



PhD-FSTM-2025-040
The Faculty of Science, Technology and Medicine

DISSERTATION

Defence held on 3 April 2025 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG EN INFORMATIQUE

by

Orestis PAPAGEORGIOU

Born on 10 September 1990 in Amarousio, Attica (Greece)

BLOCKCHAIN GOVERNANCE: PROMISES, MECHANISMS, AND BROKEN IDEALS

Dissertation defence committee

Dr Gilbert Fridgen, dissertation supervisor
Professor, Université du Luxembourg

Dr Andre Luckow
Head of Innovation and Emerging Technologies, BMW Group, München

Dr Marcus Völp, Chairman
Associate Professor, Université du Luxembourg

Dr Robert Keller
Professor, Hochschule für angewandte Wissenschaften Kempten

Dr Katrin Becker, Vice Chair
Research scientist, Université du Luxembourg

Acknowledgements

First, I would like to express my gratitude to my supervisor, Prof. Dr. Gilbert Fridgen, for granting me the opportunity to work within an interdisciplinary and dynamic research group. His guidance and encouragement through the ups and downs of my PhD journey have been invaluable.

I deeply appreciate the support from Dr. Johannes Sedlmeir, whose technical expertise and supervision have been instrumental in my learning and growth. His dedication to helping me navigate complex challenges has greatly enriched my experience and understanding.

To my colleagues, thank you for making this journey more enjoyable and for providing constant support—whether it was through sharing insights or simply tolerating my tendency to “express” every challenge.

Lastly, I would like to acknowledge the unwavering support of my family and friends. Their encouragement has been a source of strength throughout my academic journey.

This research was funded by the Luxembourg National Research Fund (FNR), in the PABLO (grant reference 16326754) and NCER22/IS/16570468/NCER-FT projects, and by PayPal, PEARL grant reference 13342933/Gilbert Fridgen. For the purpose of open access and in fulfillment of the obligations arising from the grant agreement, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

Abstract

The rapid advancement and adoption of blockchain technology have fundamentally transformed various aspects of digital interaction, leading to the emergence of novel governance frameworks that challenge traditional centralized models. At the forefront of this transformation are decentralized autonomous organizations (DAOs) and an array of other governance structures applied in both permissionless and permissioned blockchain environments. Unlike conventional organizations that rely on hierarchical authority, DAOs within permissionless systems strive to operate through decentralized networks where decision-making power is distributed among all members, facilitated by smart contracts and governance tokens. In parallel, permissioned blockchain applications, often employed by consortia or enterprises, experiment with more structured membership and delegated authority, blending decentralized principles with selective participation to maintain compliance, accountability, and operational efficiency.

These governance mechanisms, whether in DAOs or permissioned networks, are envisioned to enhance transparency, inclusivity, and autonomy. Yet, despite their idealistic promises, practical implementations have revealed significant challenges. Within DAOs, governance tokens intended to promote equitable decision-making often lead to power concentration and stakeholder inequality. Moreover, vulnerabilities in smart contract design and the absence of robust accountability frameworks have produced notable failures. In permissioned contexts, while governance structures can mitigate some of these issues through established roles and clearer recourse mechanisms, complexities arise in balancing decentralized ideals with enterprise-grade stability and oversight.

This thesis critically examines the foundational principles of blockchain-based governance—spanning from permissionless DAOs to permissioned consortia—alongside their operational realities and limitations. It explores the effectiveness of governance tokens and the vulnerabilities undermining participatory ideals. It further examines whether emerging innovations, such as quadratic voting, market-based, and NFT-based voting mechanisms, mitigate any of the identified issues.

Declaration

I, Orestis Papageorgiou, declare that I wrote this thesis by myself and that it has not been previously submitted for any other degree or professional qualification. For the jointly authored research papers, I have clearly distinguished my contributions from those of the other authors. I confirm that I have correctly referenced the work of others whenever I refer to it.

I used the AI tools Grammarly, and ChatGPT to improve the language and clarity of parts of my work. However, I carefully reviewed the output from these tools to ensure understanding and accuracy, incorporating only selected suggestions.

I confirm that I have no financial interests to disclose in relation to this research. I commit to conducting my research with transparency, integrity, and adherence to ethical principles, ensuring the reliability of my work.

Orestis Papageorgiou

Table of Contents

List of Figures	iii
List of Tables	iv
I Introduction	1
1 Motivation	1
2 Thesis Structure	3
II Governance in Permissionless Blockchains	5
1 Core Principles & Governance Process	5
2 Practical Challenges	9
III Governance in Permissioned Blockchains	16
1 Core Principles & Governance Process	16
2 Practical Challenges	19
IV Innovations in Governance Mechanisms	22
1 Quadratic Voting	22
2 NFT-Based Voting	23
3 Market-Based Voting	24
V Conclusion	26
1 Summary	26
2 Limitations and outlook	27
3 Recognition of previous and related work	28
VI References	30

TABLE OF CONTENTS

A	Appendix	43
1	Publication portfolio	43
2	Individual Contribution	44
3	Appended Research Papers	46

| List of Figures

II.1	Levels of centralization of nine different DeFi protocols according to four metrics. Higher values indicate higher levels of centralization (c.f. RP2-[11]). .	11
II.2	Minimum number of addresses required to control more than 50 % of the governance tokens (c.f. RP2-[11]).	12
II.3	Governance participation rates. Each point corresponds to a proposal (c.f. RP2-[11]).	14

| List of Tables

I.1	Overview of research publications.	4
-----	--	---

I | Introduction

1 Motivation

Since its creation by Satoshi Nakamoto [61], blockchain technology has been investigated for applications in numerous industries well beyond its initial role in powering the cryptocurrency Bitcoin. Some of the technology’s fundamental attributes (transparency, immutability, and security) [105], combined with the use of smart contracts (self-executing code that automatically enforces predefined rules without intermediaries) [2] first implemented at scale on the Ethereum blockchain [17], have resonated with organizations [7, 41, 70, 79], developers [23], and policymakers [30, 52, 80] seeking new ways to structure digital systems. Over the years, the technology shifted from the sole domain of trustless, permissionless networks, such as those in Bitcoin or Ethereum, into the domains of enterprises and traditional financial institutions. As a result, blockchain implementations can be broadly separated into two categories: permissionless blockchains, which are open to anyone without formal entry barriers, and permissioned blockchains, which require authorization to join and participate [105].

These distinctions carry significant implications for the nature of blockchain applications. In a permissionless setting, applications include, among others, decentralized finance (DeFi) protocols [1, 99] and non-fungible token (NFT) marketplaces [64], where the guiding ethos is to minimize central authorities and intermediaries, thereby reducing central control and fees [100, 101]. Conversely, permissioned blockchains, often employed by industry consortia or large enterprises, focus on use cases that demand more transparent accountability and regulatory compliance [87] with predictable and

better performance [48]. Some of these include supply chain tracking [47], financial asset tokenization [70], and digital identities [83].

Whether it occurs in a permissionless or permissioned setting, decision-making processes in blockchain systems must determine how changes are proposed, voted, and executed [85, 104]. Governance, in this context, encompasses an array of mechanisms, from updating protocol rules and utilizing treasury funds to managing infrastructure costs and settling disputes [57]. Notably, blockchain governance aspires to deviate from traditional top-down or corporate models of organizational oversight and seeks to embody values such as decentralization, transparency, and equitable participation [59]. In theory, every stakeholder should have a voice in shaping the protocol's evolution.

The concept of decentralized autonomous organizations (DAOs) epitomizes this ambition in permissionless blockchains. DAOs are designed to operate without centralized leadership, relying instead on smart contracts to enforce governance decisions [100]. In theory, DAOs promise a level of inclusivity and autonomy that goes beyond the traditional organizational hierarchy [81]. Governance within these entities mainly leverages native tokens that grant voting power (governance tokens) proportional to token holdings [11, 22]. This approach ostensibly ensures that all participants have an opportunity to influence outcomes, with the decision-making process being transparent and reliant on token-based participation by design [57, 81].

However, practice reveals substantial deviations from these ideals. Power dynamics within DAOs often skew towards large token holders RP2-[11], RP3-[10] [44, 46, 55], mirroring the inequality found in conventional corporate boards dominated by key shareholders. Rather than achieving true decentralization, many DAOs witness the emergence of unofficial leadership cliques that hold disproportionate power over decision-making. Vulnerabilities in the underlying smart contract code [26] further complicate governance processes, introducing the risk of hacks that can undermine user trust. Thus, while the architecture of DAOs differs from traditional governance frameworks, the real-world outcomes frequently replicate the concentration of influence and strategic manipulation that decentralized principles were meant to avoid while simultaneously introducing novel risks such as smart contract vulnerabilities.

In the permissioned sphere, participants in governance are known entities—regulated financial institutions, corporate partners, government agencies, etc.—and the gover-

nance structures can embed contractual obligations and legally enforceable agreements [104]. This setting allows for more conventional recourse mechanisms and transparent accountability, with decision-making mostly occurring through structured committees, weighted voting processes that consider stakeholder size or investment, or designated roles with clearly defined responsibilities [7]. These measures can help maintain network stability and compliance with relevant regulations.

Nonetheless, integrating a decentralized ethos into such settings is challenging. The requirement for trust and cooperation among known parties can constrain the degree of openness and bottom-up influence, often leading to quasi-decentralized or hybrid models [30, 77]. While permissioned models might avoid some of the security and coordination pitfalls of their permissionless counterparts, they often sacrifice the full extent of decentralization envisioned by early blockchain proponents [56].

As a result, in both permissionless and permissioned contexts, a gap persists between the narrative of blockchain governance and the reality of its operations. Despite initial enthusiasm and ongoing experimentation, neither environment has fully realized a truly equitable, transparent, and resilient governance model. Emerging innovations, such as quadratic voting, market-based, and NFT-based voting mechanisms, attempt to mitigate token-based inequities and align incentives more closely with community values. Yet even these mechanisms are subject to implementation hurdles, strategic manipulations, or misaligned incentives that limit their effectiveness.

The purpose of this work is to critically examine the foundational principles and purported benefits of blockchain governance and to contrast them with the outcomes observed in reality. By scrutinizing both permissionless DAOs and permissioned consortia, the thesis aims to illuminate not just the theoretical potential but the practical limitations of these frameworks. Through this examination, it seeks to provide insights into why blockchain governance often fails to live up to its promises.

2 Thesis Structure

The thesis relies on five research publications, as summarized in Table I.1. For clarity, these research publications will be referred to as "RP" throughout the remainder of this thesis. The thesis is organized into four main sections, followed by a concluding section

that synthesizes the findings and outlines avenues for future research and an Appendix which provides a comprehensive overview of the publications included in the dissertation, contribution statements and the full text of the research papers referenced in the dissertation.

Table I.1: Overview of research publications.

RP#	Title	Role ¹
RP1	What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains	SP
RP2	Decentralised Finance's timocratic governance: The distribution and exercise of tokenised voting rights	JP
RP3	DeFi, Not So Decentralized: The Measured Distribution of Voting Rights	JP
RP4	Pricing dynamics and herding behaviour of NFTs	JP
RP5	Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch	NP

¹ SP = Single Primary Author, JP = Joint Primary Author, NP = Non-Primary Author.

Section **I** highlighted the core tension in blockchain governance between the aspiration toward decentralization, inclusivity, and transparency versus the reality of entrenched power dynamics, strategic manipulation, and technological vulnerabilities. Section **II** examines permissionless governance and how DAOs, despite their decentralized ethos, often fall short, with large token holders exerting disproportionate influence and security vulnerabilities introducing new forms of risk. Section **III** analyses permissioned governance, exploring governance structures in closed consortium chains. It analyzes how authority and decision-making become more legally and contractually defined, potentially ensuring stability and compliance yet restricting the degree of openness and bottom-up participation. Section **IV** looks into emerging innovations in governance mechanisms that aim to close the gap between decentralized ideals and practical outcomes. Finally, Section **V** summarizes key findings, reflecting on the trade-offs in blockchain governance.

II Governance in Permissionless Blockchains

1 Core Principles & Governance Process

The concept of permissionless governance within blockchain ecosystems is deeply rooted in the early experiments that aimed to test the feasibility of decentralized decision-making. The most influential example was “The DAO,” [97], which was launched on the Ethereum network in 2016. Despite its premature failure caused by a critical smart contract vulnerability that allowed malicious actors to siphon funds [20], The DAO remains a pivotal moment in blockchain history. It served as a large-scale, real-world attempt at decentralized governance, demonstrating both the potential and the challenges of such frameworks [57]. The DAO’s core ideas continue to shape many subsequent projects, particularly within the DeFi space, where the emphasis remains on dismantling hierarchical structures in favor of trust-minimized collaboration.

Even after the failure of The DAO, four of its core principles still guide the governance mechanisms of DAOs [57, 96]. Specifically:

- **Decentralized Authority:** Rather than centralizing control in a single leader or committee, DAOs distribute decision-making power among token holders or members. This structure aims to reduce the risk of unilateral decisions and increase the collective accountability of the organization.
- **Open Participation:** DAOs allow anyone who meets certain criteria to create proposals and cast votes. This inclusivity aims to broaden the range of perspectives and expertise within the DAO, strengthening the rigor of governance.

- **Transparent Processes:** DAOs employ publicly accessible records of governance activities, where every proposal, vote, and outcome is visible to all participants. By maintaining clear, auditable records, community members can effectively validate decisions and enhance accountability.
- **Smart Contract-Based Autonomy:** At the core of DAOs are smart contracts that encode the rules and execute decisions, minimizing reliance on intermediaries. This technological backbone promotes consistency and trust by ensuring that governance mechanisms function exactly as programmed.

In practice, DAOs adopt varying governance models to operationalize their principles, which can be broadly categorized into two main approaches [59, 74]. In *on-chain governance*, all processes—including proposal submissions, voting, and execution—are fully automated and executed via smart contracts on the blockchain, ensuring transparency and minimizing the need for human intervention. Conversely, *off-chain governance* allows steps, such as proposing and voting, to occur outside the blockchain, reducing complexity and transaction fees, with execution managed manually by designated developers or community members. While both approaches aim to uphold the core principles, they differ in their reliance on automation. Regardless of the governance model, the functioning of every DAO can be distilled into three stages—proposal¹, voting, and execution [54, 85]. Each stage, in theory, aims to integrate the four core principles, ensuring alignment with the vision of DAOs.

1.1. Proposal

The proposal process is at the core of decentralized governance mechanisms in DAOs, serving as the primary method for community members to introduce changes and suggest improvements [22, 57]. In most protocols, this phase begins in an off-chain governance forum or in a platform like Telegram or Discord, where, in keeping with the principle of **open participation**, any community member, at least in principle, can put forth a proposal [74]. The forum then fosters **transparent processes** by enabling anyone to review, discuss, and refine the proposal. This inclusivity ensures that even those

¹Some scholars, such as [26], further subdivide the proposal stage into distinct phases like deliberation and proposal. Despite this nuanced categorization, the fundamental properties of the stage remain consistent.

without significant technical expertise can shape the DAO's direction, thereby reinforcing the principle of **decentralized authority** as decision-making power is distributed across the community rather than held by a select few [91].

The proposal stage serves as a counterbalance to traditional, centralized power structures by encouraging community input to shape the project's evolution. The open exchange, in theory, refines proposals and allows ideas to originate from every corner of the community rather than a select few [96].

1.2. Voting

In the voting phase, proposals that gain sufficient preliminary support are submitted for a formal vote. Reflecting the principle of **decentralized authority**, governance decisions are placed directly in the hands of DAO members who own governance tokens [100]. Each token holder has the opportunity to vote for or against a proposal, ensuring that no single entity can unilaterally dictate outcomes [25]. Governance tokens, which typically grant voting power proportional to the number of tokens held, can be acquired through decentralized or centralized exchanges. This accessibility underscores the principle of **open participation**, as anyone who holds these tokens can become an active participant in the governance process [25].

This approach aims to be both transparent and meritocratic, as it ensures that those who have invested in the governance token can directly influence its future [59]. The logic underpinning this system is that token holders are incentivized to vote in a way that promotes the long-term success of the project since acting against the best interests of the protocol could harm the value of the token and, consequently, their investments. Additionally, many DAOs establish a quorum [73] requirement to ensure that decisions are made with meaningful community participation, preventing a small minority from unilaterally guiding the project's future.

The principle of **transparent processes** is central to this phase, with vote results immutably recorded on the blockchain. This ensures that every decision is auditable and accessible to all community members, enhancing trust in the governance process [25]. Some protocols further extend inclusivity by allowing token holders to delegate their votes to other members (e.g., [14]). Delegates, typically individuals or organizations

with specialized expertise or a deeper understanding of the issues at hand, can represent the voting power of less active or less informed participants.

1.3. Execution

Once a proposal successfully passes the voting phase, it moves into the execution stage, where the decided changes are implemented according to the DAOs established rules [81]. Depending on the nature of the proposal, execution can be highly automated or may require more manual intervention [74].

In cases where the proposal involves on-chain parameters—such as adjusting interest rates in a lending protocol— **smart contract-based autonomy** typically takes center stage, with smart contracts handling the execution automatically, applying the agreed-upon changes without the need for human intermediaries [100]. This automated approach upholds **decentralized authority** by preventing any single actor from modifying the outcome unilaterally and ensures **transparent processes** through immutable blockchain records of the changes made.

However, not all proposals can be executed purely by code. Some changes may require off-chain actions, such as initiating partnerships, allocating treasury funds, or integrating new features that must be developed and tested by the protocol's core contributors [74]. To preserve **decentralized authority** even in these scenarios, many DAOs use multi-signature (multisig) [71] wallets, distributing responsibility among multiple participants.

To safeguard the integrity of the execution process, DAOs often implement protective measures to prevent malicious proposals, rash decisions, or sudden changes that could jeopardize the protocol. One widely used safeguard is the implementation of time-locks [72]—a delay between the final approval of a proposal and its execution. This waiting period allows the community to review decisions, address potential oversights, and, if necessary, nullify proposals [101], thereby reinforcing **transparent processes** and community oversight before implementation.

2 Practical Challenges

While the foundational principles of DAOs are compelling in theory, real-world implementations often diverge significantly [91, 100]. Each DAO adapts these principles to suit its specific needs, resulting in a wide array of governance models. Beyond theoretical ideals, practical constraints such as varying technical requirements and regulatory considerations lead to differing implementations across projects [57].

The thesis concentrates on DeFi protocols because the stakes for secure and robust governance are particularly high in financial applications. Some of these platforms manage assets valued in the tens of billions of dollars [18], making governance failures not merely academic concerns but genuine threats to user funds and platforms. However, because the DeFi ecosystem features an ever-expanding range of protocols, a fully comprehensive examination remains infeasible. Instead, we draw insights from prominent projects—that are widely adopted, backed by substantial capital, or frequently referenced in academic literature. Despite their distinct technical architectures and community dynamics, all these protocols strive to uphold the four core principles of DAOs outlined in the previous section [57, 101].

Nevertheless, the gap between theory and practice becomes evident when scrutinizing each step of the governance process of DAOs. Their governance models face challenges that undercut each principle. The following subsections focus on the three governance stages to explore how these issues manifest in practical implementations.

2.1. Proposal

Even though DAOs are designed in principle to enable any community member to submit new ideas through the proposal mechanism, in practice, this process is often restricted by various criteria. For instance, some protocols require holding or being delegated a substantial number of governance tokens to submit a proposal [14]. While this can deter frivolous or “spam” proposals, it also risks concentrating power in the hands of those who can afford large token positions or have strong connections with major token holders, ultimately undermining the principle of **decentralized authority**. By effectively pricing out smaller holders, such thresholds also constrain **open participation**, reducing the breadth and diversity of potential contributors.

For example, in Uniswap, the largest decentralized exchange (DEX) in the DeFi space by trading volume and total value locked (TVL) [19], the proposal mechanism exemplifies this issue. To submit a proposal, a user must be delegated 1,000,000 governance tokens [14]—equivalent to 0.1% of the token supply and approximately \$15 million as of December 9th, 2024. This high threshold effectively excludes smaller token holders from submitting proposals.

A related issue arises in DAOs where anyone can technically put forth a proposal, but only those originated by the core team become “binding” or advance to the next phase. This approach addresses the valid concern of preventing unvetted proposals from causing disruptions, yet it can reinforce a quasi-hierarchical structure in which community-driven proposals remain purely advisory. Thus, even when **transparency** is ensured by publicly displaying and discussing proposals, a dual-track system effectively consolidates final authority in the hands of a few. This imbalance dilutes **decentralized authority** by allowing core developers or a select committee to determine which proposals move forward to the voting stage. A notable example of this mechanism is SushiSwap [94], another DEX, where in recent years, proposals not initiated by a member of the core developer team have been systematically excluded from advancing to the voting stage [95], regardless of the community support they may garner during the proposal phase.

2.2. Voting

Even though DAOs conceptually place governance power in the hands of all token holders, large stake concentrations often lead to a high concentration of voting power [44, 55], RP2-[11], RP3-[10]. A few so-called “whales”—individuals or entities with disproportionately large holdings or influence—can dominate the outcome of crucial decisions, thereby undercutting the principle of **decentralized authority**. Moreover, due to the pseudonymity of the blockchain, it is often impossible to determine with certainty whether large token holdings belong to a single whale or are distributed across multiple entities RP2-[11], further complicating governance dynamics. Although voting remains technically open to all, the outsized voting power of top stakeholders effectively sidelines smaller participants, undermining **open participation**.

In RP2-[11], an analysis of governance token distribution across nine of the largest DeFi protocols utilized four metrics (Gini Coefficient, Normalized Shannon Entropy, Normalized Euclidean Distance, and Jensen-Shannon Divergenc)—to evaluate centralization. The findings revealed that, over time, the distribution of voting rights became heavily centralized in every examined protocol (Figure II.1). In some cases, the majority of votes were controlled by as few as five addresses, and in all instances, these addresses represented less than 1% of the total number of governance token holders (Figure II.2). This meant that in every examined case, fewer than 1% of users could effectively determine the outcome of the voting process.

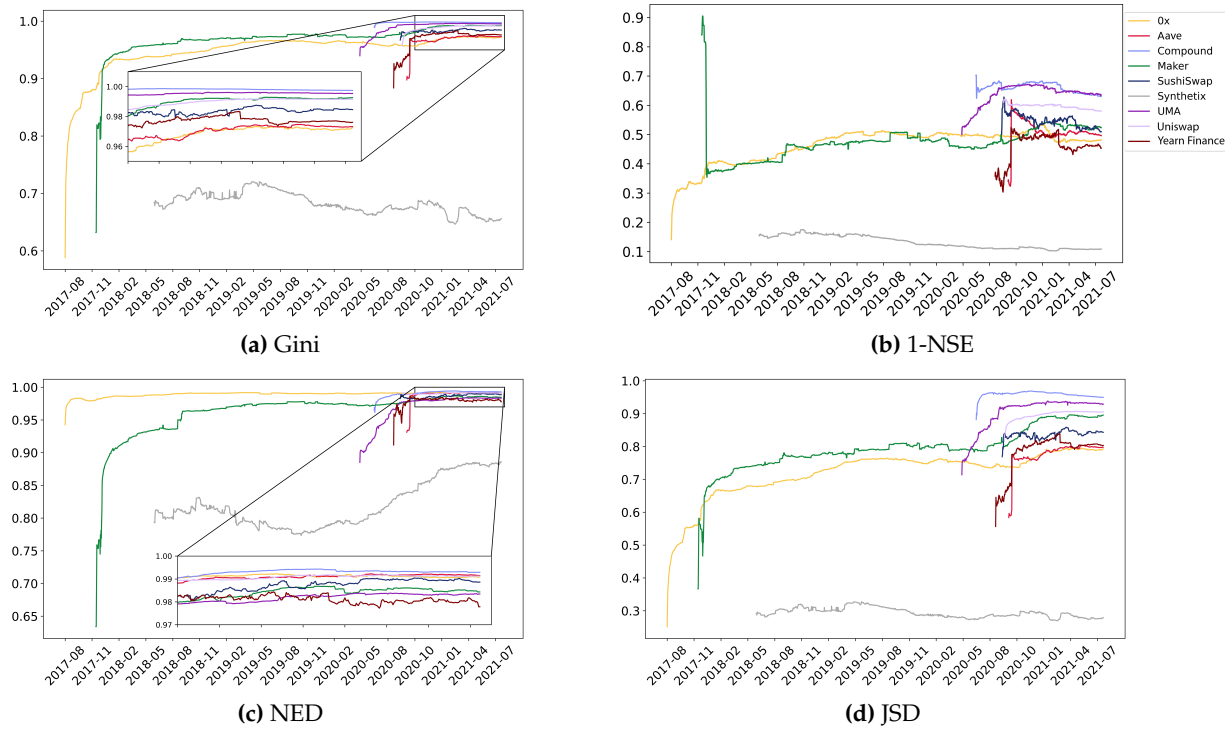


Figure II.1: Levels of centralization of nine different DeFi protocols according to four metrics. Higher values indicate higher levels of centralization (c.f. RP2-[11]).




