0.004***
						(0.001)
Female × Swear						−0.072***
						(0.023)
non-White × Swear						0.004
						(0.012)
Artist Controls	YES	YES	YES	YES	YES	YES
NFT Visual Embeddings	YES	YES	YES	YES	YES	YES
NFT Genre FE	YES	YES	YES	YES	YES	YES
NFT Media FE	YES	YES	YES	YES	YES	YES
NFT Controls	YES	YES	YES	YES	YES	YES
Transaction Controls	YES	YES	YES	YES	YES	YES
Year–Month FE	YES	YES	YES	YES	YES	YES
Num. Obs.	23 919	23 919	23 919	23 919	23 919	23 919
R2	0.669	0.669	0.669	0.669	0.670	0.669
R2 Adj.	0.668	0.667	0.667	0.668	0.668	0.668

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note. Robust standard errors are clustered at artist-level. Continuous control variables are winsorized at 1% to reduce the influence of extreme outliers.