Figure II.2: Minimum number of addresses required to control more than 50 % of the governance tokens (c.f. [RP2-\[11\]](#)).

The study further found that a significant portion—and in some cases, the majority—of the initial allocation of governance tokens was concentrated among investors or the project’s core team. For example, nearly 80% of the initial allocation of governance tokens for UMA was controlled by these groups. As a result, the investors and founders retained substantial control over the DeFi project from the very beginning, preventing it from ever achieving true decentralization. This persistent concentration of power fundamentally undercuts the principle of **decentralized authority**.

In an effort to counteract centralization, some projects have experimented with “fair launch” token distributions, aiming to distribute initial governance token supplies more evenly among participants [RP2-\[11\]](#). The term “fair launch” serves as an umbrella concept, signifying a project’s intention to allocate tokens as equitably as possible. However, it does not correspond to a single, standardized method of distribution. For example, in the case of Yearn.Finance, the term referred to governance token rewards distributed in exchange for liquidity provision in certain DEXes [\[103\]](#). More recently, in the

case of ORE, the fair launch approach is tied to a form of crypto mining, where tokens are mined based on the computational power participants are able to allocate [63].

However, RP2-[11] and RP3-[10] suggest that the fair launch approach of Yearn.Finance merely delays, rather than eliminates, the reemergence of voting power concentration. To explore whether any fair launch scenario could effectively reduce or alleviate centralization in governance token distribution, RP5-[27] employed an agent-based model (ABM) to simulate theoretical scenarios where the initial distribution of tokens is perfectly even among users. The results revealed that even under these idealized conditions, centralization inevitably reemerges over time. The only observable benefit of a perfectly even initial distribution was a delay in the onset of centralization. Overall, the findings suggest that, in the long run, fair launch strategies are unable to prevent centralization. Any benefits they provide are short-lived and do not fundamentally address the underlying dynamics that lead to concentrated governance power.

Voter apathy further exacerbates these challenges, as many token holders neither engage in discussions nor cast ballots. As a result, the pool of active voters shrinks over time, amplifying the influence of the few who do participate and further eroding the principle of **decentralized authority**. In RP2-[11], an analysis of voter apathy during the voting stage of DAOs revealed that voter participation consistently diminishes over time in all examined DeFi protocols. In most cases, just a few months after a project's launch, participation rates drop to less than 1%, with some projects experiencing rates well below 0.1% (Figure II.3).




Figure II.3: Governance participation rates. Each point corresponds to a proposal (c.f. RP2-[11]).

2.3. Execution

Although smart contracts theoretically enable **smart contract-based autonomy** by automating proposal outcomes, flawed or malicious code can cause irreparable harm to a DAO [26]. A notorious example of such risks is the early *Yam* protocol. An error in the protocol's execution code led to the unintended minting of a large number of governance tokens, causing significant inflation in the token supply. As a consequence, subsequent proposals could no longer achieve quorum, rendering the voting process ineffective and halting the implementation of any changes. This left the *Yam* protocol in a state of perpetual stagnation, effectively destroying the DAO [102]. While extreme mishaps like this are less likely today thanks to improved development practices and code audits, the *Yam* incident underscores how code-level mistakes can swiftly and irrevocably undermine **decentralized authority** by placing disproportionate control in the hands of a flawed contract.

A parallel concern arises when DAOs rely on a developer team to implement the final stage, effectively shifting the last part of execution away from automated processes. In such cases, **smart contract-based autonomy** takes a back seat to human oversight, and DAO participants must trust that the developers have neither altered the code nor deviated from the proposal's original intent [74]. While this approach may reduce the immediate risk of catastrophic code errors, it undercuts the principles of **decentralized authority** and **transparent processes** by granting a select group the power to finalize, edit, or delay proposed changes—moving control away from the broader community.

An example of the risks associated with this approach is the case of Aragon DAO. Token holders criticized the Aragon Association for underutilizing its treasury and failing to achieve its stated goals. In response, they proposed transferring control of the treasury to token governance and initiating a buyback program to redistribute value to token holders. Although the DAO voted in favor of decentralizing treasury management, the Aragon Association refused to implement the proposed changes. Instead, when faced with mounting pressure from token holders, the Association dismissed the proposals as a "coordinated attack." Furthermore, it banned dissenting voices from its communication platforms and retained control over the treasury [8].

III Governance in Permissioned Blockchains

1 Core Principles & Governance Process

In contrast to the open, permissionless nature of DAOs designed for public blockchains, permissioned blockchains operate under more stringent membership controls [79, 104]. That is, only authorized entities are allowed to participate in the network. These networks are often deployed in government, supply chain, or enterprise contexts, where organizations are required to preserve legal mandates and regulatory compliance[7, 98]. In addition, permissioned blockchains typically offer advantages such as superior performance compared to permissionless blockchains [48, 89], easier issue identification and adjustment RP1-[65], and enhanced interoperability with existing systems. Though their architectures differ substantially from DAOs, permissioned blockchains still pursue similar yet different end goals: secure collaboration, streamlined decision-making, and enhanced data integrity among multiple stakeholders [98]. In practice, the governance of permissioned blockchains revolves around clearly defined roles, detailed membership criteria, and a balance of on-chain and off-chain procedures to ensure compliance with consortium-level agreements [7].

Drawing on insights from blockchain consortia in the private sector [70, 104], as well as the public sector [79], four interrelated governance principles can be identified:

- **Selective Access Control:** In permissioned blockchains, each node is operated only by approved participants, who may hold legal or contractual obligations to the consortium. This “permission layer” maintains a strict separation of respon-

sibilities [79]. The core idea is to retain benefits of a distributed ledger—such as immutability and shared data—while limiting access to those that meet explicit compliance requirements.

- **Role-Based Transparency:** Instead of having fully public transaction logs, permissioned networks often implement granular, role-based views of the ledger [7]. Participants see only the data relevant to their legally mandated tasks [104]. This approach preserves the principle of privacy-by-design in contexts like regulated industries or federal agencies, aligning with compliance norms such as the GDPR [79].
- **Formalized Consortium Governance:** Unlike permissionless DAOs, where token holders can be mostly unknown, permissioned blockchains typically operate under *off-chain* legal agreements or memoranda of understanding that specify the rights and obligations of each consortium member [7, 98]. This formal structure, often overseen by a steering committee or governance board, is critical for multi-stakeholder collaboration in contexts such as supply chains or federal agencies.
- **Adjustable On-Chain Logic:** Although many operational rules remain off-chain, permissioned blockchains still use smart contracts to automate certain cross-organizational processes and reduce administrative overhead [66, 90]. Crucially, these smart contracts can be selectively upgraded to adapt to regulatory shifts or new business requirements.

In principle, the governance of a permissioned blockchain unfolds in three interlocked stages: Admission, which defines who is allowed to join the network; Decision-Making and Voting, where updates to smart contracts or policy are proposed and voted on by consortium members; and Execution, wherein approved proposals are implemented and continuously evaluated [79, 104].

1.1. Admission

Admission processes in permissioned blockchains aim to ensure that new participants meet necessary technical, regulatory, and business criteria. In many networks, prospective members must sign a consortium agreement or be vetted by existing members [98].

This approach stands in contrast to the **open participation** of public DAOs, where anyone can join. Here, decision-making authority is grounded in *organization-level identity* [7].

This **selective access control** fosters compliance with legal or corporate mandates. For example, a consortium might grant partial node privileges to a subcontractor whose role is strictly limited to verifying certain data. Additionally, role-based access can ensure each stakeholder only sees data germane to their function, aligning with the principle of “separation of competencies” in federally structured systems [9].

1.2. Decision-Making & Voting

After the admission stage, permissioned blockchain members collectively propose, debate, and vote on system upgrades or changes to smart contracts [79]. Yet, unlike DAOs, the forum for these discussions is often an *off-chain* consortium meeting or steering committee, which then coordinates *on-chain* votes. Typically, one or more committees composed of authorized representatives from each participating organization hold significant sway—mirroring the hierarchical structure of the real-world organizations [104].

While permissioned blockchains generally use software-based “voting modules” to finalize decisions, the threshold of approval and the distribution of voting rights are *predefined* in the consortium contract [7]. Some consortiums adopt “weighted voting,” where an organization’s stake or size determines its influence. Others opt for equal voting for all nodes, reflecting a “federal” spirit of shared governance [9]. To maintain **role-based transparency**, meeting minutes, proposals, and final ballots may be added as hashes to the blockchain, ensuring tamper-resistance and establishing an auditable record [90].

1.3. Execution

Approved proposals transition to an implementation phase, where designated developers or authorized maintainers update smart contracts or modify node software [79]. At this stage, **formalized consortium governance** may require *two-factor* or *multi-signature* validations from multiple participant organizations [98]. The code changes can then be

rolled out to all network nodes, ensuring consistency and *avoiding single points of failure* [66].

The guiding principle of **adjustable on-chain logic** allows these implementations to be updated or reversed if they conflict with new regulations or if security vulnerabilities emerge [7]. Finally, many consortia have permanent working groups or compliance teams tasked with monitoring the network, generating audits, and ensuring that real-world laws or new internal policies are reflected on-chain.

2 Practical Challenges

While the above principles and processes illustrate how permissioned blockchains are intended to function, real-world deployments often encounter obstacles.

2.1. Admission

Despite strict membership screening being central to **selective access control**, real consortia can find it hard to define uniform admission rules that satisfy each participant's legal and business requirements. Some organizations (e.g., a federal office) may be mandated by law to keep certain data fully siloed, making them hesitant to join a shared ledger [7]. Others, particularly large corporations, may demand special privileges such as the ability to veto proposals that threaten their competitive advantage. This fragmentation leads to multi-tier membership regimes, where “core” members enjoy expanded rights while “associate” members are granted minimal privileges [79]. In theory, this structure upholds *federalist* ideals of mutual respect for distinct roles, but in practice can become a quasi-hierarchical arrangement where certain members can reject new applications unilaterally [9, 79].

Moreover, in federally organized government contexts, local laws or data-protection statutes can prohibit certain agencies from sharing data outside their jurisdiction [9]. While permissioned blockchains claim to mitigate these concerns through robust role-based privacy, the administrative overhead to ensure each jurisdiction's compliance is steep, creating a labyrinth of admission protocols [7].

2.2. Decision-Making & Voting

Off-chain power dynamics can undercut nominally “egalitarian” voting. Although the consortium contract may stipulate a balanced allocation of votes, larger or more influential members can exert sway through financial leverage or political agreements [9]. In effect, a few lead stakeholders guide strategic choices, diminishing the principle of **formalized consortium governance** based on collective decision-making.

Compounding this, “one node, one vote” rarely reflects real-world realities. For instance, if a state-level agency invests considerable resources in the network, it may argue for a stronger vote. Alternatively, a smaller municipality might fear overshadowing and be reluctant to adopt the system [104]. Formalizing these power imbalances in an off-chain contract can help, but as a result, *dominant nodes* effectively set the agenda [79]. The hierarchical influence that creeps in mirrors the “whale problem” in DAOs, although it manifests through legal or bureaucratic means rather than token holdings.

2.3. Execution

While permissioned blockchains do **automate** certain cross-organizational tasks, many critical updates rely heavily on human gatekeepers—committees, consortium boards, or developer teams chosen off-chain [98]. This partial reliance on human intervention can introduce single points of failure, undermining the principle of distributed decision-making [7]. Malicious or incompetent gatekeepers might push through code that favors one member or fails to comply with specific laws, leaving other participants little recourse but to litigate off-chain.

Technical pitfalls can also arise if the code in the chain’s smart contracts is complex or poorly audited, vulnerabilities can go unnoticed until exploited [90]. Because many consortia integrate smart contracts with legacy systems, updating the chain might break existing workflows or violate local data-sharing protocols. Overextended multi-signature processes, while preventing unilateral changes, can also slow down urgent patches—especially in large consortia with many signers.

Lastly, privacy compliance adds another layer of complication. Members with distinct legal obligations must verify that on-chain data is either anonymized or accessible only to authorized participants [79]. If the privacy logic is not meticulously encoded into

each layer of smart contracts, data can unintentionally be exposed to unauthorized parties.

IV | Innovations in Governance Mechanisms

Many of the shortcomings of blockchain governance are widely recognized by both practitioners and scholars. Over the years, numerous mechanisms have been proposed to address persistent issues. Although these mechanisms have achieved varying levels of success, they collectively underscore a broader effort to reconcile the original ideals with the operational realities of blockchain projects. The following sections focus on three prominent mechanisms, focusing on how each seeks to resolve specific weaknesses in blockchain governance.

1 Quadratic Voting

One of the earliest proposed solutions to address the concentration of power in token-based governance is quadratic voting [68, 69]. Unlike the conventional “one token, one vote” mechanism, quadratic voting allocates voting power based on the square root of the number of tokens held. This diminishing returns approach limits the influence of large token holders. For instance, a user holding 100 governance tokens would receive only 10 votes under quadratic voting, compared to the 100 votes granted in a traditional system. By reducing the disproportionate impact of “whale” addresses, quadratic voting aims to strengthen the principle of **decentralized authority**. A notable example of its implementation is Synthetix [RP2-\[11\]](#), a decentralized derivatives platform that leverages quadratic voting for protocol upgrades and parameter changes.

Despite its potential to mitigate power imbalances, quadratic voting is not without its challenges. Probably the most critical vulnerability is its lack of Sybil resistance [RP2-](#)

[11]. Because the system allocates votes based on individual addresses, adversaries can exploit it by splitting their token holdings across multiple wallets. For example, a user holding 100 tokens could distribute them across 10 wallets, each containing 10 tokens. Under quadratic voting, each wallet would receive 3.16 votes, granting the adversary a total of approximately 31.6 votes. This manipulation—commonly referred to as a Sybil attack—undermines the system’s intent to limit the influence of large stakeholders.

Quadratic voting offers a compelling framework for more balanced governance, but in practice, it introduces new challenges even as it seeks to address existing issues. This underscores the persistent tension between theoretical solutions and their practical implementation in the pursuit of decentralized governance.

2 NFT-Based Voting

An alternative governance model has been explored in the past [24, 62], leveraging NFTs rather than solely relying on fungible tokens. In this framework, each eligible participant holds a unique NFT that grants specific voting rights and serves as a dynamic reputation system, capturing user actions and engagement. By anchoring votes to distinct NFTs, the model aims to go beyond the simplistic logic that more tokens necessarily equate to more votes, instead placing value on trust, accountability, and community involvement.

NFT-based voting bolsters **decentralized authority** by assigning each participant a unique identity, potentially mitigating the dominance of “whales” who accumulate large amounts of fungible tokens. Additionally, since these NFTs can reflect a user’s contributions or specific roles, the model supports **open participation** by recognizing forms of engagement overlooked in purely token-based systems. In principle, this flexible structure allows for weighted voting rules that factor in variables such as seniority, expertise, or reputation—rather than relying exclusively on financial stake.

Despite these advantages, NFT-based voting introduces its own set of challenges. RP4-[35] indicates that herding behavior is especially pronounced in NFT markets, leading to sharp inequalities and a concentration of market influence. This tendency raises concerns that NFT-based voting, if not carefully structured, might replicate the same concentration dynamics it aims to address. For instance, if governance NFTs are freely

transferable, they can be traded on secondary markets, allowing well-funded actors to replicate the concentration issues seen in fungible token models. Conversely, making them non-transferable (often called “soulbound” [92]) may strengthen Sybil resistance and avoid vote-buying but raises broader questions about user autonomy—particularly how to handle membership exits, revocations, or status changes. Even soulbound NFTs do not fully eliminate the threat of multiple identities, as adversaries can gradually build reputations across different addresses. Additionally, because the identity of the NFT owner often remains unknown, adversaries may exploit a “one-off” opportunity by building up trust over time, and at the end cast a harmful vote that personally benefits them, and then vanishing with minimal accountability. Consequently, the model still grapples with Sybil attacks and fair distribution, undercutting the ideals of **decentralized authority** and **open participation** in practice.

In theory, NFT-based governance offers a richer, context-driven approach to decision-making by incorporating more granular measures of trust and contribution. However, finding the right balance between transferability, Sybil resistance, and member autonomy remains a complex endeavor.

3 Market-Based Voting

Decentralized communities also explore *market-based* voting systems as an alternative to traditional token-based governance. These experiments often draw upon the concept of *futarchy* [45], which proposes that decisions should be made by prediction markets—functioning similarly to betting markets—rather than through direct majority votes. The primary objective of futarchy is to align stakeholder incentives with accurate forecasting where by speculating on the outcomes of proposed policies, participants are theoretically encouraged to reveal which course of action they genuinely believe will best serve the organization’s goals.

In non-blockchain contexts, futarchy has been proposed as a means to enhance decision-making in corporate governance and public policy [45] by leveraging market forces to aggregate diverse information. Proponents argue that when individuals bet on policy outcomes, rather than simply casting votes based on their personal beliefs, they gain a financial stake in the accuracy of their forecasts. According to futarchy advocates, this

setup provides a more dependable indicator of authentic preferences because individuals have a direct economic incentive to choose—and profit from—the most advantageous proposals.

For blockchain-based organizations, futarchy’s appeal lies in its potential to replace the rigid “one token, one vote” paradigm with a more flexible mechanism that rewards participants who actively back their convictions [21, 60]. Rather than tying voting power to the tokens one already holds, futarchy relies on how many tokens participants are willing to buy or sell in support of specific outcomes. If a participant believes strongly in a particular proposal, they can place larger bets on it, thereby increasing their influence on the eventual decision.

In principle, futarchy upholds the principle of **decentralized authority** by favoring informed, proactive engagement over passive token ownership. Rather than defaulting to the largest token holders, any stakeholder who is confident in a proposal’s merits can buy tokens on its likely success, thereby amplifying their governance influence. This design fosters a more meritocratic model of decision-making in which those who thoroughly assess policies gain proportional influence, instead of individuals who merely hold significant token balances. Moreover, since anyone can join the prediction market and bet according to their convictions, futarchy ostensibly supports the principle of **open participation**.

In reality, however, futarchy can reintroduce many of the same centralization issues it aims to counteract. Wealthy users may place disproportionately large bets, effectively sidelining smaller participants and replicating the power imbalances seen in other voting mechanisms. Additionally, complex market factors—such as speculative bubbles and price manipulation—can distort market signals, rendering final outcomes less representative of genuine community consensus. The need for liquidity also poses a barrier to entry, as individuals lacking sufficient capital often hesitate to participate, thus narrowing the scope of **open participation**. Arbitrageurs may further exploit price differences between the prediction market and the underlying token market purely for profit, contradicting the premise that futarchy participants prioritize accurate forecasts. Finally, prediction markets are vulnerable to strategic manipulation—such as coordinated betting—raising additional questions about the extent to which futarchy fosters a truly fair and decentralized governance.

V | Conclusion

1 Summary

The thesis set out to scrutinize the discrepancy between the theoretical aspirations of blockchain governance and the practical realities in which these systems operate, examining both permissionless and permissioned contexts. In permissionless settings, DAOs were hailed as the epitome of decentralized decision-making, designed to uphold core principles like open participation, transparent processes, and decentralized authority. In theory, self-executing smart contracts would reduce the reliance on centralized intermediaries, and token-based voting would empower token-holders to collaboratively shape the trajectory of decentralized protocols. Yet, numerous cases reveal the persistence and reinforcement of hierarchical power dynamics, where a small cluster of influential token holders, wields disproportionate influence. Voter apathy adds to this concentration by allowing a small minority of engaged actors to steer protocol outcomes unilaterally. Beyond power concentration, practical execution challenges have arisen, including code vulnerabilities that can immobilize entire projects. Off-chain coordination and human oversight, meant to mitigate these risks, paradoxically reintroduce forms of centralized control.

In permissioned blockchains, governance mechanisms were found to be more formalized, reflecting real-world organizational structures and regulatory constraints. Their governance processes rely on contracts, committees, and hierarchical decision-making bodies that incorporate legal and business mandates, thus trading off some decentralization to maintain compliance, privacy, and accountability. Although permissioned networks do simplify the resolution of disputes and provide legal recourse, they repli-

cate many of the oligarchic patterns seen in permissionless systems. Prominent members—such as large corporations or government agencies—often hold veto power, while smaller partners or new entrants are relegated to less influential roles. Consequently, despite employing distributed ledger technology, many consortia revert to traditional top-down hierarchies.

The gaps between theoretical ideals and operational realities have prompted a range of innovations. Quadratic voting, NFT-based voting, and futarchy-based approaches each offer ways to address power imbalances and enhance the quality of collective decisions. Yet these approaches come with their own set of vulnerabilities. Collectively, these solutions illustrate the unresolved tension between limiting undue influence and preserving open participation, as well as the difficulty of simultaneously achieving both efficiency and inclusiveness.

Overall, the thesis finds that while blockchains and DAOs bring novel governance possibilities, they have yet to fully deliver on their ideals. Real-world applications consistently run into issues of power asymmetry, voter disinterest, technical fragility, and re-centralization. Although these challenges do not negate the potential of blockchain technology, they underscore the pressing need for more nuanced, context-specific governance designs that can adapt to varying stakeholder requirements.

2 Limitations and outlook

Although this work offers a critique of how blockchain governance falls short of its theoretical aspirations, it centers predominantly on instances where implementation problems prevent the realization of these ideals. An equally important—but less explored—avenue lies in evaluating whether the prevailing governance designs themselves are fit for purpose. In other words, beyond finding that systems fail to meet ideals, one must also examine whether those ideals and the corresponding designs are aligned with the diverse operational needs of different protocols.

The current one-size-fits-all approach to DAOs governance attempts to impose universal governance model for every conceivable scenario, treating all proposals equally regardless of their complexity, urgency, or risk profile. This slow-moving decision-making processes may be well suited for substantial strategic changes or treasury allocations,

but can be ill-suited for situations demanding rapid responses. For instance, many DeFi platforms must tweak interest rates or implement security updates on short notice to protect user assets. Forcing every decision through a slow governance cycle, in which discussions and voting can take days if not weeks, can leave protocols exposed to hacks or market volatilities. Similarly, short-term operational tasks rarely necessitate the same level of deliberation as major policy shifts. The monolithic governance model often applied in DAOs fails to differentiate between such scenarios, leading to substantial delays.

Another factor that warrants more scrutiny is the reinforcing tension between timocratic decision-making (as identified in [RP2-\[11\]](#)) and the reality of execution. While voting outcomes may reflect the preferences of large token holders, practical implementation of these decisions frequently resides with a small team of core developers—who function, in effect, as (benevolent) dictators. Users must trust these developers to adhere faithfully to governance decisions, raising questions about how “decentralized” these systems truly are. While the trust in experienced developers can foster network stability and ensure timely upgrades, it conflicts with the premise that permissionless governance should bypass centralized custodians of authority.

In sum, these limitations highlight that the thesis tackled only part of the puzzle by focusing on implementation pitfalls. The broader question of whether current designs are conceptually sound for all use cases remains open. Addressing this question will likely involve devising more adaptive governance models that can align the necessary speed and expertise for specific proposals with the principles that underpin the broader blockchain ethos. Without such a deeper analysis of design principles, blockchain governance risks stagnating in a realm where both its ideals and real-world practices remain perpetually out of sync.

3 Recognition of previous and related work

Research thrives through collaboration, drawing on shared insights and exchanges among experts in various fields. Reflecting this spirit, this thesis builds on a broad body of work to analyze blockchain governance and decentralized decision-making. The inherent complexity of blockchain technology necessitates a holistic approach, examining

both its foundational concepts and practical applications to fully grasp its opportunities and constraints. While the core of this work relies on research specifically focused on blockchain governance, it also draws from a multitude of studies that address the broader blockchain landscape, ensuring a comprehensive understanding of this evolving domain.

Specifically, this work relies on a set of works that focuses on the fundamentals of blockchain technology, smart contracts, tokenization, and DeFi. For the technical aspect of the blockchain technology, the thesis relies on [3, 38, 39, 82, 86], for the influence of smart contracts on blockchain-based consortia and workflows on [4, 5, 36, 40, 75], and for the evolving landscape of token economies, ICOs, and DeFi protocols on [12, 13, 15, 16, 42, 51, 78, 84, 88, 93]. Collectively, these works offer a foundation for understanding the intersections of decentralized architectures and shifting power structures, which are core ideas that shape this thesis.

Additionally, the thesis leverages contributions that examine the implications of blockchain technology within public-sector and identity-driven contexts. These include studies on blockchain-enabled cross-organizational processes, governmental services, and GDPR-related matters [6, 28, 29, 31, 32, 33, 34, 37, 43, 49, 50, 53, 58, 67, 76]. While they largely address specialized use cases, these references collectively demonstrate how blockchain's governance features, data protection requirements, and stakeholder engagement practices permeate a variety of sectors, further underscoring the broader relevance of decentralized decision-making explored in this thesis.

VI | References

- [1] *Aave Protocol Overview*. URL: <https://docs-aave.vercel.app/docs/> (visited on 12/15/2024).
- [2] M. Abdelhamid and G. Hassan. “Blockchain and Smart Contracts”. In: ICSIE ’19. New York, NY, USA: Association for Computing Machinery, 2019, 91–95. DOI: 10.1145/3328833.3328857.
- [3] I. Abellan Alvarez, V. Gramlich, and J. Sedlmeir. “Unsealing the secrets of blockchain consensus: A systematic comparison of the formal security of proof-of-work and proof-of-stake”. English. In: Ávila, Spain, 2024. DOI: 10.1145/3605098.3635970.
- [4] S. Albrecht, S. Reichert, J. Schmid, J. Strüker, D. Neumann, and G. Fridgen. “Dynamics of Blockchain Implementation : A Case Study from the Energy Sector”. English. In: *51st Hawaii International Conference on System Sciences*. 2018. URL: <https://eref.uni-bayreuth.de/49870/>.
- [5] J. Amend, C. van Dun, G. Fridgen, F. Köhler, A. Rieger, A. Stohr, and A. Weninger. “Using Blockchain to Coordinate Federal Processes: The Case of Germany’s Federal Office for Migration and Refugees”. English. In: *Digitalization Cases Vol. 2: Mastering Digital Transformation for Global Business*. Springer International Publishing, 2021. DOI: 10.1007/978-3-030-80003-1_5.
- [6] J. Amend, M. Federbusch, G. Fridgen, F. Köhler, A. Rieger, V. Schlatt, J. Sedlmeir, A. Stohr, and C. van Dun. *Digitization of certification processes in the asylum procedure by means of digital identities: A feasibility study by Germany’s Federal Office for Migration and Refugees*. English. Tech. rep. 2021. URL: <https://www.bamf.de>

- /SharedDocs/Anlagen/EN/Digitalisierung/blockchain-whitepaper-2021.pdf?__blob=publicationFile&v=2.
- [7] J. Amend, G. Fridgen, A. Rieger, T. Roth, and A. Stohr. “The Evolution of an Architectural Paradigm - Using Blockchain to Build a Cross-Organizational Enterprise Service Bus”. In: *Hawaii International Conference on System Sciences*. 2021. URL: <https://scholarspace.manoa.hawaii.edu/items/dcbcd94a-2ac4-48f8-9a85-7c36c6046946>.
- [8] *Aragon’s \$200m activist battle ignites DAO debate: ‘Lining up soldiers doesn’t mean an attack’*. URL: <https://www.dlnews.com/articles/defi/aragon-dao-activist-highlights-peril-of-community-control/> (visited on 12/05/2024).
- [9] A. Auer. “The constitutional scheme of federalism”. In: *Journal of European Public Policy* 12.3 (2005), pp. 419–431. DOI: 10.1080/13501760500091166.
- [10] T. Barbereau, R. Smethurst, O. Papageorgiou, A. Rieger, and G. Fridgen. “DeFi, Not So Decentralized: The Measured Distribution of Voting Rights”. In: *Proceedings of the 55th Hawaii International Conference on System Sciences*. AIS, 2022, pp. 6043–6052. DOI: 10.24251/HICSS.2022.734.
- [11] T. Barbereau, R. Smethurst, O. Papageorgiou, J. Sedlmeir, and G. Fridgen. “Decentralised Finance’s timocratic governance: The distribution and exercise of tokenised voting rights”. In: *Technology in Society* 73 (2023), p. 102251. DOI: <https://doi.org/10.1016/j.techsoc.2023.102251>.
- [12] T. J. Barbereau and B. Bodó. “Beyond Financial Regulation of Crypto-asset Wallet Software: In Search of Secondary Liability”. English. In: *Computer Law Security Review* 49 (2023). DOI: 10.1016/j.clsr.2023.105829.
- [13] T. J. Barbereau, R. Smethurst, J. Sedlmeir, G. Fridgen, and A. Rieger. “Chapter 8: Tokenization and Regulatory Compliance for Art and Collectible Markets: From Regulators’ Demands for Transparency to Investors’ Demands for Privacy”. English. In: *Blockchains and the Token Economy: Studies in Theory and Practice*. Palgrave Macmillan, 2022. DOI: 10.1007/978-3-030-95108-5_8.
- [14] *Beginners’ Guide to Voting | Uniswap*. URL: <https://docs.uniswap.org/concepts/governance/guide-to-voting> (visited on 12/05/2024).

- [15] M. Brennecke, T. Guggenberger, A. SACHS, and B. Schellinger. "The Human Factor in Blockchain Ecosystems: A Sociotechnical Framework". English. In: Nuremberg, Germany, 2022. URL: https://aisel.aisnet.org/wi2022/finance_and_blockchain/finance_and_blockchain/3/.
- [16] M. Brennecke, T. Guggenberger, B. Schellinger, and N. Urbach. "The De-Central Bank in Decentralized Finance: A Case Study of MakerDAO". English. In: Virtual, United States - Hawaii, 2022. DOI: 10.24251/HICSS.2022.737.
- [17] V. Buterin. *Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform*. 2014. URL: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- [18] *DefiLlama - DeFi Dashboard*. URL: <https://defillama.com/> (visited on 12/05/2024).
- [19] *Dexes TVL Rankings*. URL: <https://defillama.com/protocols/Dexes> (visited on 12/05/2024).
- [20] V. Dhillon, D. Metcalf, and M. Hooper. "The DAO Hacked". In: *Blockchain Enabled Applications: Understand the Blockchain Ecosystem and How to Make it Work for You*. Apress, 2017, pp. 67–78. DOI: 10.1007/978-1-4842-3081-7_6.
- [21] W. W. Ding, X. Liang, J. Hou, G. Wang, Y. Yuan, J. Li, and F.-Y. Wang. "Parallel Governance for Decentralized Autonomous Organizations enabled by Blockchain and Smart Contracts". In: *2021 IEEE 1st International Conference on Digital Twins and Parallel Intelligence (DTPI)*. 2021, pp. 1–4. DOI: 10.1109/DTPI52967.2021.9540069.
- [22] Y. El Faqir, J. Arroyo, and S. Hassan. "An overview of decentralized autonomous organizations on the blockchain". In: *Proceedings of the 16th International Symposium on Open Collaboration*. OpenSym '20. Association for Computing Machinery, 2020. DOI: 10.1145/3412569.3412579.
- [23] Electric Capital. *Developer Report*. 2024. URL: <https://www.developerreport.com/reports/devs/2024?s=developer-report>.

- [24] M. Eltuhami, M. Abdullahe, B. A. Talip, and N. A. N. Ahmad. “Blockchain Based Voting Framework With Non-Transferable Non-Fungible Tokens”. In: *2023 International Conference on Information Technology (ICIT)*. 2023, pp. 378–383. DOI: 10.1109/ICIT58056.2023.10225885.
- [25] Y. Fan, L. Zhang, R. Wang, and M. A. Imran. “Insight Into Voting in DAOs: Conceptual Analysis and a Proposal for Evaluation Framework”. In: *IEEE Network* 38.3 (2024), pp. 92–99. DOI: 10.1109/MNET.137.2200561.
- [26] R. Feichtinger, R. Fritsch, L. Heimbach, Y. Vonlanthen, and R. Wattenhofer. *SoK: Attacks on DAOs*. 2024. arXiv: 2406.15071 [cs.CR]. URL: <https://arxiv.org/abs/2406.15071>.
- [27] J. D. Fernandez, T. Barbereau, and O. Papageorgiou. “Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch”. In: *ACM Trans. Manage. Inf. Syst.* 15.1 (2024). URL: <https://doi.org/10.1145/3649318>.
- [28] G. Fridgen, N. Guggenberger, T. Hoeren, W. Prinz, N. Urbach, J. Baur, H. Brockmeyer, W. Gräther, E. Rabovskaja, V. Schlatt, A. Schweizer, J. Sedlmeir, and L. Wederhake. *Chancen und Herausforderungen von DLT (Blockchain) in Mobilität und Logistik*. English. Federal Ministry of Transport and Digital Infrastructure, 2019. URL: <https://eref.uni-bayreuth.de/55550/>.
- [29] G. Fridgen, N. Guggenberger, T. Hoeren, W. Prinz, N. Urbach, J. Baur, H. Brockmeyer, W. Gräther, E. Rabovskaja, V. Schlatt, A. Schweizer, J. Sedlmeir, and L. Wederhake. *Opportunities and Challenges of DLT (Blockchain) in Mobility and Logistics*. English. Tech. rep. 2019. URL: <https://eref.uni-bayreuth.de/55481/>.
- [30] G. Fridgen, F. Guggenmos, J. Lockl, A. Rieger, and A. Schweizer. “Developing an Evaluation Framework for Blockchain in the Public Sector: The Example of the German Asylum Process”. In: *Proceedings of 1st ERCIM Blockchain Workshop 2018*. European Society for Socially Embedded Technologies (EUSSET), 2018. DOI: 10.18420/blockchain2018_10.
- [31] G. Fridgen, F. Guggenmos, J. Lockl, and A. Rieger. “Challenges and Opportunities of Blockchain-based Platformization of Digital Identities in the Public Sector”. English. In: *Workshop on Platformization in the Public Sector, 26th European*

- Conference on Information Systems (ECIS)*. 2018. URL: <https://eref.uni-bayreuth.de/44302/>.
- [32] G. Fridgen, F. Guggenmos, J. Lockl, A. Rieger, A. E. Schweizer, and N. Urbach. "Developing an Evaluation Framework for Blockchain in the Public Sector : The Example of the German Asylum Process". English. In: *1st ERCIM Blockchain Workshop 2018*. 2018. URL: <https://eref.uni-bayreuth.de/44347/>.
- [33] G. Fridgen, F. Guggenmos, J. Lockl, A. Rieger, and N. Urbach. *Support of communication and cooperation in the asylum process with the help of blockchain: A feasibility study by the Federal Office for Migration and Refugees*. English. Tech. rep. 2018. URL: <https://eref.uni-bayreuth.de/46480/>.
- [34] G. Fridgen, F. Guggenmos, J. Lockl, A. Rieger, and N. Urbach. *Supporting communication and cooperation in the asylum procedure with Blockchain technology : A proof of concept by the Federal Office for Migration and Refugees*. English. Tech. rep. 2019. URL: <https://eref.uni-bayreuth.de/48037/>.
- [35] G. Fridgen, R. Kräussl, O. Papageorgiou, and A. Tugnetti. "Pricing dynamics and herding behaviour of NFTs". In: *European Financial Management* (). DOI: <https://doi.org/10.1111/eufm.12506>.
- [36] G. Fridgen, M.-F. Körner, J. Sedlmeir, and M. Weibelzahl. "(How) Can Blockchain Contribute to the Management of Systemic Risks in Global Supply Networks?" English. In: *14. Internationale Tagung Wirtschaftsinformatik*. 2019. URL: <https://eref.uni-bayreuth.de/51932/>.
- [37] G. Fridgen, J. Lockl, S. Radszuwill, A. Rieger, A. Schweizer, and N. Urbach. "A Solution in Search of a Problem : A Method for the Development of Blockchain Use Cases". English. In: *24th Americas Conference on Information Systems (AMCIS)*. 2018. URL: <https://eref.uni-bayreuth.de/44279/>.
- [38] G. Fridgen, W. Prinz, T. Rose, and N. Urbach. "Blockchain Lab : Design, Implementation and Evaluation of Innovative Business and Process Models". English. In: *ERCIM News 110* (2017). URL: <https://eref.uni-bayreuth.de/38956/>.
- [39] G. Fridgen, S. Radszuwill, A. E. Schweizer, and N. Urbach. "Development of Disruptive Innovations with Blockchain: The Way to the Right Use Case". In:

- Wirtschaftsinformatik Management* (2017). URL: <https://link.springer.com/article/10.1007/s35764-017-0108-0>.
- [40] G. Fridgen, S. Radszuwill, A. Schweizer, and N. Urbach. "Blockchain Won't Kill the Banks: Why Disintermediation Doesn't Work in International Trade Finance". English. In: *Communications of the Association for Information Systems* 49 (2021). DOI: 10.17705/1CAIS.04932.
- [41] G. Fridgen, S. Radszuwill, N. Urbach, and L. Utz. "Cross-Organizational Workflow Management Using Blockchain Technology - Towards Applicability, Auditability, and Automation". In: *Hawaii International Conference on System Sciences*. 2018. URL: <https://www.fim-rc.de/Paperbibliothek/Veroeffentlicht/696/wi-696.pdf>.
- [42] G. Fridgen, F. Regner, A. Schweizer, and N. Urbach. "Don't Slip on the Initial Coin Offering (ICO) : A Taxonomy for a Blockchain-enabled Form of Crowdfunding". English. In: *26th European Conference on Information Systems (ECIS)*. 2018.
- [43] G. Fridgen, J. Röhlen, T. Guggenberger, V. Schlatt, and M. Schulze. *Überprüfung der Machbarkeit eines offenen und dezentralen Mobilitätssystems (OMOS)*. English. Tech. rep. 2019. URL: <https://eref.uni-bayreuth.de/55752/>.
- [44] R. Fritsch, M. Müller, and R. Wattenhofer. "Analyzing voting power in decentralized governance: Who controls DAOs?" In: *Blockchain: Research and Applications* 5.3 (2024), p. 100208. DOI: <https://doi.org/10.1016/j.bcra.2024.100208>.
- [45] *Futarchy: Vote Values, But Bet Beliefs*. URL: <https://mason.gmu.edu/~rhanson/futarchy.html> (visited on 12/11/2024).
- [46] *Governance | DeFi Navigator*. URL: <https://ccaf.io/defi/protocol-deep-dive/governance> (visited on 01/10/2025).
- [47] T. Guggenberger, A. Schweizer, and N. Urbach. "Improving interorganizational information sharing for vendor managed inventory: Toward a decentralized information hub using blockchain technology". In: *IEEE Transactions on Engineering Management* 67.4 (2020), pp. 1074–1085.

- [48] T. Guggenberger, J. Sedlmeir, G. Fridgen, and A. Luckow. "An In-Depth Investigation of the Performance Characteristics of Hyperledger Fabric". In: *Computers & Industrial Engineering* 173 (2022), p. 108716. ISSN: 0360-8352. DOI: <https://doi.org/10.1016/j.cie.2022.108716>.
- [49] F. Guggenmos, J. Lockl, A. Rieger, and G. Fridgen. "Blockchain in der öffentlichen Verwaltung - Unterstützung der Zusammenarbeit im Asylprozess". English. In: *Informatik-Spektrum* (2019). ISSN: 0170-6012. DOI: 10.1007/s00287-019-01177-y.
- [50] F. Guggenmos, J. Lockl, A. Rieger, and G. Fridgen. "How to Develop a GDPR-Compliant Blockchain Solution for Cross-Organizational Workflow Management: Evidence from the German Asylum Procedure." English. In: *Proceedings of the Hawaii International Conference on System Sciences 2020*. Maui, Hawaii, United States, 2020. DOI: 10.24251/HICSS.2020.492. URL: <http://hdl.handle.net/10125/64234>.
- [51] E. Hartwich, P. Ollig, G. Fridgen, and A. Rieger. "Probably something: a multi-layer taxonomy of non-fungible tokens". In: *Internet Research* 34.1 (2024), pp. 216–238. DOI: 10.1108/INTR-08-2022-0666.
- [52] A. Hoess, D. H. Maclennan, A. Rieger, E. Ermolaev, G. Fridgen, and T. Roth. *Issuing and verifying digital diplomas with the European Blockchain Services Infrastructure - Insights from the EBSILUX project*. 2022. URL: <https://orbilu.uni.lu/bitstream/10993/59132/1/EBSILUX%20Whitepaper.pdf>.
- [53] A. Höß, V. Schlatt, A. Rieger, and G. Fridgen. "The Blockchain Effect: From Inter-Ecosystem to Intra-Ecosystem Competition". English. In: *Proceedings of the Twenty-Ninth European Conference on Information Systems (ECIS)*. Association for Information Systems, 2021. URL: https://aisel.aisnet.org/ecis2021_rp/36.
- [54] R. Innovation. *DAOs Explained: Complete Guide to Decentralized Autonomous Organizations*. URL: https://www.rapidinnovation.io/post/daos-explained-ultimate-guide-to-decentralized-autonomous-organizations?utm_source=chatgpt.com (visited on 12/05/2024).

- [55] J. R. Jensen, V. von Wachter, and O. Ross. *How Decentralized is the Governance of Blockchain-based Finance: Empirical Evidence from four Governance Token Distributions*. 2021. arXiv: 2102.10096 [q-fin.GN]. URL: <https://arxiv.org/abs/2102.10096>.
- [56] O. Konashevych. "Why 'Permissioned' and 'Private' are not Blockchains". In: (2019). DOI: 10.2139/ssrn.3496468. URL: <https://ssrn.com/abstract=3496468>.
- [57] Y. Liu, Q. Lu, L. Zhu, H.-Y. Paik, and M. Staples. "A systematic literature review on blockchain governance". In: *Journal of Systems and Software* 197 (2023), p. 111576. DOI: <https://doi.org/10.1016/j.jss.2022.111576>.
- [58] J. Lockl, A. Rieger, G. Fridgen, M. Röglinger, and N. Urbach. "Towards a Theory of Decentral Digital Process Ecosystems : Evidence from the Case of Digital Identities". English. In: *Workshop Blockchain Research ? Beyond the Horizon*. 2018. URL: <https://eref.uni-bayreuth.de/44282/>.
- [59] J. Ma, M. Jiang, J. Jiang, X. Luo, Y. Hu, Y. Zhou, Q. Wang, and F. Zhang. *Demystifying the DAO Governance Process*. 2024. arXiv: 2403.11758 [cs.SE]. URL: <https://arxiv.org/abs/2403.11758>.
- [60] *MetaDAO - Futarchy Overview*. URL: <https://docs.metadao.fi/futarchy/overview> (visited on 12/11/2024).
- [61] S. Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. URL: bitcoin.org/bitcoin.pdf.
- [62] *NFT-Driven Governance*. URL: <https://members.delphidigital.io/learn/nft-driven-governance> (visited on 12/05/2024).
- [63] ORE [@ORESsupply]. Tweet. 2024. URL: <https://x.com/ORESsupply/status/1775176586154889653> (visited on 01/06/2025).
- [64] *Overview | Blur Foundation*. URL: <https://docs.blur.foundation> (visited on 12/15/2024).
- [65] O. Papageorgiou, L. Börtzler, E. Ermolaev, J. Kumari, and J. Sedlmeir. "What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains". In: Big Island, United States - Hawaii, 2025. URL: <https://hdl.handle.net/10125/109740>.

- [66] A. B. Pedersen, M. Risius, and R. Beck. “A Ten-Step Decision Path to Determine When to Use Blockchain Technologies”. In: *MIS Q. Executive* 18 (2019), p. 3.
- [67] N. Pocher. “The Internet of Money between Anonymity and Publicity: Legal Challenges of Distributed Ledger Technologies in the Crypto Financial Landscape”. English. In: *DC JURIX 2019. Proceedings of the Seventh JURIX 2019 Doctoral Consortium*. RWTH Aachen University, 2020. URL: <http://ceur-ws.org/Vol-2598>.
- [68] E. Posner and G. Weyl. “Quadratic Vote Buying as Efficient Corporate Governance”. In: *The University of Chicago Law Review* 81 (2013).
- [69] E. A. Posner and E. G. Weyl. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton University Press, 2018. DOI: 10.2307/j.ctvc77c4f.
- [70] A. Prager, K. Moudgal, and N. Sharma. *Project EPIC, Kinexys by J.P. Morgan*. 2024. URL: <https://www.jpmorgan.com/kinexys/documents/JPMC-Kinexys-Project-Epic-Whitepaper-2024.pdf>.
- [71] I. Prokopiiev. *DAOs and Multi-Signature Wallets*. 2024. URL: <https://www.prokopiievlaw.com/post/daos-and-multi-signature-wallets> (visited on 12/04/2024).
- [72] *Protect Your Users With Smart Contract Timelocks - OpenZeppelin blog*. URL: <https://blog.openzeppelin.com/protect-your-users-with-smart-contract-timelocks> (visited on 12/04/2024).
- [73] *Quorum’s Role in DAO Voting: Key to Inclusive Decisions*. 2024. URL: <https://blog.colony.io/quorum-in-dao-voting-ensuring-decisive-and-representative-decisions> (visited on 12/04/2024).
- [74] W. Reijers, I. Wuisman, M. Mannan, et al. “Now the Code Runs Itself: On-Chain and Off-Chain Governance of Blockchain Technologies”. In: *Topoi* 40 (2021), pp. 821–831. DOI: 10.1007/s11245-018-9626-5.
- [75] A. Rieger, F. Guggenmos, J. Lockl, G. Fridgen, and N. Urbach. “Building a Blockchain Application that Complies with the EU General Data Protection Regulation”. English. In: *MIS Quarterly Executive* 18.4 (2019). DOI: 10.17705/2msqe.00020.

- [76] A. Rieger, A. Stohr, A. Wenninger, and G. Fridgen. “Reconciling Blockchain with the GDPR: Insights from the German Asylum Procedure”. English. In: *Blockchain and the Public Sector: Theories, Reforms, and Case Studies*. Springer Nature Switzerland, 2021.
- [77] O. Rikken, M. Janssen, and Z. Kwee. “Blockchain and the Public Sector”. In: *Information Polity* 24.4 (2019), pp. 397–417. DOI: 10.3233/IP-190154.
- [78] T. Roth, A. Rieger, and A. Höß. “From Mutualism to Amensalism: A Case Study of Blockchain and Digital Identity Wallets”. English. In: *Lecture Notes in Business Information Processing*. Springer Nature Switzerland, 2024. DOI: 10.1007/978-3-031-64073-5_10.
- [79] T. Roth, A. Stohr, J. Amend, G. Fridgen, and A. Rieger. “Blockchain as a driving force for federalism: A theory of cross-organizational task-technology fit”. In: *International Journal of Information Management* 68 (2023), p. 102476. DOI: <https://doi.org/10.1016/j.ijinfomgt.2022.102476>.
- [80] T. Roth, M. Utz, F. Baumgarte, A. Rieger, J. Sedlmeir, and J. Strüker. “Electricity powered by blockchain: A review with a European perspective”. In: *Applied Energy* 325 (2022), p. 119799. DOI: <https://doi.org/10.1016/j.apenergy.2022.119799>.
- [81] C. Santana and L. Albareda. “Blockchain and the emergence of Decentralized Autonomous Organizations (DAOs): An integrative model and research agenda”. In: *Technological Forecasting and Social Change* 182 (2022), p. 121806. ISSN: 0040-1625. DOI: <https://doi.org/10.1016/j.techfore.2022.121806>.
- [82] V. Schlatt, A. Schweizer, N. Urbach, and G. Fridgen. *Blockchain: Fundamentals, Applications and Potential*. Tech. rep. 2016. URL: https://www.fit.fraunhofer.de/content/dam/fit/de/documents/Blockchain_WhitePaper_Grundlagen-Anwendungen-Potentiale.pdf.
- [83] V. Schlatt, J. Sedlmeir, S. Feulner, and N. Urbach. “Designing a Framework for Digital KYC Processes Built on Blockchain-Based Self-Sovereign Identity”. In: *Information Management* 59.7 (2022), p. 103553. DOI: <https://doi.org/10.1016/j.im.2021.103553>.

- [84] A. Schweizer, V. Schlatt, N. Urbach, and G. Fridgen. “Unchaining Social Businesses : Blockchain as the Basic Technology of a Crowdlending Platform”. English. In: *38th International Conference on Information Systems (ICIS)*. 2017. URL: <https://aisel.aisnet.org/icis2017/TransformingSociety/Presentations/8/>.
- [85] L. Schädler, M. Lustenberger, and F. Spychiger. “Analyzing decision-making in blockchain governance”. In: *Frontiers in Blockchain* 6 (2023). DOI: 10.3389/fbloc.2023.1256651.
- [86] Schütte, G. Fridgen, W. Prinz, T. Rose, N. Urbach, T. Hoeren, N. Guggenberger, C. Welzel, S. Holly, A. Schulte, P. Sprenger, C. Schwede, B. Weimert, B. Otto, M. Dalheimer, M. Wenzel, M. Kreutzer, M. Fritz, U. Leiner, and A. Nouak. *Blockchain and smart contracts : Technologies, research issues and applications*. English. Tech. rep. 2018. URL: <https://eref.uni-bayreuth.de/45018/>.
- [87] J. Sedlmeir, J. Lautenschlager, G. Fridgen, and N. Urbach. “The Transparency Challenge of Blockchain in Organizations”. In: *Electronic Markets* 32 (3 2022), pp. 1779–1794. DOI: 10.1007/s12525-022-00536-0.
- [88] J. Sedlmeir, J. Lautenschlager, G. Fridgen, and N. Urbach. “The transparency challenge of blockchain in organizations”. English. In: *Electronic Markets* 32.3 (2022). ISSN: 1019-6781. DOI: 10.1007/s12525-022-00536-0.
- [89] J. Sedlmeir, P. Ross, A. Luckow, J. Lockl, D. Miehle, and G. Fridgen. “The DLPS: A New Framework for Benchmarking Blockchains”. In: 2021. DOI: 10.24251/HI CSS.2021.822.
- [90] S. Seebacher and R. Schüritz. “Blockchain Technology as an Enabler of Service Systems: A Structured Literature Review”. In: *Exploring Services Science*. Springer International Publishing, 2017, pp. 12–23.
- [91] A. Sims. “Blockchain and Decentralised Autonomous Organisations (DAOs): The Evolution of Companies?” In: *New Zealand Universities Law Review* 28 (2019), pp. 423–458. DOI: 10.2139/ssrn.3524674. URL: <https://ssrn.com/abstract=3524674>.
- [92] *Soulbound*. URL: <https://vitalik.eth.limo/general/2022/01/26/soulbound.html> (visited on 12/11/2024).

- [93] A. Sunyaev, N. Kannengießer, R. Beck, H. Treiblmaier, M. Lacity, J. Kranz, G. Fridgen, U. Spankowski, and A. Luckow. “Token Economy”. English. In: *Business and Information Systems Engineering* (2021). DOI: 10.1007/s12599-021-00684-1.
- [94] *SushiSwap Governance Model*. URL: <https://docs.sushi.com/concepts/governance/model> (visited on 12/05/2024).
- [95] *SushiSwap Proposals*. URL: <https://snapshot.box/#/s:sushigov.eth/proposals> (visited on 12/05/2024).
- [96] J. Z. Tan, M. Langenkamp, A. Weichselbraun, A. Brody, and L. Korpas. *The Constitutions of Web3*. 2024. arXiv: 2403.00081 [cs.CY]. URL: <https://arxiv.org/abs/2403.00081>.
- [97] *The DAO - Manifesto*. 2016. URL: <http://web.archive.org/web/20160501142416/https://daohub.org/manifesto.html> (visited on 12/05/2024).
- [98] E. Toufaily, T. Zalan, and S. B. Dhaou. “A framework of blockchain technology adoption: An investigation of challenges and expected value”. In: *Information Management* 58.3 (2021), p. 103444. DOI: <https://doi.org/10.1016/j.im.2021.103444>.
- [99] *Uniswap Docs | Uniswap*. URL: <https://docs.uniswap.org/> (visited on 12/15/2024).
- [100] S. Wang, W. Ding, J. Li, Y. Yuan, L. Ouyang, and F.-Y. Wang. “Decentralized Autonomous Organizations: Concept, Model, and Applications”. In: *IEEE Transactions on Computational Social Systems* 6.5 (2019), pp. 870–878. DOI: 10.1109/TCSS.2019.2938190.
- [101] S. Werner, D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. Knottenbelt. “SoK: Decentralized Finance (DeFi)”. In: *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. AFT ’22. Association for Computing Machinery, 2023, 30–46. DOI: <https://doi.org/10.1145/3558535.3559780>.
- [102] *YAM Post-Rescue Attempt Update*. URL: <https://medium.com/yam-finance/yam-post-rescue-attempt-update-c9c90c05953f> (visited on 12/05/2024).
- [103] *Yearn Finanace - Data from Flipsidecrypto*. URL: <https://gov.yearn.fi/t/data-from-flipsidecrypto/2018> (visited on 12/05/2024).

- [104] L. Zavolokina, R. Ziolkowski, I. Bauer, and G. Schwabe. “Management, Governance, and Value Creation in a Blockchain Consortium”. In: *MIS Quarterly Executive* 19.1 (2020), Article 3. URL: <https://aisel.aisnet.org/misqe/vol19/iss1/3>.
- [105] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang. “Blockchain challenges and opportunities: a survey”. In: *Int. J. Web Grid Serv.* 14.4 (2018), 352–375.

A | Appendix

This appendix contains three main sections:

Section 1: This section provides a comprehensive overview of the publications included in this dissertation.

Section 2: This section contains contribution statements that detail specific roles and contributions to each publication incorporated into this dissertation.

Section 3: This section is a repository for the papers referenced in the dissertation.

1 Publication portfolio

The research papers (RP) included in this dissertation are ranked according to the 2023 Scopus metrics¹, while the conferences are ranked based on the 2021 GII-GRIN-SCIE (GGS) standards².

- **RP1** - O. Papageorgiou et al. [65]: O. Papageorgiou, L. Börtzler, E. Ermolaev, J. Kumari, J. Sedlmeir. "What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains", In: Proceedings of the 58th Hawaii International Conference on System Sciences, (2025), pp. 7431-7440, URL: <https://hdl.handle.net/10125/109740>.
GGS Ranking: A
- **RP2** - Barbereau et al. [11]: T. Barbereau, R. Smethurst, O. Papageorgiou, J. Sedlmeir, and G. Fridgen. "Decentralised Finance's timocratic governance: The

¹<https://www.scopus.com/sources.uri>

²<https://scie.lcc.uma.es/>

distribution and exercise of tokenised voting rights”. In: *Technology in Society* 73 (2023), p. 102251. DOI: <https://doi.org/10.1016/j.techsoc.2023.102251>.

Scopus Ranking: 99%

- **RP3** - Barbereau et al. [10] - T. Barbereau, R. Smethurst, O. Papageorgiou, A. Rieger, G. Fridgen, "DeFi, not so Decentralized: The Measured Distribution of Voting Rights", In: *Proceedings of the 55th Hawaii International Conference on System Sciences*, (2022), pp. 6043–6052, <https://doi.org/10.24251/HICSS.2022.734>. GGS Ranking: A
- **RP4** - Fridgen et al. [35] - G. Fridgen, R. Kräussl, O. Papageorgiou, and A. Tugnetti. "Pricing dynamics and herding behaviour of NFTs". In: *European Financial Management*. DOI: <https://doi.org/10.1111/eufm.12506>. Scopus Ranking: 83%
- **RP5** - Delgado Fernández et al. [27] - J. D. Fernandez, T. Barbereau, and O. Papageorgiou. "Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch". In: *ACM Trans. Manage. Inf. Syst.* 15.1 (2024). URL: <https://doi.org/10.1145/3649318>. Scopus Ranking: 79%

2 Individual Contribution

Contribution statements based on the CRediT system are provided, detailing the specific roles and responsibilities undertaken for each publication included in this dissertation.

RP1: What Blocks My Blockchain’s Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains

Orestis Papageorgiou: Conceptualization, Methodology, Data curation, Formal analysis, Writing – original draft, Software, Writing – review & editing, Visualization.

RP2: Decentralised Finance’s timocratic governance: The distribution and exercise of tokenised voting rights

Orestis Papageorgiou: Conceptualization, Methodology, Data curation, Writing – original draft, Software, Writing – review & editing, Visualization.

RP3: DeFi, not so Decentralized: The Measured Distribution of Voting Rights

Orestis Papageorgiou: Conceptualization, Methodology, Data curation, Writing – original draft, Software, Writing – review & editing, Visualization.

RP4: Pricing dynamics and herding behaviour of NFTs

Orestis Papageorgiou: Conceptualization, Methodology, Data curation, Writing – original draft, Software, Writing – review & editing, Visualization.

RP5: Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch

Orestis Papageorgiou: Conceptualization, Methodology, Writing – review & editing.

3 Appended Research Papers

3.1. Research Paper 1: *What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains*

What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains

Orestis Papageorgiou
SnT – University of Luxembourg
orestis.papageorgiou@uni.lu

Lasse Börtzler
Karlsruhe Institute of Technology
lasse.boertzler@student.kit.edu

Egor Ermolaev
SnT – University of Luxembourg
egor.ermolaev@uni.lu

Jyoti Kumari
SnT – University of Luxembourg
jyoti.kumari@uni.lu

Johannes Sedlmeir
SnT – University of Luxembourg
johannes.sedlmeir@uni.lu

Abstract

Permissioned blockchains have been proposed for various use cases where a certain degree of decentralization is necessary yet enterprise IT requirements must be met. However, their throughput remains considerably lower than that of established centralized systems. Previous studies that address permissioned blockchains' performance remain blockchain-specific, lacking a generalizable approach for locating and understanding bottlenecks. This paper presents a unified, graphical method for identifying bottlenecks in permissioned blockchains. We augment the DLPS – an open-source benchmarking tool – with graphical evaluation functionalities and use them to identify performance bottlenecks of Hyperledger Fabric and Quorum, two widely used permissioned blockchains with distinct architectural designs. Our work provides researchers and practitioners with a toolkit, guidelines on blockchain performance data analytics, and insights that assist with the bottleneck identification and improvement of permissioned blockchains.

Keywords: Blockchain, Distributed Ledger, Performance Evaluation, Hyperledger Fabric, Quorum

1. Introduction

Since its inception by Nakamoto (2008), blockchain technology has been explored across various industries far beyond its use in the cryptocurrency Bitcoin. Researchers and practitioners have analyzed its potential in a variety of applications where a neutral platform is desirable (Sedlmeir et al., 2022b), such as supply chain management (Queiroz et al., 2020) and the streamlining of cross-organizational workflows (Fridgen et al., 2018). Organizations looking to implement blockchain-based

information systems often opt for permissioned blockchains as they restrict participation in consensus, reduce data visibility and latency, and allow for better throughput. However, transitioning projects based on permissioned blockchains from pilot stages to business applications still presents many technical challenges (Toufaily et al., 2021). One major obstacle is the technology's inferior performance compared to established centralized systems, stemming from the resource-intensive nature of replication and consensus (Sedlmeir et al., 2022a). As a result, a substantial part of research focuses on examining the performance characteristics of permissioned blockchains (Fan et al., 2020). Blockchain benchmarking research has primarily examined high-level performance indicators such as throughput, with less emphasis on identifying performance-limiting factors. While assessing throughput is useful for comparing different blockchains and assessing deployment parameters (Guggenberger et al., 2022), it provides limited insights into crucial aspects like node resource utilization, which are usually reported only as aggregated data (Fan et al., 2020), offering a limited understanding of the inner workings of blockchain performance.

This paper addresses this shortcoming by analyzing resource-related metrics of blockchain nodes and their impact on throughput in a systematic and graphical way, providing a general approach that can be used to detect bottlenecks. We survey related work to ground our method and then conduct *exploratory data analysis* (EDA) to determine the performance bottlenecks of Hyperledger Fabric (Fabric) and Quorum, using an extended version of the distributed ledger performance scan (DLPS) (Sedlmeir et al., 2021).

2. Background

2.1. Hyperledger Fabric

Fabric has become one of the industry's leading permissioned blockchains (Guggenberger et al., 2022), known for its unique architecture that provides ample opportunities for finetuning performance (Androulaki et al., 2018). In Fabric, nodes are grouped into organizations, and a node can take at least one of the roles of a peer node (*peer*) or an orderer node (*orderer*) (Androulaki et al., 2018). Peers maintain an append-only ledger and a corresponding running aggregate (state), whereas orderers create and broadcast blocks. Fabric relies on the *execute-order-validate* paradigm that involves three phases. In the *Execution Phase*, a client sends a signed transaction proposal to the peers. Peers *simulate* the transaction, i.e., they run the required smart contract ("chaincode") on their current version of the state without updating it. Peers then respond to the client with a signed *endorsement* containing the peer's digital certificate, the transaction's read-write set, and the simulation outcome (Androulaki et al., 2019). During the *Ordering Phase*, and once a client has collected sufficient endorsements according to a chaincode's policy, it packs them into a transaction and forwards it to the ordering service. Orderers use a consensus protocol, such as RAFT (Ongaro et al., 2014), to sort transactions, group them into a batch ("block"), and sign this block without evaluating the transactions' validities. Subsequently, they broadcast the block to a subset of peers (to one *anchor peer* per organization) for validation. In the *Validation Phase*, upon receiving a block – either directly from an orderer or via a fellow peer – a peer validates the included transactions in three steps (Thakkar et al., 2018). First, through parallel verification, *validation system chaincode* (VSCC) ensures transactions have the required endorsements and consistent execution results. Next, valid transactions undergo *multi version concurrency control* (MVCC) – a sequential check of whether the simulations were conducted on compatible ledger versions by comparing the respective read-write sets. Finally, each peer commits the transaction to their local ledger (including a flag for its validity) and – if both checks have passed – updates its state accordingly.

2.2. Quorum

Quorum is another blockchain that has emerged as a significant player among the industry's permissioned blockchains because of its similarity to the public Ethereum blockchain. Unlike Fabric, Quorum relies on the *order-execute* paradigm. It supports four

different consensus mechanisms, including RAFT. In the *Ordering Phase*, clients send signed transactions to the nodes, which perform preliminary validations (e.g., correct nonce, syntax, and signatures). Verified transactions are shared with other nodes via a gossip protocol and added to their unconfirmed transaction pools ("mempool"). The RAFT leader sorts and batches valid transactions from its mempool, compiles them into a block, and disseminates the block to follower nodes. Followers attach the received block to their ledger and send acceptance messages to the leader. After receiving acceptance messages from the majority of nodes, the block becomes the new head of the blockchain ("2-phase commit"). During the *Execution Phase*, each node then updates its state deterministically according to each included transaction.

2.3. Distributed Ledger Performance Scan

The DLPS is an open-source, end-to-end benchmarking framework where users can define specifications for various blockchain and client network configurations using a single configuration file (Sedlmeir et al., 2021). We simplified its complex deployment process by dockerizing both the deployment and experiment handlers and extended its graphical capabilities for analyzing performance and resource utilization data.¹ The benchmarking follows a recursive localization of maximum throughput by gradually increasing the request rate to determine the network's maximum throughput. The first run starts by sending (asynchronous) requests at a base rate. The DLPS measures each component's (e.g., nodes, clients) resource utilization, as well as the average request frequency (f_{req}) and response frequency (f_{resp}), of successful transactions by collecting request and response timestamps from all clients. To achieve this, the DLPS leverages system monitoring tools available by the operating system, such as `vmstat`, `mpstat`, `sar`, and `ping`, to gather comprehensive performance metrics on CPU usage, memory, network activity, and I/O operations. If f_{resp} does not deviate from f_{req} by more than a given threshold, the next run increases f_{req} . Otherwise, a certain number of retries is performed. If f_{resp} fails to get close to f_{req} also in the retries, the ramp-up sequence is terminated and the maximum throughput is set to the highest average f_{resp} in any of the runs of the previous ramp-up sequence. After the experiment, the DLPS stores fine-granular data in CSV format for further analysis and generates summary figures for illustrative purposes.

¹The source code and higher resolution figures are available at: https://github.com/orepapas/What_Blocks_My_Blockchains_Throughput_Data.

3. Related work

To gain an overview of the academic literature on bottleneck identification in permissioned blockchains, we conducted a systematic literature review. We used the broad search string (*blockchain OR “distributed ledger technology”*) AND (*performance OR throughput OR latency*) AND (*benchmarking OR measurement OR evaluation OR analysis*) on Google Scholar, ACM Digital Library, IEEE Xplore, and arXiv. Our search yielded 4,248 results. After manually reviewing titles and abstracts and removing duplicates, 57 publications remained. For these, we performed a full-text screening and excluded articles that did not provide empirical results on blockchain performance. Additionally, we excluded papers that lacked insights into potential bottlenecks, such as papers focusing on comparisons of different configurations or comparisons between different blockchains. For Fabric, we excluded research on v0.6 since it still used the order-execute architecture. We ended up with six relevant publications.

In Fabric v1.0, Androulaki et al. (2018) identify the validation phase and, in particular, the VSCC as a major bottleneck. Thakkar et al. (2018) find three major bottlenecks of v1.0, which are related to the validation phase and were addressed in subsequent Fabric versions. Ruan et al. (2020), using Fabric v1.3, also point to the validation phase as the bottleneck, especially when many unserializable transactions are included in the ledger. Wang et al. (2020) find that in v1.4, the VSCC remains the bottleneck due to limited parallelization. For the same version, Chacko et al. (2021) trace transaction failures to systemic issues, with the validation phase being the bottleneck. Specifically, MVCC read conflicts result in transaction failure, necessitating a return to the execution phase for a new round of endorsements and endorsement policy failures, which significantly slow down the VSCC and transaction processing. We could not find research focusing on identifying bottlenecks in Fabric v2.0 or higher. For Quorum, Mazzoni et al. (2021) posit that a potential bottleneck lies with the node’s remote procedure call (RPC) server buffers being capped at 128KB. While this limitation suggests a maximum transaction size of 128KB, it is improbable to be the main bottleneck, as average transaction sizes are only a few hundred bytes.

4. Method and Results

We developed our approach by analyzing the experiments of the papers mentioned in Section 3, using

the DLPS to collect data, and EDA (Chatfield, 1986) to identify bottlenecks. We collected extensive data on factors influencing node performance, ensuring a comprehensive overview of resource utilization. After cleaning and validating the data, we used EDA to detect performance irregularities. We then examined the relevant metrics to determine potential root causes. We divided our analysis into two major parts. The first part identifies potential bottlenecks by analyzing the different node resources that can impact blockchain performance. In the second part, we examine the relationship between these candidates and the blockchain’s throughput in varying degrees of resolution, such as different time windows and component selections.

Since many enterprise applications have focused on Fabric and Quorum, and Section 3 indicates that their bottleneck analysis is intricate, we detail our bottleneck identification method for both blockchains. We selected an experiment for Fabric v2.0, based on the findings of (Guggenberger et al., 2022) to determine a configuration that seems robust under modifications and extended it to Quorum v23.4. For Fabric, the network configuration comprised 16 clients, 8 peers, and 4 orderers, with four organizations comprised of four clients, two peers and one orderer each. The Quorum configuration consisted of 16 clients and 8 nodes. In both configurations, we selected RAFT as consensus mechanism due to its minimal overhead, as previous publications suggest that consensus is not the bottleneck in this case (Guggenberger et al., 2022; Mazzoni et al., 2021). We conducted the experiment on the AWS cloud platform (Amazon EC2), where each node was configured in an independent EC2 instance allocated with 16 vCPUs, 64 GB of RAM, and 1 Gbps of bandwidth running Ubuntu Server 18.04 LTS (HVM) – a typical configuration for enterprise blockchain nodes.

4.1. Fabric: Resource utilization

The initial analysis focuses on the impact of incremental increases in request rate on the resource utilization of each node. We plot data points for each node every second during a 14-second period within a 20-second experiment, excluding the first and last three seconds to avoid distortions related to the discontinuity of f_{req} (Figure 1). This approach aims to uncover trends and correlations between resource usage and increased f_{req} . We directly see that Fabric peers’ central processing unit (CPU) and network utilization exhibit clear plateaus, indicating they could be bottlenecks.




Figure 1: Fabric – Nodes' key resource utilizations.




Figure 2: Fabric – CPU utilization of all cores.

4.1.1. CPU The analysis begins by examining the average CPU utilization across a 14-second period. We observe that the linear relationship between the request rate and CPU utilization of peers breaks down at $f_{req} \approx 1600 \text{ s}^{-1}$, indicating a potential saturation point or limitation in CPU capacity. Looking into the CPU usage across individual cores (Figure 2) reveals a limited utilization for orderers and clients that is not plateauing at higher request rates. Focusing on peers, we observe a significant variance in CPU utilization, ranging from 25 % to 80 % at higher request rates. This variability suggests an uneven distribution of resources, with some peers bearing a heavier workload than others or an imbalanced allocation of tasks within some peers' cores. First, we investigate the CPU utilization per peer in Figure 3a, which indicates an equitable distribution of computational resources among the peers, with peer 5 falling behind slightly. Analyzing the mean CPU utilization of individual cores for a single peer (peer 0 in this case) in Figure 3b reveals that this also is not the cause of the high fluctuations in CPU usage since all cores show similar utilization. Because of the inconclusive results of both possible explanations, our analysis progresses to evaluate the temporal evolution of CPU usage across individual cores at the highest




Figure 3: Fabric – peer CPU utilizations.




Figure 4: Fabric – Peer 0 CPU utilization for $f_{req}=1600 \text{ s}^{-1}$.

f_{req} before the utilization plateaus (Figure 4). This f_{req} represents the peak stress on the network before any performance degradation. Here, we observe that individual cores' CPU utilization can fluctuate as much as 30 % within a single run. We are unable to deduce the reasons behind these fluctuations from CPU utilization data alone. However, as these fluctuations balance out over time according to the narrowly distributed mean CPU utilization, they are likely not the reason for the plateau. It is worth noting that Figure 3 indicates that average CPU usage plateaus at around 50 % across all cores on all peers, an improvement over previous Fabric versions. However, in scenarios utilizing a higher number of vCPUs (16 in our case), Fabric v2.0 still demonstrates a mediocre mean utilization, leaving room for further improvement (see also Thakkar et al. (2018)).

4.1.2. Network The network-related analysis begins by looking into the mean network utilization, distinguishing between inbound and outbound traffic for the different types of nodes. Orderers' traffic does not plateau at high request rates, suggesting the ordering service is not the bottleneck (Figure 5a). Notably, orderer 0 broadcasts a disproportionate amount of traffic compared to the rest, which indicates that orderer 0 is the RAFT leader and, as a result, has the additional task of broadcasting new blocks to each following orderer. Detailed traffic analysis allows its decomposition into individual components, such as the traffic generated by block propagation. According to




Figure 5: Fabric – orderer mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

Fabric’s architecture, the outbound traffic of follower orderers predominantly consists of sending blocks to the peers. From Figure 5a, and in line with expectations, consistency of inbound and outbound traffic of the followers is observed: followers receive each block once from the leader, and each orderer forwards a received block only once to an anchor peer. On the other hand, when comparing followers’ outbound traffic with the leader’s outbound traffic (Figure 5b), the ratio consistently stands around four, as the RAFT leader sends each block to each follower (three in our case) and to one peer. Our analysis excludes traffic generated by consensus-related messages, such as appended entries and heartbeat, because it is difficult to distinguish them from block propagation traffic. Nonetheless, the consensus-related messages generate significantly less traffic compared to block dissemination, making the outbound traffic of follower orderers a viable approximation for traffic related to block propagation. Regarding peers (Figure 6), we see that inbound traffic scales linearly with f_{req} for all of them, indicating it is not a bottleneck. Concerning outbound traffic, we observe that it plateaus for some peers while remaining unaffected for others (Figure 6a). Thus, we classify the peers into two main clusters, color-coded as blue and orange. We infer that blue peers act as the gossip leaders of their respective organizations, each with one follower.

The primary distinction in outbound traffic among the two types of peers stems from block propagation between them. To confirm that blue peers are gossip leaders, we deduct the traffic associated with block propagation (as obtained from the ordering service analysis) from the outbound traffic of blue peers (Figure 6b). The resulting traffic is relatively uniform, which is in line with our hypothesis that the blue peers are the gossip leaders. Furthermore, we observe a significant overlap in outbound traffic among peers, particularly at lower request frequencies, with some discrepancies at higher rates. This is




Figure 6: Fabric – peer mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).




Figure 7: Fabric – Client mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

expected, as traffic from consensus-related messages increases at higher request rates, making our outbound traffic approximation less accurate. Peer outbound traffic consists of sending requested endorsements and transaction confirmations back to clients. Since the peer outbound traffic plateaus at high f_{req} , it suggests that these components are potential bottlenecks.

We focus on client traffic to pinpoint the specific bottleneck (Figure 7). The inbound traffic of clients, which drops at higher request rates, is generated by the same components as peers’ outbound traffic, leaving us with the same potential bottlenecks. Client outbound traffic includes transaction proposals submitted to peers and endorsed transactions sent to the orderers. Since the traffic does not plateau, it suggests that these are not the bottleneck. Overall, the traffic from endorsed transactions sent to the ordering service and blocks sent to peers never plateaus, indicating that the execution and ordering phases are not the bottleneck. Since the endorsements that peers send back to the clients come in between the execution and ordering phases, they are also not the bottleneck. This leaves transaction confirmations from peers to clients, sent after the validation phase, as the likely candidate for a bottleneck. In other words, the validation phase seems to remain the bottleneck even in Fabric v2.0.

4.1.3. Memory & Hard Drive Starting with memory usage, from Figure 1, we observe that all node types exhibit non-plateauing usage levels, suggesting




Figure 8: Fabric – Rolling mean of throughput with different window sizes.

that memory constraints are not the bottleneck. For hard drive utilization, only peers' usage appears to plateau. However, peers' I/O operations (such as ledger updates or block processing) occur during or after the validation phase, which does not provide new insights. Additionally, with the peak utilization at approximately 2 %, it is evident that hard drive usage is far from reaching capacity. This indicates that constraints lie within a different component, which in turn limits hard drive utilization. As mentioned in Section 2, the validation phase is comprised of three steps: VSCC, MVCC, and each peer updating its state. As the peers' hard drive utilization is minimal, this leaves only VSCC and MVCC as the potential bottlenecks.

4.2. Fabric: Throughput

We start by gaining an overview of how the request rate affects throughput. This is achieved by plotting the throughput as a rolling mean across two window sizes (Figure 8). Using a one-second window, each data point is plotted individually, revealing significant fluctuations in throughput beyond, $f_{req}=1200\text{ s}^{-1}$, with fluctuations in f_{resp} reaching up to 1000 s^{-1} . We increase the window size to three seconds to observe the overall network performance trend. We selected this interval as it matches the average time it takes for a transaction to be committed to the blockchain under high request rates (see also Guggenberger et al. (2022)). This is significant because queuing effects become prominent at elevated f_{req} , and opting for a shorter time window could underestimate throughput. We see that, on average, the system keeps up with the request rate until reaching approximately $f_{req}=1600\text{ s}^{-1}$.

Next, we explore the correlation between Fabric's throughput and the components and resources highlighted as potential bottlenecks in the first part of the analysis, namely peer CPU utilization and peer network traffic. Figure 9 plots the two resources against the network's throughput using the three-second window. For CPU usage, we observe an initial linear




Figure 9: Fabric – Throughput against key resources.




Figure 10: Fabric – Throughput per peer.

increase with throughput until the characteristic plateau. Additionally, we observe that the instability in CPU utilization starts at around $f_{resp}=1200\text{ s}^{-1}$, which is also the point at which the correlation between f_{req} and f_{resp} starts breaking down (Figure 8a). Consequently, CPU usage correlates more closely with f_{resp} than f_{req} . Similarly, although network traffic increases with throughput, it does so at a lower rate and begins to exhibit instability at $f_{resp}=600\text{ s}^{-1}$ already, where the throughput still manages to keep up with f_{req} .

Next, we examine the throughput of individual peers (Figure 10). We observe that every peer contributes similarly to throughput, with minor variations at high f_{req} . Given the significant differences in network traffic between gossip leaders and followers, network traffic is likely not a significant factor in determining throughput. If it were, we would expect noticeable differences in throughput between gossip leaders and followers. Thus, peer CPU utilization appears to be the primary factor behind the leveling off of throughput. The crucial role of peer CPU utilization is in line with our conclusion in Section 4.1 that VSCC and MVCC are the only candidates for bottlenecks in Fabric as both depend on peer CPU. However, given that the checks in MVCC are sequential and Figure 4 shows similar mean core utilization for peers, it is unlikely that MVCC is the bottleneck. This leaves VSCC as the only bottleneck candidate. This hypothesis is further supported by the fact that the validations executed in VSCC are parallelized, and we have noted that the parallelization capacity of Fabric is limited.




Figure 11: Quorum – Key resource utilization for different request rates f_{req} .

4.3. Quorum: Resource utilization

Our benchmarking experiment for Quorum increases the request rate in increments of 100 requests per second. The Quorum bottleneck analysis based on this series of measurements also starts by gaining an overview of the four resources in relation to f_{req} (Figure 11). Similar to the Fabric case, it is apparent that CPU and network utilization are closely correlated with the request rate. Memory and hard drive usage are limited, with I/O operations exhibiting high fluctuations but generally showing less than 1% usage. Across all resources, client utilization appears to be minimal, and it is either unaffected or grows linearly with the request rate. Therefore, we again focus on the resource utilization of nodes.

4.3.1. CPU We begin the node CPU analysis by examining the utilization across all cores, where – as for the case of Fabric – we note significant fluctuations among the cores of the nodes as request rates increase (Figure 12). Examining the mean node CPU utilization (Figure 13a), we see that nodes can be grouped into three distinct categories. Node 0 (green) exhibits the highest utilization, indicating it is the RAFT leader, which is responsible for additional operations in consensus, such as transaction ordering and block composition. Nodes 1, 2, and 3 (blue) display slightly lower utilization levels, as their role in consensus is receiving and pre-validating transactions. The remaining nodes (orange) exhibit significantly lower CPU usage as they do not have additional responsibilities beyond appending blocks to their local ledgers and applying them to their state.

Examining the utilization per core of node 0




Figure 12: Quorum – CPU utilization of all cores.

(Figure 13b), we observe that one core (core 15) bears a higher workload across all request rates. This is because, except pre-validation, all other tasks of the leader are executed sequentially, leading to a disproportionate strain on one core. Following this observation, we examine the temporal evolution of CPU usage across individual cores at $f_{req}=2300\text{ s}^{-1}$ (Figure 14). Here, we see fluctuations by as much as 20% for the RAFT leader and followers, with disparities of up to 10% between individual cores when excluding core 15 for node 0. The main difference between the mean utilization of node 0 and nodes 1, 2, and 3 comes only from core 15. These results suggest that parallel processing in Quorum is even more limited than in Fabric, with average core usage not exceeding 40% and specific leader tasks overburdening one core. The CPU utilization of nodes 0, 1, 2, and 3 reaches a plateau at $f_{req}=2400\text{ s}^{-1}$. Additionally, we notice the first break in the linear relationship between CPU usage and the request rate at $f_{req}=1800\text{ s}^{-1}$. While this does not necessarily indicate a bottleneck, it could provide clues for identifying factors behind the decline in CPU utilization. From Figure 13a, we infer that node 0 and nodes 1, 2, and 3 have similar behavior after reaching a plateau. Therefore, transaction ordering and block building, the main unique operations performed by the leader, are likely not the bottlenecks. Examining Figure 13b, we see a rapid decline in the utilization of core 15 at high request rates. Considering that the remaining sequential operations, such as transaction ordering and block propagation, require minimal CPU resources, this sharp drop cannot be justified, and it appears that another component limits CPU utilization.

4.3.2. Network Focusing on mean network utilization (Figure 15), we can categorize the nodes into three groups. Due to the complexity of Quorum network traffic, we cannot accurately decompose it into individual components, so we rely on the architectural design to identify bottlenecks. Starting with the blue nodes, outbound traffic levels off at $f_{req}=2400\text{ s}^{-1}$ while




Figure 13: Quorum – Mean CPU utilization.




Figure 14: Quorum – CPU utilization for $f_{req}=2300$ s⁻¹.

signs of plateauing in their inbound traffic appear at $f_{req}=1800$ s⁻¹. The inbound traffic primarily consists of transactions received either directly from clients or through gossip and blocks from the leader, while outbound traffic originates from the dissemination of pre-validated transactions. This suggests that these operations are the limiting factors at their respective request rates. While we see similar patterns for the orange nodes, the patterns of the leader are essentially the opposite of those of other nodes. The inbound traffic is comprised of all the transactions that are broadcasted to the network and reach the leader through gossip or directly from the clients, and plateaus at $f_{req}=2400$ s⁻¹. The outbound traffic involves mainly block dissemination to the other nodes and plateaus at $f_{req}=1800$ s⁻¹. Since their inbound traffic plateaus at $f_{req}=2400$ s⁻¹, the leader likely keeps receiving the transactions from the blue nodes normally up until that point, leaving us only with block dissemination as the main bottleneck at $f_{req}=1800$ s⁻¹ and transaction propagation as the main issue for $f_{req}=2400$ s⁻¹.

Considering our CPU utilization findings, we posit that the rapid drop in CPU utilization for core 15 at $f_{req}=2400$ s⁻¹ (Figure 13b) is caused by the leader not receiving enough transactions from the other nodes. At $f_{req}=1800$ s⁻¹, the decline could be attributed to either the block propagation or one of the processes preceding it, such as transaction pre-validation and adding the block to the chain. Considering that appending the block to a node's local ledger is not CPU intensive,




Figure 15: Quorum – Node mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

the primary issues likely lie with block propagation or pre-validation. Using the same line of reasoning for $f_{req}=2400$ s⁻¹, the bottleneck appears to be either transaction propagation or pre-validation, as it is the only operation that precedes propagation.

4.3.3. Memory & Hard Drive Starting with memory utilization, we see consistent behavior across all nodes, with no signs of plateauing (Figure 11). As such, memory does not seem to contribute to performance degradation. Despite observing significant peaks in hard drive utilization (Figure 11), the mean utilization remains mostly below 1%, indicating it is not a constraining factor. The large fluctuations are probably related to the writing of the block into each node's database, but since it is improbable that it leads to a bottleneck, we do not examine it further.

4.4. Quorum: Throughput

Analyzing the correlation between f_{resp} and f_{req} , we see that even for the one-second window size, throughput remains stable but starts to exhibit higher fluctuations at $f_{req}=1800$ s⁻¹. Examining the three-second window, throughput plateaus at around $f_{resp}=2100$ s⁻¹, significantly lower than the maximum request rate. This suggests that the overall performance of the blockchain started to decline before reaching the highest request rate. This indicates that the performance degradation noted at $f_{req}=1800$ s⁻¹ for both CPU and network utilization may be more critical in identifying the bottleneck.

Examining throughput against the CPU and network utilization (Figure 17), we see similar patterns. Both resources keep up with throughput initially and experience higher fluctuations around $f_{resp}=1800$ s⁻¹. Beyond this point, differences in throughput between request rates become less pronounced (even close




Figure 16: Quorum – Rolling mean of throughput with different window sizes.




Figure 17: Quorum – Throughput against key resources.

to indistinguishable near $f_{resp}=2100\text{ s}^{-1}$). Given these similar patterns, we search for a potential common source behind the blockchain’s performance degradation. Since block and validated transaction propagation are network-related, if they were the bottleneck, they would mainly impact network traffic, and, as a result, they are less likely to be the bottleneck.

This observation leaves only transaction pre-validation as the possible constraining factor, which impacts both CPU and network traffic when nodes are not receiving enough transactions. According to Figure 11, network utilization of clients increases linearly with the request rate, which suggests they send the proper number of transactions to the nodes. Examining further the interaction of nodes with the incoming transactions, we look into the number of rejected transactions (Figure 18a). Rejected transactions are those that nodes decline to propagate, leading to clients receiving nearly immediate (within 50 ms) notifications of transaction failure. Initially, the count of rejected transactions is minimal but begins to surge at $f_{req}=1800\text{ s}^{-1}$, culminating in approximately 7000 rejections by $f_{req}=2500\text{ s}^{-1}$. This corresponds to a rejection rate of 20 %, prompting further investigation into the underlying causes.

In Quorum, transactions are categorized as either executable or non-executable. Executable transactions can be immediately included in a block, while non-executable transactions are out of nonce order and must wait for preceding transactions with lower nonce to execute first. Clients are limited to keeping




Figure 18: Quorum – Transaction metrics.

16 executable and 500 non-executable transactions in the pool at any given time. Figure 18b illustrates that at $f_{req}=1800\text{ s}^{-1}$, the number of transactions in the pool from six clients begins to exhibit instability, diverging from the previous linear relationship with the request rate, and most clients diverge at $f_{req}=2400\text{ s}^{-1}$. This pattern, along with the fact that the number of client transactions in the pool never exceeds 500, suggests that clients are reaching the limit of executable transactions in the pool, leading to more rejections. This hypothesis is further supported by the number of rejected transactions, which increases sharply at $f_{req}=1800\text{ s}^{-1}$, coinciding with the initial reduction in CPU and network utilization and the subsequent plateau at $f_{req}=2400\text{ s}^{-1}$. This indicates that the main bottleneck for Quorum is the clients’ limit of executable transactions in the pool.

5. Conclusion

This paper introduces a general illustrative method for blockchain bottleneck identification, demonstrated through an analysis of a 12-node Fabric network and an 8-node Quorum network. Our method leverages EDA to analyze blockchain performance metrics, highlighting their specific characteristics and bottlenecks. By employing a combination of proportional analysis and the study of plateau-shaped trends in resource utilization versus transaction metrics, we uncover performance anomalies. This approach allows us to narrow down the reasons for bottlenecks by comparing the correlation between data trends, the request rates (f_{req}), and response rates (f_{resp}).

For Fabric, we identify the validation phase as the main bottleneck, even for v2.0, with VSCC being the most likely component behind the bottleneck. We were also able to showcase the average moderate degree of parallelization within Fabric, which leaves ample room for improvement. In this sense, our findings align with previous studies (see Section 3). For Quorum, we posit that the bottleneck stems from the restriction

on the number of executable transactions a client can have in the transaction pool, leading to many rejections. Additionally, our findings illustrate Quorum's relatively limited capacity for parallel processing.

Acknowledgements

Funded by the Luxembourg National Research Fund (FNR), grant reference 16326754 and NCER22/IS/16570468/NCER-FT, and by PayPal, PEARL grant reference 13342933/Gilbert Fridgen. To meet grant obligations, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

References

- Androulaki, Elli et al. (2018). "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains". In: *Proceedings of the 13th EuroSys Conference*. ACM. DOI: 10.1145/3190508.3190538.
- Androulaki, Elli et al. (2019). "Endorsement in Hyperledger Fabric". In: *Proceedings of the International Conference on Blockchain*. IEEE. DOI: 10.1109/Blockchain.2019.00077.
- Chacko, Jeeta Ann et al. (2021). "Why Do My Blockchain Transactions Fail? A Study of Hyperledger Fabric". In: *Proceedings of the International Conference on Management of Data*. ACM. DOI: 10.1145/3448016.3452823.
- Chatfield, Chris (1986). "Exploratory Data Analysis". In: *European Journal of Operational Research*. DOI: 10.1016/0377-2217(86)90209-2.
- Fan, Caixiang et al. (2020). "Performance Evaluation of Blockchain Systems: A Systematic Survey". In: *IEEE Access* 8. DOI: 10.1109/ACCESS.2020.3006078.
- Fridgen, Gilbert et al. (2018). "Cross-Organizational Workflow Management Using Blockchain Technology – Towards Applicability, Auditability, and Automation". In: *Proceedings of the 51st Hawaii International Conference on System Sciences*. DOI: 10.24251/hicss.2018.444.
- Guggenberger, Tobias et al. (2022). "An In-Depth Investigation of the Performance Characteristics of Hyperledger Fabric". In: *Computers & Industrial Engineering* 173. DOI: 10.1016/j.cie.2022.108716.
- Mazzoni, Marco et al. (2021). "Performance Evaluation of Permissioned Blockchains for Financial Applications: The ConsenSys Quorum Case Study". en. In: *Blockchain: Research and Applications*. DOI: 10.1016/j.bcra.2021.100026.
- Nakamoto, Satoshi (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. URL: bitcoin.org/bitcoin.pdf.
- Ongaro, Diego et al. (2014). "In Search of an Understandable Consensus Algorithm". In: *USENIX Annual Technical Conference*. URL: raft.github.io/raft.pdf.
- Queiroz, Maciel M et al. (2020). "Blockchain and Supply Chain Management Integration: A Systematic Review of the Literature". In: *Supply Chain Management: An International Journal* 25. DOI: 10.1108/SCM-03-2018-0143.
- Ruan, Pingcheng et al. (2020). "A Transactional Perspective on Execute-Order-Validate Blockchains". In: *Proceedings of the ACM SIGMOD International Conference on Management of Data*. DOI: 10.1145/3318464.3389693.
- Sedlmeir, Johannes et al. (2021). "The DLPS: A New Framework for Benchmarking Blockchains". In: *54th Hawaii International Conference on System Sciences*. URL: hdl.handle.net/10993/45620.
- Sedlmeir, Johannes et al. (2022a). "A Serverless Distributed Ledger for Enterprises". In: *55th Hawaii International Conference on System Sciences*. URL: <http://hdl.handle.net/10125/80228>.
- Sedlmeir, Johannes et al. (2022b). "The Transparency Challenge of Blockchain in Organizations". In: *Electronic Markets* 32 (3). DOI: 10.1007/s12525-022-00536-0.
- Thakkar, Parth et al. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *Proceedings of the 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*. IEEE. DOI: 10.1109/MASCOTS.2018.00034.
- Toufaily, Elissar et al. (2021). "A Framework of Blockchain Technology Adoption: An Investigation of Challenges and Expected Value". In: *Information & Management* 58.3. DOI: 10.1016/j.im.2021.103444.
- Wang, Canhui et al. (2020). "Performance Characterization and Bottleneck Analysis of Hyperledger Fabric". en. In: *Proceedings of the 40th International Conference on Distributed Computing Systems*. DOI: 10.1109/ICDCS47774.2020.00165.

**3.2. Research Paper 2: *Decentralised Finance's timocratic governance:
The distribution and exercise of tokenised voting rights***



Decentralised Finance's timocratic governance: The distribution and exercise of tokenised voting rights

Tom Barbereau^{a,*}, Reilly Smethurst^a, Orestis Papageorgiou^a, Johannes Sedlmeir^b, Gilbert Fridgen^a

^a SnT — Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Luxembourg City, Luxembourg

^b FIM Research Centre, University of Bayreuth, Bayreuth, Germany

ARTICLE INFO

Keywords:

Blockchain

DAO

DeFi

Ethereum

Platform governance

Voting rights

ABSTRACT

Ethereum's public distributed ledger can issue tokenised voting rights that are tradable on crypto-asset exchanges by potentially anyone. Ethereum thus enables global, unincorporated associations to conduct governance experiments. Such experiments are crucial to Decentralised Finance (DeFi). DeFi is a nascent field of unlicensed, unregulated, and non-custodial financial services that utilise public distributed ledgers and crypto-assets rather than corporate structures and sovereign currencies. The inaugural Bloomberg Galaxy DeFi Index, launched in August 2021, included nine Ethereum-based projects – non-custodial exchanges as well as lending and derivatives platforms. Each project is governed, at least in part, by unregistered holders of tokenised voting rights (also known as *governance tokens*). Token-holders typically vote for or against coders' *improvement proposals* that pertain to anything from the allocation of treasury funds to a collateral's risk parameters. DeFi's governance thus depends on the distribution and exercise of tokenised voting rights. Since archetypal DeFi projects are not managed by companies or public institutions, not much is known about DeFi's governance. Regulators and law-makers from the United States recently asked if DeFi's governance entails a new class of “shadowy” elites. In response, we conducted an exploratory multiple-case study that focused on the tokenised voting rights issued by the nine projects from Bloomberg's inaugural Galaxy DeFi index. Our mixed methods approach drew on Ethereum-based data about the distribution, trading, staking, and delegation of voting rights tokens, as well as project documentation and archival records. We discovered that DeFi projects' voting rights are highly concentrated, and the exercise of these rights is very low. Our theoretical contribution is a philosophical intervention: minority rule, not “democracy”, is the probable outcome of token-tradable voting rights and a lack of applicable anti-concentration laws. We interpret DeFi's minority rule as *timocratic*.

1. Introduction

If the global financial system's governing elites are responsible for the crisis of 2007–2008, what we need is a new form of community-based governance – an alternative to both public institutions and private companies. This is a claim espoused by proponents of blockchain and Decentralised Finance (DeFi) [1–3]. DeFi projects are typically governed by coders together with unregistered holders of tokenised voting rights (also known as *governance tokens*) [4]. Coders issue a diverse array of *improvement proposals* that are subjected to voting rounds. Votes are cast by the token-holders. DeFi's tokenised voting rights can be bought and sold in a peer-to-peer manner or via

crypto-asset exchanges. Anyone on Earth can henceforth purchase a right to vote, at a price determined by the global market. The mantra is not *one person, one vote*; individuals are free to purchase multiple voting rights and accumulate power over time, without needing to register their legal identity information.

Unlike shareholders' voting rights, issued by registered companies, DeFi's tokenised voting rights are typically issued by global entities that are unincorporated and non-liaable [5,6]. Blockchain developers refer to these ambiguous global entities as Decentralised Autonomous Organisations (DAOs). DeFi's tokenised voting rights, issued by DAOs, are supposed to offer a more democratic and inclusive alternative to corporate governance [7–9]. According to Pandian et al. [9],

* Corresponding author.

E-mail addresses: tom.barbereau@uni.lu (T. Barbereau), reilly.smethurst@uni.lu (R. Smethurst), orestis.papageorgiou@uni.lu (O. Papageorgiou), johannes.sedlmeir@fim-rc.de (J. Sedlmeir), gilbert.fridgen@uni.lu (G. Fridgen).

<https://doi.org/10.1016/j.techsoc.2023.102251>

Received 22 June 2022; Received in revised form 27 March 2023; Accepted 8 April 2023

Available online 19 April 2023

0160-791X/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

“governance tokens ensure DeFi’s democratic and decentralised governance”; but early empirical research reveals highly concentrated voting rights and an absence of applicable anti-concentration laws [10,11]. This tension between “democracy” and concentrated governance power piqued our curiosity.

The nascent field of DeFi mostly consists of unlicensed, unregulated, and non-custodial financial services that exist thanks to public distributed ledgers like the Ethereum Mainnet [6,12–15]. In January 2021, a non-custodial crypto-asset exchange named Uniswap became the first DeFi project to reach a total trading volume of USD \$100 billion. In response to interest from investors and financial analysts, Bloomberg launched the Galaxy DeFi Index in August 2021. The DeFi projects from Bloomberg’s index are not governed by private companies or public institutions. Each project is Ethereum-based, and each project is governed (at least in part) by a DAO.

In other words, the DeFi projects from Bloomberg’s index involve unregistered token-holders’ voting rights, not registered shareholders’ voting rights. Unlike shares, voting rights tokens are not legal contracts, and they do not entitle holders to a share of a registered company’s profits [6]. A community of voting rights token-holders – a DAO [16] – is not typically positioned as a *juridical person* [5]. The distinctions are clear.

In the context of DeFi’s governance, *decentralised* specifically denotes independence from regulators, registered companies, and investor registration processes [6,10,13]. Hence the Democratic Senator Elizabeth Warren [17] referred to DeFi as “the Wild West of our financial system”. The Financial Action Task Force (FATF), the Chair of the United States Securities and Exchange Commission (SEC), and Professor Kenneth Rogoff from Harvard University issued similar remarks [18–21]. According to Warren, DeFi replaces an old form of centralised control – the traditional financial system – with a new form of centralised control that is shadowy and unregistered [22]. She therefore associates DeFi not with a uniform, democratic distribution of power and resources but rather with “new concentration risks” [23; p. 3]. The Republican House Majority Whip, Tom Emmer [24], advanced an opposing view. According to him, *de-centralised* finance “shifts economic power from *centralised* institutions back into the hands of the people”. The tension is obvious: DeFi is oligarchic according to Warren, and it is democratic according to Emmer.

Early socio-technical studies about DeFi’s governance support Warren’s assertion [10,11,14,25]. The new concentration risks and oligarchic governance structures are broadly labelled “Fake-DeFi” [26] and “the illusion of decentralisation” [27]. Barbereau et al. [10] studied the tokenised voting rights issued by five Ethereum-based DeFi projects: Uniswap, Maker, SushiSwap, Yearn Finance, and UMA. The authors discovered that, for all five cases, the distributions of tokenised voting rights are highly concentrated. Their study did not, however, examine the actual *exercise* of voting rights. In response to this notable gap and DeFi’s contested *democratic/oligarchic* governance, we formulated three research questions.

1. How distributed (or, rather, concentrated) are DeFi’s tokenised voting rights?
2. What portion of DeFi’s token-holders actually exercise their voting rights?
3. What are the entitlements and the initial distribution strategies that pertain to DeFi’s voting rights tokens?

DeFi’s token-holder governance is a novel topic that cannot be subsumed under the well-known categories of corporate governance or State governance. It also cannot be subsumed under the blockchain researchers’ esoteric category of *off-chain governance* [28–31] – specifically, the governance of an entire distributed ledger by parties such as the Bitcoin Core Developers and the Ethereum Foundation.

Motivated by the topic’s novelty as well as the tension between “democracy” and concentrated power, we conducted an exploratory,

longitudinal, multiple-case study [32–34]. The nine cases we selected are all Ethereum-based, non-custodial DeFi projects that include tokenised voting rights: Uniswap, Aave, Maker, Compound, SushiSwap, Synthetix, Yearn Finance, Ox, and UMA. As noted, these cases are all included in Bloomberg’s inaugural Galaxy DeFi Index.

We quantitatively examined the “physical” artefact [34], namely the Ethereum Mainnet’s ledger, which records all transactions that determine the possession and delegation of tokenised voting rights. We employed statistical methods to determine for each case: (a) the level of *decentrality* (uniformity) achieved by the token distribution strategy over time, and (b) how often voting rights are exercised over time. We also sourced qualitative data – project documentation, white-papers, and grey literature – to derive knowledge about voting rights tokens’ entitlements and distribution strategies.

The article is primarily addressed to the critical sub-field of Information Systems research, which focuses on power relations instead of common theories about technology acceptance or transaction costs [35–37]. Hence the study’s main subject is *voting power*, and its cases are all hosted on an *information system* (Ethereum). The article is also addressed to critical, interdisciplinary researchers that study socio-technical topics [4,38,39]. Our theoretical discussion is a philosophical intervention [40,41]: in response to platitudes about DeFi’s “democratic” governance, we depict it instead as *timocratic* [42]. This is because the DeFi cases we examined each exhibit minority rule (specifically, concentrated voting rights and low participation rates in voting rounds), “shadowy” unregistered entities, and large crypto-asset treasuries. The theoretical contribution is novel for two additional reasons: (1) it does not entertain the false equivalence between corporate governance and DeFi’s DAO-based governance (see Ref. [43]), and (2) it does not opportunistically import theories from New Institutional Economics, apply them to DAOs, and assume that technology principally causes or determines institutional structures (see Refs. [44–46]). Our treatment of DeFi’s governance as unorthodox and “novel” follows Vergne [47; p. 18]: he held that “theories premised on a manager- or shareholder-centric view are of limited usefulness” when it comes to understanding governance by DAOs.

2. Background

The governance of Decentralised Finance (DeFi), as the name denotes, is tied to *decentralisation* – a topic that is irreducibly political, legal, and technical [39,48]. *Decentralisation* is an equivocal, polysemous, and often confusing word that is frequently used by crypto-asset developers and researchers [49,50]. The word variably denotes socio-political ideals [28], the elusion of centralised authorities and detachment from legal legitimacy [5,38,44], and the technical features of distributed computing systems [51]. DeFi’s governance by DAOs is linked to each of these points; therefore a nuanced understanding of decentralisation is required.

Bitcoin’s blockchain network is an exemplar of decentralisation [52]. So, too, is the Ethereum Mainnet [2] – the predominant host of DeFi projects. Bitcoin’s native crypto-currency, BTC, is the fastest asset in history to reach a \$1 trillion market capitalisation [53]. It achieved this feat with no legal entity attached to it as a manager or majority owner. Bitcoin transactions are not registered on central servers; the Bitcoin network is instead distributed across a global computing grid that consists of voluntary participants [54]. Because the Bitcoin network is not wholly grounded in any particular jurisdiction, it entails challenges for regulators [55,56]. The Bitcoin network is, in the legal parlance, an “unincorporated distributed ledger system” [6; p. 13]. The network can alternatively be described as a distributed financial market infrastructure (dFMI), since it can operate autonomously and it does not rely on conventional financial market infrastructures (FMIs) such as central banks, auditors, or intermediary payment systems [57].

Decentralisation is often treated as a synonym for *distribution*, but this is not always appropriate. There is no necessary link, for instance,

between socio-political ideals of decentralisation (*qua* democratisation or egalitarianism) and an actualised, uniform distribution of capital (BTC) and governance power within the Bitcoin network [58–60]. As Kostakis and Giotitsas [52; p. 437] put it, “in theory you have equipotential individuals (that is, everyone can potentially participate in a project), but in practice what one gets is concentrated capital and centralised governance.” The 2015 *block-size debate* (and subsequent revision to the Bitcoin protocol) revealed the governance power that is concentrated among the Bitcoin Core Developers and the Lead Developer [28,60].

The Ethereum Mainnet, inspired by the Bitcoin network, launched in 2015. Ethereum introduced a Turing-complete program execution layer – the Ethereum Virtual Machine (EVM) – on top of a blockchain layer, together with a native crypto-currency (ETH). The EVM can execute persistent scripts – so-called *smart contracts* or computerised transaction protocols [61,62]. Developers can use smart contracts to create DAOs [63] or other types of decentralised governance protocols [64]. Hassan and De Filippi [16; p. 2] defined a DAO as “a blockchain-based system that enables people to coordinate and govern themselves, mediated by a set of self-executing rules deployed on a public blockchain, and whose governance is decentralised”.

This new form of governance by DAOs received attention in 2020, thanks to DeFi projects like Compound, SushiSwap, and Uniswap [12, 13,15]. In the legal parlance, DAO-based governance is *non-labile*, simply because a DAO is not a liable legal entity [6,13]. DeFi projects that are non-custodial and governed by non-labile DAOs are thus distinct from custodial crypto-asset exchanges and lending platforms that are owned by liable companies like Coinbase and Celsius [15]. DeFi projects governed by non-labile DAOs are virtually free from applicable law and regulation. Proponents note that this reduces compliance costs [65]; whereas legal scholars claim that DeFi projects can potentially “undermine traditional forms of accountability and erode the effectiveness of traditional financial regulation and enforcement” [6; p. 1].

Regulators and law-makers deploy similar, critical rhetoric about DeFi’s unorthodox governance [19,66,67]. Japan’s FinTech Innovation Hub [68], for example, labelled DeFi a threat to the ability of the Financial Services Agency (FSA) to enforce existing regulations. Commissioner Dan Berkovitz from the United States Commodity Futures Trading Commission (CFTC) warned that DeFi could become a massive, “unregulated shadow financial market” [69]. In spite of regulators’ and law-makers’ concerns, the market capitalisation of DeFi’s voting rights tokens (and consequently the tokens’ prices) increased dramatically from 2020 to early 2021 (Fig. 1).

DeFi’s voting rights tokens allow holders to vote for or against a diverse array of improvement proposals. Improvement proposals pertain to a project’s rules, parameters, or features [4,70]. Ethereum-based voting rights tokens are typically ERC-20 format; so, like other ERC-20 tokens, they can be traded on non-custodial exchanges like Uniswap

and SushiSwap or on custodial exchanges like Coinbase and Binance. Some DeFi projects allow token-holders to delegate their voting power to other wallet addresses. The delegation of voting power can be reversed by the original token-holder at any point.

DeFi’s token-holder governance is a compelling research topic, especially following the concerns expressed by regulators and law-makers about DeFi’s concentration risks and its new class of unregistered, “shadowy” elites [22]. These concerns contradict common, pre-theoretical assumptions about decentralised governance as democratic, inclusive, or empowering [2,9,24].

DeFi’s token-holder governance is a novel topic as well. Klimon [71] reviewed governance options for organisational structures that historically preceded DAOs, such as unincorporated associations, trusts, trade associations, and other membership organisations. “Membership”, wrote Klimon [71; p. 8], “is essentially the right to vote for the governing body”. In contrast to DAOs, conventional membership organisations do not involve a global community of token-holders; but more importantly, they usually involve registered voting rights that are not purchased for speculative reasons. This reinforces the distinction made by regulators and law-makers: DeFi’s voting rights are unorthodox, because they are tradable and unregistered, and their distribution is not confined to any particular jurisdiction.

Since our study focuses on DeFi’s tokenised voting rights, it is distinct from prior research that measured the concentration of wealth or the concentration of mining power in the Bitcoin and Ethereum networks (see Refs. [51,58,72,73]). The closest precedent text is a conference paper by Barbereau et al. [10].¹ As noted, Barbereau et al. [10] studied five out of nine cases that Bloomberg eventually chose for their inaugural Galaxy DeFi Index: Uniswap, Maker, SushiSwap, Yearn Finance, and UMA. The authors used multiple statistical metrics to analyse the five cases’ distributions of tokenised voting rights; but they did not, we repeat, examine the exercise of voting rights over time.

Our contribution is unique, since we studied all nine cases from Bloomberg’s inaugural Galaxy DeFi Index, and we examined both the concentration and exercise of tokenised voting rights. Following Barbereau et al. [10], we chose a complementary array of statistical metrics (or, simply put, we did not exclusively rely on the Gini coefficient), a combination of qualitative and quantitative sources [74,75], and a longitudinal multiple-case study design [34].

3. Methods

Since DeFi’s governance is a novel topic for critical, socio-technical research [6,12,13,76], we designed an exploratory multiple-case study [32,33,77]. To account for the distribution and exercise of tokenised voting rights over time, our multiple-case study is longitudinal [34]. Our total data set spans a period of 1,466 days. It draws on both quantitative and qualitative sources to comprehensively address our three research questions [74].

To reach a representative sample of our target site, we created a list of 25 tokens from the field of DeFi, sorted by market capitalisation. From this list, we selected cases that satisfy three conditions: (1) the token is issued exclusively in Ethereum’s ERC-20 format, (2) the token attributes voting rights (or, more specifically, the token denominates voting units), and (3) the project that issued the token is governed by a DAO (at least in part). The first criterion excludes voting rights tokens that are issued on blockchains other than Ethereum, namely Terra Luna, BNB Chain, and THORChain. These blockchains are not as widely adopted as Ethereum, so we excluded them. The second criterion excludes crypto-assets like tokenised collateral and stable-coins; hence it allows us to focus on tokenised voting rights. The third criterion excludes only Nexus Mutual. Nexus Mutual is primarily governed by a registered company, and its




Fig. 1. Prices of voting rights tokens, together with BTC and ETH (data retrieved from CoinGecko on 13 August 2021).

¹ The secondary precedent texts are pre-prints by Sun and Stasinakis [11], Jensen et al. [25], and Nadler and Schär [14].

tokenised voting rights uphold the principle of *one person, one vote*. Nexus Mutual's token-holders are required to register and complete a Know Your Customer (KYC) check. We excluded Nexus Mutual because of its proximity to conventional, corporate governance and its registered voting rights.

Table 1 presents the case selection. Nine cases qualified (highlighted in grey): Uniswap, Aave (formerly named ETHlend), Maker, Compound, SushiSwap, Synthetix (formerly named Havven), Yearn Finance, Ox, and UMA. Incidentally, the inaugural version of Bloomberg's Galaxy DeFi Index consists of these same nine cases.²

3.1. Data collection

There are six typical evidence sources for case studies: interviews, documentation, direct observations, participant observations, archival records, and "physical" artefacts [34]. We collected data from three typical sources: documentation, archival records, and a "physical" artefact. This allowed us to triangulate findings [75].

Documentation provides knowledge about the particular entitlements of each case's voting rights tokens and each case's token distribution strategy. This knowledge is required to address our third research question. Documentation also provides details about each case's improvement proposals [4]. We retrieved documentation from the cases' websites and file repositories (GitBook). Some of our cases have changed and developed significantly since their date of inception; hence we also considered archival records. Archival records allowed us to gauge the differences between project versions.

We retrieved documentation and archival records via a structured, two-stage process. First, in a scouting phase, we mapped out the available data sources for various, open-source DeFi projects: blogs, forums, developer repositories (GitHub), and Wikis. We evaluated these sources based on the number of available documents, descriptive richness, and technical depth. In the second stage, we selected the qualitative data sources that describe the project's governance.

Our third source of evidence is a "physical" artefact that is common to all nine cases: Ethereum's public ledger. Previous studies treated public ledgers as "physical" artefacts: for example, research dedicated to blockchains' throughput, fees, and transaction volumes [78–80], Bitcoin miners [81], and the performance of Uniswap's Automated Market Maker (AMM) protocol [82]. Additional examples include: studies about the feasibility of de-anonymisation [83,84], transaction *front-running* [85,86], and blockchain performance or improvement opportunities [87,88].

Ethereum's public ledger records a variety of transactions that pertain to each case's voting rights tokens, from their creation and initial distribution to buying, selling, and staking. Ethereum's public ledger is thus an eligible and crucial data source for our multiple-case study [34]. We used Dune's analytics platform and Structured Query Language to extract data from Ethereum's public ledger. Table 2 presents the quantitative data collection.

3.2. Data preparation

Following the data extraction process, we used Etherscan to identify relevant token-holders' wallet addresses and to prepare our quantitative data sets. The quantitative data sets help us answer our first two research questions, since they account for: (1) the distribution of voting rights over time, and (2) the exercise of voting rights over time (Fig. 2).

Gochhayat et al. [51] refer to statistical analyses of token distributions as assessments of *decentrality*. We retrieved all the Ethereum wallet

addresses that hold relevant, tokenised voting units at 24-hour intervals. For the cases that allow wallet addresses to deposit tokens in smart contracts (e.g., SushiSwap's MasterChef staking contract) but retain their voting power, we traced these tokens back to the original holders' addresses.

We discovered that wallets controlled by custodial crypto-asset exchanges and lending platforms like Binance and Celsius hold a significant portion of our nine cases' voting rights tokens. These parties did not exercise any of their voting rights during our analysis period; but it is worth noting that, on 19 October 2022, Binance accidentally delegated a substantial number of Uniswap voting rights [89]. The accident demonstrated that Binance at least has the technical capacity to delegate voting power, if not the political or professional will to do so. For the sake of a comparison, we created a second data set that excludes wallets controlled by custodial exchanges and lending platforms. This second data set accounts for a hypothetical scenario in which companies like Binance and Celsius preserve the current trend and choose to *not* exercise or delegate their voting rights. We excluded the following addresses from both data sets.

1. We excluded automata's smart contract addresses. Although automata's smart contract addresses can hold voting rights tokens and potentially participate in governance processes [90], we did not observe this in our nine cases. Hence we do not consider the automata's smart contract addresses relevant to our study.
2. We excluded wallet addresses that hold tokenised voting units (or portions of voting units) valued at less than \$1. These wallet addresses very rarely participate in governance processes. This is not surprising, given the fact that the transaction fees required to exercise votes are frequently greater than \$10.
3. We excluded addresses like 0x000...0000. These addresses are used to *burn* tokens, which renders tokens inaccessible and unable to attribute voting rights. We identified these addresses using the public labels from Etherscan.

We created a sum of the voting units held by all these excluded addresses, then we subtracted this sum from the overall token supply. For the first data set, we also bundled together addresses that have the same well-known controller (e.g., a custodial crypto-asset exchange or lending platform). The first data set yields the upper boundary for each case's decentrality. (Appendix B provides a mathematical proof of this argument.) It assumes that addresses not controlled by known entities (like crypto-asset exchanges) are controlled by unique individuals. (This assumption may not be correct, of course, because unregistered individuals can control more than one wallet address.)

Our second research question requires us to examine the exercise of voting rights over time. We considered on-chain voting activity as well as activity that occurs via off-chain voting mechanisms like Snapshot. We excluded activity that pertains to minor improvement proposals, such as proposals that do not have binding effects. To examine the voting activity provoked by major improvement proposals, we retrieved every wallet address that cast at least one vote. We excluded wallet addresses that only voted for or against anomalous improvement proposals that were cancelled in the middle of the designated voting period. Unsurprisingly, cancellations correspond to extremely low rates of exercised voting rights. We excluded this data, so that it would not distort our perception of the token-holders' normal participation rates.

To account for projects that permit the delegation of voting rights, we created a third data set to represent the delegates' active participation plus a fourth data set to represent the delegators' passive participation. For the third data set, the exercised voting power is treated as if it were held by the delegates' addresses. For the fourth data set, the exercised voting power is treated as if it were held by the delegators' addresses. The third data set also represents the active participation of token-holders from projects that do not permit the delegation of voting rights.

² We almost included Bancor's BNT token and SushiSwap's xSUSHI token in our case selection; but strictly speaking, these tokens do not denominate voting units. The respective voting units are denominated by two different tokens: vBNT and SUSHI.

Table 1

Case selection overview (market data derived from CoinGecko on 9 August 2021).

#	Name	Token	Token price	Token market cap.	DAO-governed	Token format	Denominates voting units
1	Uniswap	UNI	\$28.42	\$14,573,644,345	Yes	ERC-20	Yes
2	Chainlink	LINK	\$24.27	\$10,660,779,664	No	ERC-20	No
3	Terra	LUNA	\$13.96	\$5,730,458,699	Yes	Native	Yes
4	Dai	DAI	\$1.00	\$5,532,087,205	Yes	ERC-20	No
5	Compound USD Coin	cUSDC	\$0.022	\$5,137,739,177	Yes	ERC-20	No
6	Aave	AAVE	\$374.75	\$4,817,712,614	Yes	ERC-20	Yes
7	Compound Ether	cETH	\$63.38	\$4,659,745,677	Yes	ERC-20	No
8	Compound Dai	cDAI	\$0.022	\$4,456,149,064	Yes	ERC-20	No
9	PancakeSwap	CAKE	\$18.37	\$3,786,938,059	Yes	BEP-20	Yes
10	The Graph	GRT	\$0.708	\$3,163,453,402	No	ERC-20	No
11	Maker	MKR	\$3357.32	\$2,989,256,505	Yes	ERC-20	Yes
12	Amp	AMP	\$0.061	\$2,912,536,380	Yes	ERC-20	No
13	Compound	COMP	\$468.40	\$2,540,146,333	Yes	ERC-20	Yes
14	Lido Staked Ether	stETH	\$3135.77	\$2,317,693,291	Yes	ERC-20	No
15	THORChain	RUNE	\$6.83	\$1,889,788,372	Yes	Native	No
16	SushiSwap	SUSHI	\$9.82	\$1,885,006,157	Yes	ERC-20	Yes
17	1inch	1INCH	\$2.77	\$1,800,821,120	Yes	BEP-20, ERC-20	Yes
18	Synthetix	SNX	\$10.16	\$1,690,978,613	Yes	ERC-20	Yes
19	Yearn Finance	YFI	\$33,865	\$1,197,635,194	Yes	ERC-20	Yes
20	Bancor	BNT	\$3.99	\$937,888,541	Yes	ERC-20	No
21	xSUSHI	xSUSHI	\$11.55	\$820,466,370	Yes	ERC-20	No
22	Nexus Mutual	NXM	\$120.30	\$808,684,814	No	ERC-20	Yes
23	0x	ZRX	\$0.954	\$800,108,188	Yes	ERC-20	Yes
24	UMA	UMA	\$11.49	\$700,893,630	Yes	ERC-20	Yes
25	Perpetual Protocol	PERP	\$15.94	\$694,385,701	Yes	BEP-20, ERC-20	Yes

Table 2

Overview of quantitative data collection.

Case	Extracted addresses	Extraction period	Addresses used in the analysis	Addresses with (passive) governance participation
Uniswap	527,258	2020-09-14 – 2021-08-15	486,421	1,113 (9,012)
Aave	182,556	2020-10-02 – 2021-08-15	171,424	660 (724)
Maker	144,599	2017-11-25 – 2021-08-15	135,369	–
Compound	274,803	2020-03-04 – 2021-08-15	253,861	970 (6,519)
SushiSwap	143,915	2020-08-28 – 2021-08-15	126,370	5,458
Synthetix	164,003	2018-06-11 – 2021-08-15	161,426	391
Yearn Finance	96,227	2020-07-17 – 2021-08-15	86,752	7,566
0x	385,257	2017-08-11 – 2021-08-15	359,067	–
UMA	34,098	2020-01-09 – 2021-08-15	32,297	370

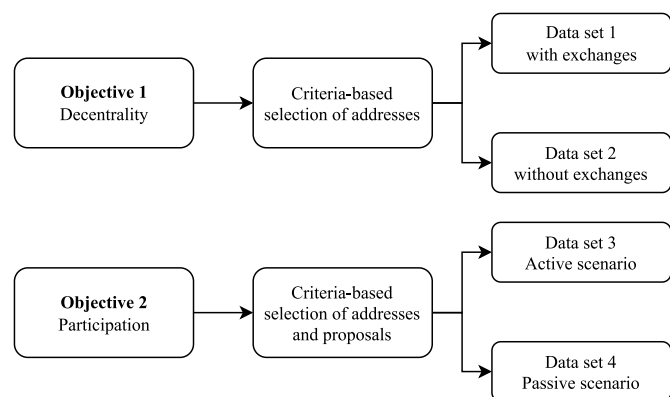
To reiterate, we created four data sets. The first and second data sets represent wallet addresses that hold voting rights tokens; the third and fourth data sets represent wallet addresses that exercise voting rights. We retrieved a total of 1,527,643 distinct addresses.

3.3. Data analysis

Data about voting rights tokens are irreducibly social and technical, electoral and financial; hence we used a mixed methods approach to analyse our three sources [74,77]. For the documentation and archival records, we used qualitative analysis techniques; and for the “physical” artefact, we used multiple statistical measures to assess decentrality and governance participation. We coded the project documentation and archival records by following the two-stage process proposed by Miles et al. [91]. During the initial stage, we considered each case’s documentation separately and assigned codes using the software MAXQDA [92]. As a group of researchers, we regularly reviewed the emerging concepts and ensured consistency in the coding system [93]. During the second stage, we clustered codes and assigned them to higher level themes that emerged contingently via data collection (inductive coding) or else pertained to an existing hypothesis (deductive coding). Since our study is not normative, we did not interpret the cases’ improvement proposals and judge whether or not they successfully uphold stakeholders’ values.

To analyse the “physical” artefact and examine decentrality for our first and second data sets, we selected statistical measures from two categories: (1) dimensionless dispersion measures, and (2) distance measures. We identified the possible candidates following a review of key texts [94–100]. These texts provide a list of multiple measures and their properties.

We first eliminated the measures that do not apply to our first and second data sets. The coefficient of variation was excluded, for example, because our data sets have heavy outliers. From families of measures (namely, the Minkowski distance family and the f -divergence family), we selected just one measure per family. From groups or pairs of measures that exhibit strong correlation by design, we again selected just one measure. Since, for example, the normalised Euclidean distance (NED) measure and the cosine similarity measure are strongly correlated, we selected just the NED. To determine the level of correlation by design, we calculated the measures using pseudo-random data sets and

**Fig. 2.** Four data sets prepared in accordance with two objectives.

the Pearson correlation coefficient [101]. We consider the measures to be strongly correlated when the absolute value of the Pearson correlation coefficient is larger than 0.7. At the end of these selection processes, we had two dispersion measures and two distance measures. These four measures are described below.

3.3.1. Gini coefficient (G)

Originally developed to assess the income and wealth inequality of different countries, the Gini coefficient [102] has found applications in multiple areas from chemistry [103] and education [104] to blockchain networks [51,73]. The coefficient takes values in $[0,1]$. A higher value indicates higher inequality. This allows for efficient interpretation and comparison of the results. The Gini coefficient should not, however, be treated as a single source of truth about inequality, as it is not possible to reduce information generated by thousands or millions of data points to a single value without losing relevant information. For our cases, the Gini coefficient is a suitable estimator of the decentrality of voting rights tokens, because an electronic voting system in which few wallet addresses hold a large portion of the tokens will exhibit high inequality. The Gini coefficient's results should not, however, be interpreted in isolation [105]. The Gini coefficient is given by

$$G = \frac{\sum_{i=1}^N \sum_{j=1}^N |p_i - p_j|}{2N \cdot \sum_{j=1}^N p_j}, \quad (1)$$

where p_i corresponds to the share of voting rights tokens held by address i and N is the total number of addresses. It is maximised through the Dirac distribution δ_{i_0} , i.e., $p_{i_0} = 1$ for some $i_0 \in \{1, \dots, N\}$ and $p_i = 0$ for all $i \neq i_0$, and minimised through the uniform distribution, i.e., $p_i = \frac{1}{N}$ for all i .

3.3.2. Normalised Shannon's entropy (NSE)

Shannon [106] initially designed the NSE to quantify lost information in phone signals. More broadly, the NSE determines the unpredictability of a distribution. The normalised version of the measure has an upper boundary, which allows for easier visualisation and interpretation of the results. The NSE takes values in $[0,1]$. Higher values indicate higher unpredictability. We assume that a globally distributed voting system can potentially exhibit high unpredictability, due to the fact that multiple individuals influence the voting outcomes. The difference between the Gini coefficient and the NSE is described in [Appendix A](#). The NSE is given by

$$\text{NSE} = - \sum_{i=1}^N \frac{p_i \log(p_i)}{\log N}, \quad (2)$$

where $0 \log(0) \equiv 0$ by convention since $\lim_{p \rightarrow 0} p \log(p) = 0$. It is 0 for δ_{i_0} and 1 for the uniform distribution (i.e., the extremes are interchanged when compared to the Gini coefficient).

3.3.3. Normalised Euclidean distance (NED)

The NED is a commonly used distance measure. It compares two distributions and measures the shortest distance between them. The NED allows us to compare each case's distribution of voting rights tokens against a hypothetical, uniform distribution in which every wallet address holds equal voting rights. If there is a small distance between a case's distribution of voting rights tokens and the hypothetical, uniform distribution, we assume that this result would be interpreted by various parties as equitable, egalitarian, democratic, or decentralised. The NED is given by

$$\text{NED} = 2^{-\frac{1}{2}} \left\| \frac{p}{\|p\|_2} - \frac{s}{\|s\|_2} \right\|_2, \quad (3)$$

where $p = (p_1, p_2, \dots, p_N)$ and $s = (s_1, s_2, \dots, s_N)$ with $s_i = \frac{1}{N}$ for every $i \in \{1, \dots, N\}$.

3.3.4. Jensen-Shannon divergence (JSD)

The JSD belongs to the f -divergence family. It measures differences in data distributions [107]. We use the JSD, like the NED, to compare each case's distribution of voting rights tokens against the hypothetical, uniform distribution. If there is a small divergence between a case's distribution of voting rights tokens and the hypothetical, uniform distribution, we assume that this result would be interpreted as decentralised. The JSD is given by

$$\text{JSD}(P \| S) = \frac{1}{2} (D(P \| M) + D(S \| M)), \quad (4)$$

where $D(P \| M) := \sum_{i=1}^N p_i \cdot \log_2 \left(\frac{p_i}{m_i} \right)$ and $D(S \| M) := \sum_{i=1}^N s_i \cdot \log_2 \left(\frac{s_i}{m_i} \right)$ with $m_i := \frac{1}{2} (p_i + s_i)$.

To analyse our third and fourth data sets and to assess the governance participation rates, we considered voting rounds for each case's major improvement proposals. We calculated the number of wallet addresses that exercised voting rights during each round, then we compared this with the number of wallet addresses that held voting rights tokens at the beginning of each round. We thus compared the potential voting rights with the exercised voting rights over time for each case.

4. Findings

4.1. Documented entitlements of voting rights tokens

The voting rights tokens from our study are case-specific. They have only three properties in common: the tokens are fungible and Ethereum-based (ERC-20), they grant holders the right to vote on community proposals, and they can be traded on both custodial and non-custodial crypto-asset exchanges. There are not many entitlements or powers pre-determined by developers. In a figurative sense, voting rights tokens grant the holder permission to enter the Senate Floor and cast votes on the measures of the day [10].

[Table 3](#) provides a reduced summary of the token-holders' entitlements that are: (a) specified in the project documentation, and (b) common to at least two cases. The entitlements can change at any time, if each project's community votes to implement changes to their token's entitlements. The table does not, therefore, represent every possible or actual entitlement.

Unsurprisingly, all the voting rights tokens entitle holders to cast votes about improvement proposals. Improvement proposals consist of executable code that can either be implemented or rejected, in accordance with the voting results. For all cases except Synthetix, token-holders can directly cast votes. Synthetix's SNX token-holders can only vote indirectly, that is, by nominating a representative from the Synthetix Council. For some cases, the power to vote and/or select what proposals are put up for a vote can be delegated to Ethereum wallet addresses that hold either no voting rights tokens or else an insufficient number of tokens. Only four cases' tokens offer a reward when they are staked: AAVE, SNX, SUSHI, and ZRX. The three lending platforms' tokens – AAVE, MKR, and COMP – entitle holders to manage the collateral types and the associated risk parameters.

There are two entitlements that happen to be unique; hence they are not mentioned in [Table 3](#). UNI entitles token-holders to manage a community treasury, and UMA entitles token-holders to cast dispute resolution votes. Uniswap's developers created the community treasury and allocated to it 43% of the total token supply (430 million UNI). The developers ceded control of the community treasury to UNI token-holders on 18 October 2020. As of this date, UNI token-holders could "vote to allocate UNI towards grants, strategic partnerships, governance

Table 3

Documented entitlements of voting rights tokens.

Case	Token	Cast votes	Choose proposals	Manage collateral	Staking reward	Delegation
Uniswap	UNI	×	×			×
Aave	AAVE	×	×	×	×	×
Maker	MKR	×	×	×		×
Compound	COMP	×	×	×		×
SushiSwap	SUSHI	×	×		×	
Synthetic	SNX	×			×	×
Yearn	YFI	×	×			
Ox	ZRX	×			×	×
UMA	UMA	×				

initiatives, [...] and other programs". The documentation's language is purposefully unambiguous, because it is not possible to predict what decisions the community will make over time about the treasury funds. As for UMA token-holders, they can vote to resolve disputes using the Voter dApp. The work required to resolve disputes generates new UMA tokens (as payment for the work). The UMA token supply is thus inflationary. There are two grounds for disputes: incorrect crowd-sourced (off-chain) pricing information, and contracts liquidated for an improper reason. UMA token-holders can assess the disputes and verify the pricing data and liquidation data. UMA employs this system, named the Data Verification Mechanism, instead of an oracle's data feed.

4.2. Comparison between trading volume and delegation volume

Four cases allow token-holders to delegate their voting rights: Ox, AAVE, Compound, and Uniswap. For these cases, we can compare the cumulative delegation volume with the cumulative trading volume. Fig. 3 illustrates that the number of traded tokens is much larger than the number of delegated tokens for each of the four cases. In short, these voting rights tokens are primarily traded; they are secondarily used for delegation purposes.

4.3. Token supply policies and distribution strategies

We collected data about two economic factors that are considered important by *tokenomics* researchers [70,108,109]: the *quasi-monetary policy* that determines the supply of tokens, and the *distribution strategy* or initial allocation of the tokens.

Table 4 provides a high-level overview of our cases' quasi-monetary policies. The emission type is usually inflationary, but in the cases of Aave and Maker, it is deflationary. Aave has a Buy-Back-and-Burn policy; Maker has a Burn-and-Mint policy. These policies are motivated by the assumed link between economic value and resource scarcity, hence they are expected to affect the value of AAVE and MKR tokens over time [70]. There is an important caveat about Maker: although the supply of MKR tokens is deflationary, Maker's DAO once elected to mint and sell additional MKR tokens in order to cover debts [110]. This explains why the fully diluted supply of MKR tokens is larger than the initial supply of MKR tokens.

Distribution strategies are important for our study of voting rights tokens, since they determine how many wallet addresses initially exercise control over a project and how much voting power these wallet addresses possess. The initial allocations of tokens are case-specific (Fig. 4). After an initial allocation is made, the tokens can be traded bi-laterally and on crypto-asset exchanges. These token transactions, recorded on Ethereum's public ledger, change the distributions of voting rights over time.

The initial allocations of MKR tokens and AAVE tokens each warrant an aside. The initial allocation of MKR tokens is not included in our study. Maker's founder, Rune Christensen, admitted that Maker's governing bodies "don't have the precise data" about the initial allocation of MKR tokens [111]. Maker's founders hold an undisclosed number of MKR tokens, and participants in three private funding rounds from 2017

to 2019 received an undisclosed number of MKR tokens. To account for the initial allocation of AAVE tokens, we had to briefly examine an early version of the Aave project from 2017 named ETHlend. ETHlend issued LEND tokens. LEND token-holders later received an opportunity to *migrate* from LEND tokens to Aave's new AAVE tokens (at a ratio of 100 LEND to 1 AAVE). The total supply of AAVE tokens is 16 million. 13 million AAVE tokens could be redeemed by LEND holders. The three million remaining AAVE tokens were allocated to Aave's treasury. We examined ETHlend's initial allocation of LEND tokens; but other details about the LEND tokens are outside the scope of our study.

Tokens are initially distributed to stakeholders from five distinguishable categories: the beachhead, team, early investors, treasury, and incentives. The beachhead is a small, influential community that promotes the project and supports the technology's adoption. The team receives remuneration or commissions for project development or related work. Early investors and venture capitalists receive tokens in response to their provision of project development funds. A project's treasury can act as a fund, overhead, or reserve. Incentives, finally, are distributed to a project's users, contributors, or prospective participants: Uniswap's *airdrop* of UNI tokens, for example, rewarded the exchange's traders and liquidity providers.

Three cases allocated a considerable portion of their tokens to the beachhead, team, and early investors: UMA (63.5%), UNI (55%), and SNX (40%). In two other cases, the team and investors received 49.95% of COMP tokens and 35% of ZRX tokens. Yearn Finance and SushiSwap did not allocate tokens to these same categories. They each branded their distribution strategy a *fair launch*. Yearn Finance initially allocated 100% of the YFI tokens to users and contributors. In the case of SushiSwap, the beachhead and early investors did not receive an initial allocation of SUSHI tokens. Prospective participants initially received 90% of SUSHI tokens. The SushiSwap treasury received the remaining 10%. At the time of this initial allocation, only developers could access the SushiSwap treasury. SushiSwap's and Yearn Finance's *fair launch* policies are thus not the same.

4.4. Decentrality: distributions of voting rights

According to our decentrality measurements, all cases' distributions of voting rights tokens are highly concentrated. The most concentrated distributions are those of COMP, UMA, and UNI. All three distributions started with a high degree of concentration, then they became even more concentrated over a short period of time. For about a year afterwards, the concentration levels remained relatively steady. Notably, throughout the entire period of our analysis, 50% of the active supply of UMA tokens was never held by more than five wallets.

MKR and ZRX tokens are the oldest. Both launched in late 2017, which is years before the launch of our other cases' tokens (aside from the aforementioned LEND tokens). The distributions of MKR and ZRX each began with a very low level of centrality that took some months to

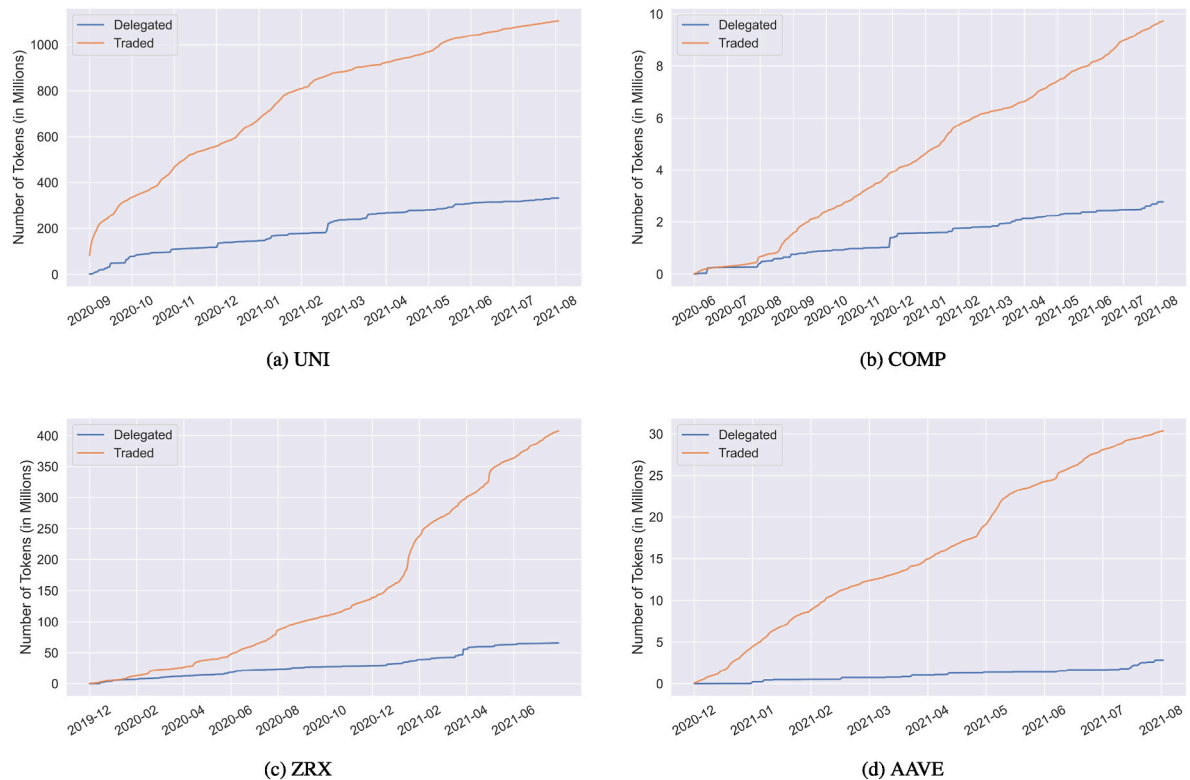


Fig. 3. Number of traded tokens against the number of delegated tokens.

Table 4

The nine cases' quasi-monetary policies.

Case	Token	Emission type	Fully diluted supply	Initial supply
Uniswap	UNI	Inflationary	1,000,000,000	1,000,000,000
Aave	AAVE	Deflationary	16,000,000	16,000,000
Maker	MKR	Deflationary	1,005,577	1,000,000
Compound	COMP	Inflationary	10,000,000	10,000,000
SushiSwap	SUSHI	Inflationary	250,000,000	0
Synthetix	SNX	Inflationary	212,424,133	100,000,000
Yearn Finance	YFI	Inflationary	36,666	30,000
Ox	ZRX	Inflationary	1,000,000,000	1,000,000,000
UMA	UMA	Inflationary	101,172,570	100,000,000

increase significantly.³ As of 15 August 2021, both the MKR and ZRX distributions are highly concentrated.

AAVE is the newest token. Its distribution began with a relatively low level of centrality, but it rapidly reached high values that are comparable with those of YFI and ZRX. For the active supply of AAVE, at any given point in time, no more than 53 addresses control more than 50%. Although this is not as concentrated as the distribution of UMA tokens, it is still highly concentrated.

As noted, SUSHI and YFI tokens were distributed via *fair launches*. The *fair launch* policies did not noticeably affect the decentrality measurements of either SUSHI or YFI. The SNX distribution exhibits the lowest level of centrality, and it did not involve a *fair launch*. As of 15 August 2021, more than 4,400 addresses (approximately 5% of the total number of addresses holding SNX) are required in order to control more than 50% of the SNX voting rights. What distinguishes Synthetix from the other eight cases is its *quadratic voting* mechanism. Whereas other




Fig. 4. Initial token allocations (for each case except Maker).

projects adopted a one-to-one connection between tokens and voting power, Synthetix's quadratic voting mechanism calculates an address's voting power as the square root of the number of tokenised voting rights that it holds. The quadratic voting mechanism is designed to limit the voting power of an address with a significant number of SNX tokens, such that an address that holds 10,000 SNX tokens, for example, would possess the power to cast only 100 votes. Furthermore, the voting power's rate of reduction *increases* in relation to the number of SNX tokens that an address holds. There is, however, an important caveat: quadratic voting is not Sybil-resistant [112], which means that individuals can divide their SNX tokens across multiple wallet addresses and obtain greater voting power than they otherwise would have if they held all their SNX tokens in a single wallet address.

Aside from SNX, the examined distributions of voting rights tokens all exhibit a very high level of centrality (Fig. 5), with fewer than 100 addresses controlling more than 50% of the voting power. This is true even when we exclude the tokens held by custodial exchanges and lending services (Fig. 6). The initial token allocation appears to have a

³ Under the NSE, the initial distribution of MKR tokens appears to be highly concentrated. This is, however, a metric artefact caused by the fact that, up until 13 December 2017, only 11 addresses held any MKR tokens, and just three of these addresses controlled more than 97% of the active supply.

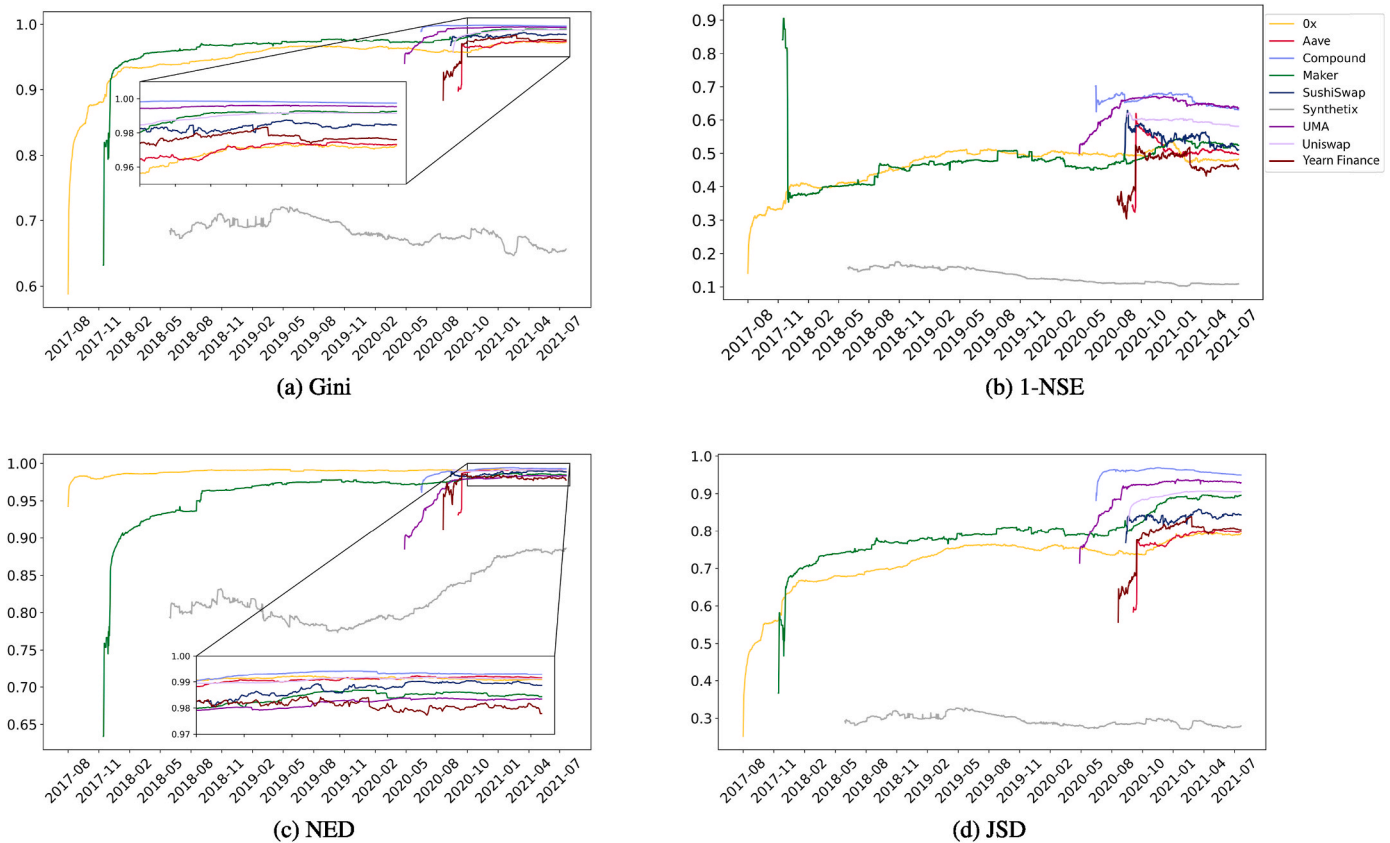


Fig. 5. Levels of decentrality (when custodial exchanges and lending platforms are included). Higher values indicate higher levels of centrality.

short-term effect on a project's decentrality. This means that, even if a case deploys a *fair launch*, the centrality level increases rapidly following the launch date, then later, the centrality seems to stabilise at a high level. Quadratic voting is seemingly the only factor from our cases that affects the decentrality measurements.

Finally, we determined that the number of voting rights tokens held by custodial exchanges and lending platforms only has a negligible impact on most cases' decentrality measurements (see Fig. 7). This is despite the fact that custodial exchanges and lending platforms hold a considerable number of tokens for each project.

4.5. Participation: exercise of voting rights

To assess the voting participation rates for each of our cases, we determined the proportion of wallet addresses that exercised voting rights in relation to the total number of wallet addresses that held eligible voting rights (see Section 3.2). For all our cases, the voting participation rates were very low. In the final months of the analysis period, only three cases witnessed more than 1% of the eligible wallet addresses engaged in voting rounds. Fig. 8 presents the voting participation rates. The active scenario accounts for voting rights that were exercised from non-delegators' wallet addresses and delegates' wallet addresses. For the sake of a comparison, the passive scenario accounts for voting rights that were exercised from non-delegators' wallet addresses and delegates' wallet addresses as well as voting rights that were delegated. In other words, the active scenario only accounts for votes that were actually cast, whereas the passive scenario also counts delegation-transactions as evidence of participation.

According to our analysis, Yearn Finance's governance participation rates are the highest. More than 60% of YFI token-holders voted for/against the first improvement proposal. This can probably be attributed to the fact that, at the time, only a small number of wallet addresses held

voting rights tokens that were earned from liquidity mining programs. This rate of participation did not last long. Shortly after crypto-asset exchanges listed YFI tokens, Yearn Finance's governance participation rates diminished drastically. Both Yearn Finance and SushiSwap have higher governance participation rates than the cases that permit delegation.⁴ In the majority of recent voting rounds, approximately 1%–2% of Yearn Finance's and SushiSwap's token-holders exercised their voting rights. Yearn Finance's rate reached as high as 7% in one of the recent voting rounds.

Aave, Compound, and Uniswap each permit delegation. In recent voting rounds, these cases' participation rates are lower than 0.1% in the active scenario, and they are consistently lower than 0.5% in the passive scenario. Aave and Compound have the lowest participation rates in the active scenario; and in the passive scenario, Aave's participation rates are the lowest by a substantial margin. As in the case of Yearn Finance and in the case of SushiSwap, Compound's initial participation rates were high, then they soon declined.

UMA's governance participation rates are notable, since they are the third highest and the most consistent. Other cases' rates exhibit significant fluctuations over time. UMA incentivises token-holders to exercise their voting rights by minting and distributing 0.05% of its token supply to wallet addresses that actually cast votes. It also offers gas rebates to cover the voting transaction costs. This effectively eliminates a drawback of on-chain voting. In spite of these incentives (and the fact that it ranks third among our cases), UMA's participation rates are not high. They are usually around 0.5%.

We were unable to include the MKR, SNX, and ZRX tokens during this stage of our analysis. In accordance with Maker's Continuous

⁴ The current version of Yearn Finance allows token-holders to delegate their voting power via Snapshot, not via an Ethereum transaction. This type of delegation is outside the scope of our data collection.

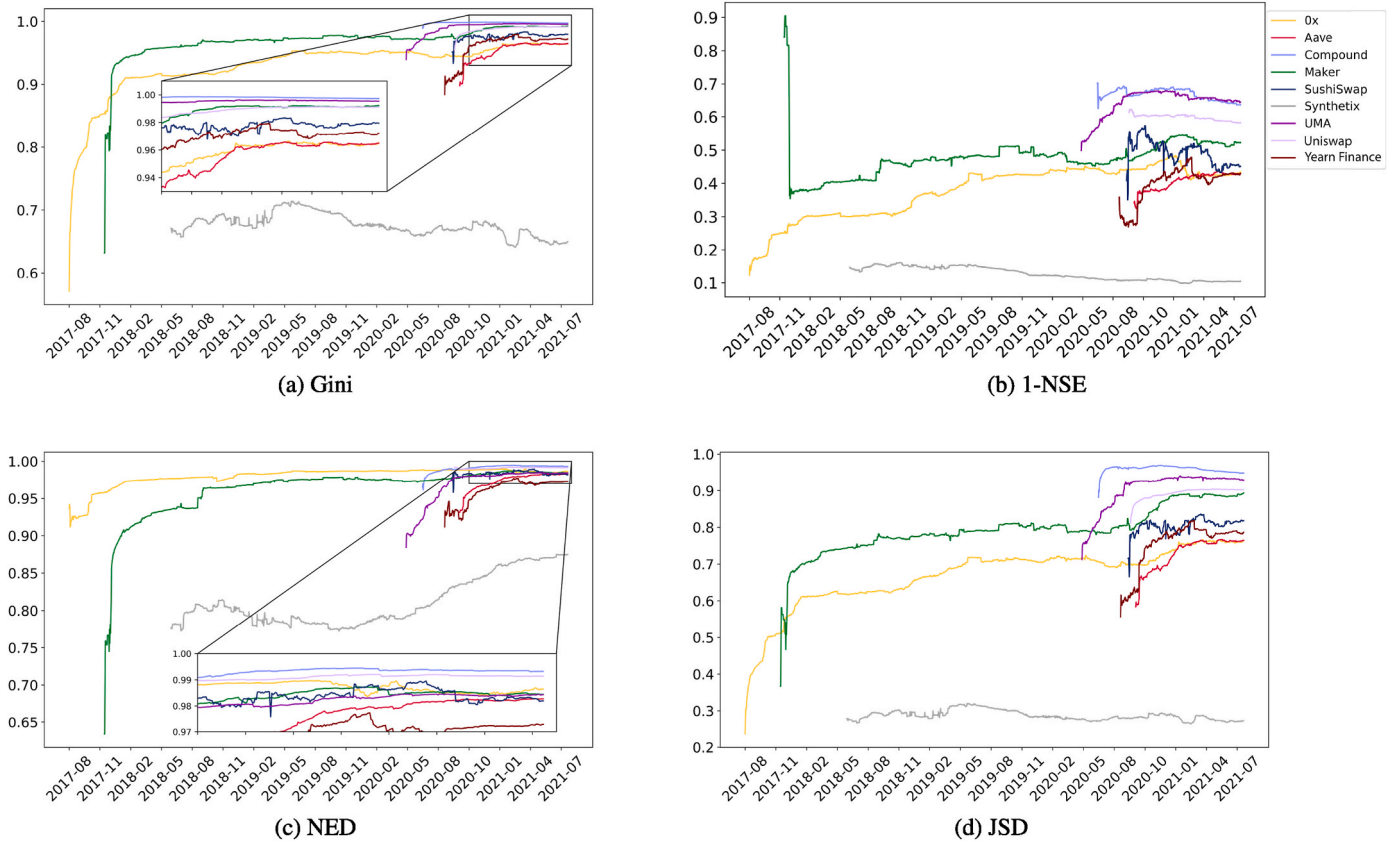


Fig. 6. Levels of decentrality (when custodial exchanges and lending platforms are excluded). Higher values indicate higher levels of centrality.

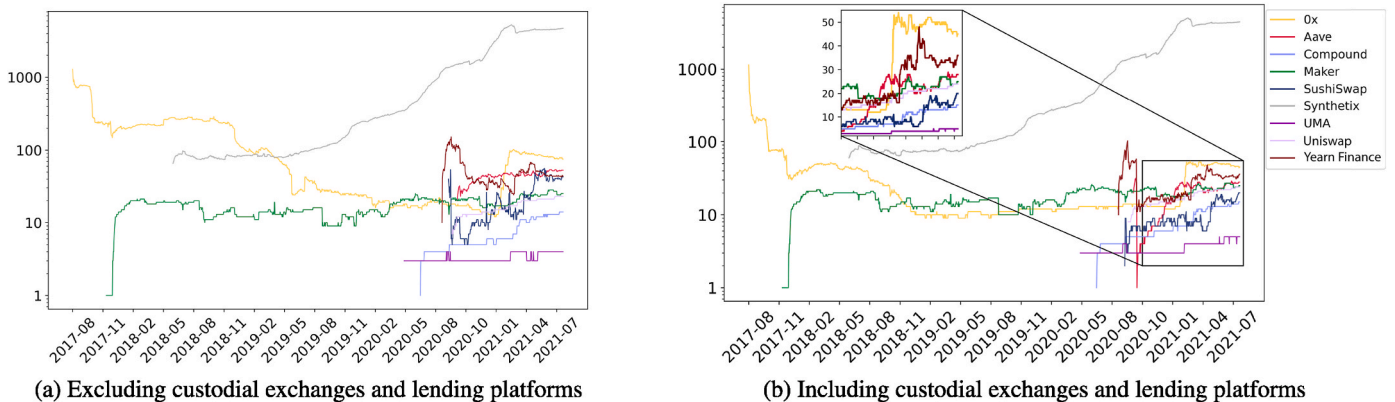


Fig. 7. Minimum number of addresses required to accumulate more than 50% of the active supply.

Approval Voting model, voting rounds for improvement proposals remain open indefinitely. We decided to not compare participation data about Maker's *indefinite* voting rounds with participation data about the other cases' *defined* voting rounds. According to a pre-print by Sun et al. [113], the participation rates in Maker's indefinite voting rounds are low. Most proposals attract votes from fewer than 60 wallet addresses (less than 0.04% of the total eligible wallet addresses). Only four proposals attracted votes from more than 100 wallet addresses. Synthetix's SNX tokens are excluded, because we could not retrieve participation data for all four Spartan Council elections that occurred during the analysis period. We could only collect data about the third and fourth proposals. These proposals attracted votes from no more than 356 wallet addresses (0.22% of the total eligible wallet addresses). As for 0x, we could not include the ZRX tokens, because 0x uses its own off-chain

mechanism for voting rounds, from which we could not access data. According to 0x's "Governance Update #1" [139], just 600 wallet addresses (0.17% of the total eligible wallet addresses) cast votes across a sum of 10 improvement proposals.

5. Discussion

Some findings are consistent across all nine cases; hence we can formulate a general definition of DeFi's voting rights tokens. Established texts about concentrated governance power, *bearer shares*, and unregulated markets account for portions of our findings; but the existing literature does not precisely account for the *type* of minority rule that is evidenced by our cases.

Repeated points from blockchain literature about DeFi's democratic

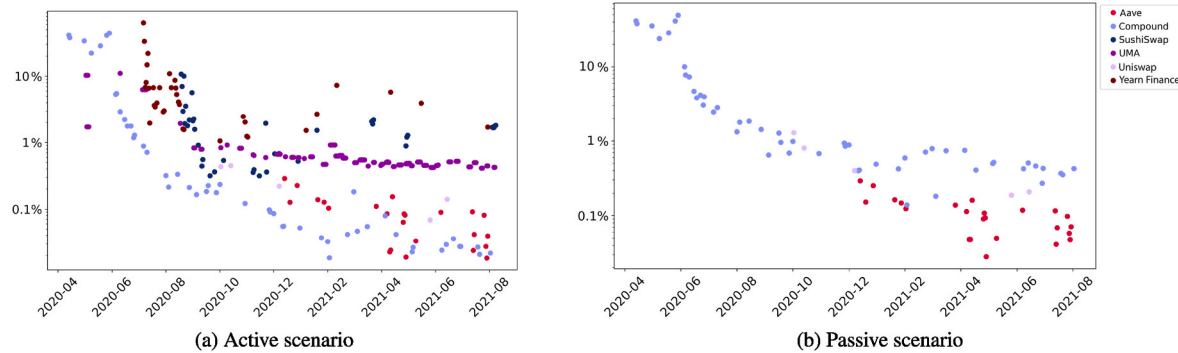


Fig. 8. Governance participation rates. Each point corresponds to an improvement proposal's voting round.

and inclusive governance are especially at odds with our findings. Hence we discuss DeFi's token-holder governance as *timocratic*, not democratic. The discussion's sceptical style is justified by a "rich data-driven inquiry" [114; pp. 271–273]. It satisfies recent calls from critical Information Systems researchers for philosophical interventions [40,41].

5.1. Definition of voting rights tokens

For all our cases, voting rights tokens are not backed by anything physical; therefore they are "digital assets" [70]. The tokens are not legal contracts, and the token-holders' legal identities are not mandatorily registered. The tokens do not entitle the holders to a share of a registered company's profits, and none of our cases' DAOs are registered as limited liability companies [27].

For the four cases that permit delegation (for actual voting purposes), the cumulative number of trading events is much greater than the number of delegation events. We thus classify voting rights tokens as *primarily* tradable. We classify voting rights tokens *secondarily* as "usage-based" [70; p. 10], since only a small minority of token-holders actually use the tokens for something other than trading. For most cases in the final months of our analysis period, this minority is less than 1%. We classify voting rights tokens *thirdly* as "distributed", in a weak sense, for even though token-holders are scattered across the globe, the token distributions exhibit low decentrality. The token distributions are in fact highly concentrated. In sum, we propose a three-part definition of voting rights tokens. Voting rights tokens are:

1. Tradable crypto-assets without registered holders.
2. "Usage-based" (but used only by a small minority).
3. "Distributed" (but far from equi-distributed).

Our first classification is corrective with respect to the intuitive assumption that voting rights tokens are primarily used for voting purposes. Our second classification is qualified: voting rights tokens are potentially utilised (for something other than trading), but only a small minority of voting rights tokens are actually utilised. Our third classification is determinedly sceptical: voting rights tokens are indeed distributed, but their distributions are highly concentrated. These classifications are crucial to our anti-platitudinous theory of DeFi's governance as a type of minority rule.

5.2. Governance discussions from legal and political texts

Since DeFi's DAOs are non-liaible and holders of tokenised voting rights are unregistered, analogies with corporate governance norms and registered shareholders' voting rights are often distant [5,6]. A close analogy can be made, however, between unregistered token-holders' voting rights and *bearer shares* that happen to attribute voting rights [115–118]. This is because they both attribute voting rights to whoever is the bearer – to whoever holds them – without requiring the disclosure

of legal identity information. A proof of legal identity is not required when DeFi's tokens are used to cast votes, just as it is not required when the tokens are sold on non-custodial exchanges like Uniswap and SushiSwap. If it were possible for global, unincorporated associations to issue bearer share certificates in a digital form, these certificates would closely resemble DeFi's tokenised voting rights.

Beyond bearer shares, corporate governance literature offers three distant analogies. First, the concentration of DeFi's voting rights among developers and venture capitalists is roughly akin to "the decline of shareholder democracy and the normalisation of founder primacy", exemplified by Snap, Inc. [119], or alternatively "the curious turn toward board primacy" [120; p. 2071]. Second, the recognition that token-tradable voting rights are not democratically distributed is loosely relevant to governance theories that propose democratic alternatives to tradable voting shares [121]. Third, DeFi's non-democratic governance can be placed in a context that includes *controlled companies*, which are governed by just one individual or legal entity. There are over 100 controlled companies listed in the Standard & Poor's Composite 1500 index [122].

Seminal works by Vilfredo Pareto [123,124] and Alexis de Tocqueville [125] help situate DeFi's governance in a modern political context. Pareto's work [123,124] accounts for the emergence of innovative elites following a crisis that affects established elites. To borrow Pareto's animal imagery, DeFi's innovative elites are foxes. They promote decentralisation in a cunning and devious manner, in order to accumulate power and wealth without public displays of force. The traditional financial system's elites are lions. They govern centralised institutions, they aim to conserve power, and they exercise force in public [126]. Pareto [124] held that minority rule is inevitable, since one class of elites follows another class of elites. Pareto's notion of recycled elites complements a maxim by Tocqueville [125]: after power is dissolved and decentralised, it is reassembled and re-centralised. As Tocqueville stated, "all authorities by nature lean towards unity"; hence, if power is decentralised, "it automatically hurtles towards centralisation". These seminal texts prompted us to specify DeFi's type of minority rule.

5.3. Philosophical intervention: DeFi's minority rule is timocratic

All nine DeFi cases exhibited concentrated voting power, which became more concentrated over time rather than less concentrated; and for most cases in the final months of our analysis period, fewer than 1% of eligible token-holders participated in governance proceedings. We deduce that minority rule is the probable consequence of tradable voting rights, ineffectual *fair launch* distribution strategies, and no applicable anti-monopoly or anti-concentration laws. A sceptical-empirical discussion of DeFi's minority rule is therefore appropriate; and, following platitudes about the democratic, inclusive, or empowering nature of DeFi and DAOs (see Refs. [7–9,46,127]), it is desirable as well.

On 7 December 2021, the founder of Yearn Finance, Andre Cronje, wrote, "Time to retire 'decentralised finance'. We aren't decentralised".

This is due to an “old guard” of DeFi insiders [128]. A few days prior, the crypto-asset researcher Ryan Selkis [129; p. 153] named “benevolent dictators”, “coin voting” (voting rights tokens), “voter apathy”, and oligarchic “collusion” as likely problems for DAOs in the year 2022. On 7 October 2021, the founder of Synthetix, Kain Warwick, identified “plutocracy” as a risk to DeFi. In Warwick’s opinion, DeFi’s current governance options are “not very generalisable” and “pretty terrible” [130]. This includes ineffectual *fair launches*. Brekke et al. [4], likewise, acknowledged claims that tokenised voting rights lead to “plutocracy” and “governing by the wealthy”. *Plutocratic* is a plausible term for DeFi’s governance; but in response to our findings, we instead claim that DeFi’s governance is timocratic. This reflects the fact that DeFi’s elites are not simply wealthy; they are sometimes “shadowy” and unidentified as well.

Timocracy is a type of stakeholder governance that dates back to the Solonian Constitution from the sixth century BC. It involves both producers (developers) and property owners (token-holders) that pursue a mixture of particular interests and common interests. Timocracy’s main risk is effectively the same as the risk identified by Warwick [131,132]: if a timocracy degenerates, then it becomes an oligarchy (or a plutocracy).

As discussed in Plato’s *Republic* [42; 545a–550c], a timocracy is a conflicted mix of community-oriented virtues and private wealth accumulation. Timocracy’s emblematic rulers hide their wealth in “treasuries and strongrooms”, which implies that timocratic power is *crypto-* (“concealed or secret”). “Victory and honour”, in a timocratic regime, are obtained via calculated risks and courageous conquests, not via harmonious rationality or dialectical reason. Timocratic rulers accumulate wealth in subtle ways that are odds with the commons, and they “run away from the law like children running away from their father” [42]. Socrates’ contemporary Sparta is an example of a timocratic regime, and it enjoys an eponymous relationship with Synthetix’s Spartan Council. Incidentally, Lysurgus of Sparta first implemented a separated powers model, which Synthetix and Yearn Finance later adopted [133].

A timocracy degenerates into an oligarchy when its constitution is rewritten to explicitly reserve power for the wealthy, and when contracts are created to purposefully exacerbate the division between a wealthy class and a poor class [42]. Although our nine DeFi cases are tacitly governed by wealthy, above-average token-holders (whales), average and below-average token-holders (minnows) are allowed to participate in the governance procedures. The minnows’ participation is not as effective as the whales’ participation; but it is not strictly precluded. Furthermore, developments like Synthetix’s quadratic voting are intended to stop the inequality gap from widening.

To reiterate, DeFi’s governance is timocratic. DeFi’s timocratic rulers – the unregistered, above-average token-holders – accumulate more power over time by purchasing even more tokens. The concentration of power is gradual, subtle, and putatively neutral, because crypto-asset marketplaces are open to all [52]. DeFi’s timocratic governance could degenerate and become oligarchic if above-average token-holders vote for so-called improvement proposals that significantly diminish the power of average and below-average token-holders, or if they vote for changes that make it more difficult for newcomers to acquire and exercise voting rights [42; 550d–551b]. Such actions would further strengthen the “old guard”.

6. Summary and outlook

The governance of DeFi projects by DAOs is putatively democratic and inclusive (see Refs. [7–9,127]). DAOs are global, unincorporated associations – euphemistically referred to as online communities – with unregistered token-holders. When we examined both the distribution and exercise of tokenised voting rights, issued by nine DeFi projects, we discovered the following:

1. DeFi’s tokenised voting rights are highly concentrated.

2. A very small portion of DeFi’s token-holders exercise their voting rights.
3. DeFi’s voting rights tokens let holders decide whether or not to implement a contingent array of improvement proposals.

These points directly answer our three research questions.

6.1. Theoretical contribution

A philosophical intervention is required to dispel spurious, platitudinous notions of “democracy” and to precisely account for DeFi’s *type* of minority rule. We drew on Plato’s *Republic* [42; 545a–550c] and argued that DeFi’s minority rule is timocratic. We also played on the association between Synthetix’s Spartan Council and the timocratic governance model implemented by Lysurgus of Sparta.

6.2. Limitations and future research topics

Our study bears two notable limitations. First, the selected cases can change by the day. This is especially true for the projects’ documentation [129]; hence the validity of our qualitative data is limited to the period of extraction. Our findings are potentially undermined by any community-formulated improvement proposals that were implemented after this article’s initial submission date. Synthetix’s governance structures in 2023, for example, differ substantially from the governance structures that we examined in 2020 and 2021. Simply put, our findings are context-specific. Second, the metrics we used to assess decentrality are imperfect, and different metrics can yield different outcomes. Hence we employed a variety of metrics from different categories. We are reasonably confident that our measured outcomes reflect reality.

DeFi’s governance is a “shadowy” and fascinating research topic that is likely to become more complicated in future. Already, DAOs can optionally register as unincorporated non-profit associations [134], limited liability companies in States like Wyoming [27], or non-grantor purpose trusts in the jurisdiction of Guernsey [135]. From 2024 onwards, a DAO can be recognised by the State of Utah as a *juridical person* that is distinct from a limited liability company [136] – a “LLD” for short, not a “LLC”. There are also options to “wrap” less exotic legal forms around the operations of DAOs [137]. We believe that DeFi offers multiple avenues for future research and theory development.

6.3. Coda

DeFi is both decentralised and centralised. It is *de*-centralised, in a stereotypical sense, if its governing bodies are *non*-liable, the holders of its voting rights are *un*-registered, its distributed ledgers are *non*-localisable and *un*-incorporated, its assets are *not* custodied, and its exchanges are *un*-regulated. At the same time, DeFi is politically centralised, because its governance is timocratic.

Credit author statement

Tom Barbereau: Conceptualisation, Method, Data curation, Formal analysis, Writing – original draft, Writing – review & editing. **Reilly Smethurst:** Conceptualisation, Method, Data curation, Formal analysis, Writing – original draft, Writing – review & editing. **Orestis Papa-georgiou:** Conceptualisation, Method, Software, Data curation, Formal analysis, Visualisation, Writing – original draft, Writing – review & editing. **Johannes Sedlmeir:** Conceptualisation, Method, Data curation, Supervision, Validation, Writing – review & editing, Funding acquisition. **Gilbert Fridgen:** Supervision, Writing – review & editing, Funding acquisition.

Data availability

The data that has been used is confidential.

Acknowledgments

This research was funded by the Luxembourg National Research

Fund (FNR) and PayPal (PEARL grant reference 13342933/Gilbert Fridgen). The authors thank Udo Zander, Svenja Miltner, and the anonymous reviewers for their detailed feedback.

Appendices

Appendix A. Gini coefficient compared to the NSE

Even though the Gini coefficient and the NSE can be used to evaluate the same attribute, there are important differences between the two measures. Unlike the NSE, the Gini coefficient takes into consideration the ranks of the data points in a distribution. This can be seen in Figure A.9: according to the Gini, distribution (2) has lower inequality than distributions (1) and (3); whereas according to the NSE, the three distributions are effectively the same. The Gini produces a different result, because the three data points on distribution (2) have the same rank. It is important to highlight that this is not indicative of the Gini coefficient's poor performance, as both interpretations of inequality can be valid depending on the situation.




Fig. A.9. The values of the Gini coefficient and NSE for three different distributions.

Appendix B. Gini coefficient's behaviour with many little-funded addresses

For $X \equiv X^{(0)} = (x_1, x_2, \dots, x_N)$ representing a distribution of holdings where $0 \leq x_1 \leq x_2 \leq \dots \leq x_N \leq 1$ and $\sum_{i=1}^N x_i = 1$, the Gini coefficient $G \equiv G^{(0)} \equiv G(X^{(0)})$ can be expressed by [138].

$$G = \frac{2 \sum_{i=1}^N i x_i}{N} - \frac{N+1}{N} \Rightarrow 2 \sum_{i=1}^N i x_i = NG + N + 1. \quad (B.1)$$

Let us define a modified distribution by adding a single address that holds tokens corresponding to an $\epsilon^{(1)}$ share of the new total supply (old supply + tokens of the new address), which is not greater than the shares held by any of the previously existing addresses after redistribution:

$$X_i^{(1)} := \begin{cases} \epsilon^{(1)}, & i = 1 \\ x_{i-1} - \epsilon^{(1)}, & 2 \leq i \leq N+1, \end{cases} \quad X_i^{(1)} = X_{i-1}^{(0)} - \epsilon^{(1)},$$

where $\sum_{i=1}^{N+1} X_i^{(1)} = \epsilon^{(1)}$, $\epsilon^{(1)} \leq x_1 - \epsilon^{(1)}$, and $x_{i-1} - \epsilon^{(1)} \leq x_i - \epsilon^{(1)}$ for all $i \in \{2, \dots, N\}$. Then $X^{(1)}$ also satisfies the conditions required for (B.1) to hold, and

$$2 \sum_{i=1}^{N+1} i X_i^{(1)} = 2 \sum_{i=1}^N (i+1)(x_i - \epsilon^{(1)}) + 2\epsilon^{(1)} = 2 \sum_{i=1}^N i x_i + 2 - 2 \sum_{i=1}^N i \epsilon^{(1)} = NG + N + 3 - 2 \sum_{i=1}^N i \epsilon^{(1)}.$$

Let

$$D_1 := G^{(1)} - G^{(0)} = G^{(1)} - G = \frac{2 \sum_{i=1}^{N+1} i X_i^{(1)}}{N+1} - \frac{N+2}{N+1} - G = \frac{2 \sum_{i=1}^{N+1} i X_i^{(1)} - N - 2 - NG - G}{N+1} = \frac{1 - 2 \sum_{i=1}^N i \epsilon^{(1)}}{N+1} - G$$

and define inductively $X^{(k)}$ by adding an address with holdings $\epsilon^{(k)}$ to the distribution described by $X^{(k-1)}$ and $D_k := G^{(k)} - G = G(X^{(k)}) - G$. It follows that

$$D_k = G^{(k)} - G = \sum_{j=1}^k \frac{1 - 2 \sum_{i=1}^{N+j-1} i \epsilon^{(j)}}{N+j} - G^{(j-1)}, \quad G^{(j)} = \frac{1 - 2 \sum_{i=1}^{N+j-1} i \epsilon^{(j)} + (N+j-1)G^{(j-1)}}{N+j}.$$

We also have

$$X_i^{(j)} = \begin{cases} e^{(j-i+1)} - \sum_{m=1}^{i-1} e^{(j-m+1)}_{i-m}, & i \leq j, \\ x_{i-j} - \sum_{m=1}^j e^{(j-m+1)}_{i-m}, & i \geq j+1. \end{cases}$$

For our multiple-case study, we exclude the addresses that hold token balances valued at less than \$1, and we subtract the balances of these excluded addresses from the total supply. Consequently, in our special case, $e_i^{(1)} = x_i e^{(1)}$ and

$$0 \leq e_i^{(j)} = X_i^{(j)} e^{(j)} = \begin{cases} \left(e^{(j-i)} - \sum_{m=1}^{i-1} e^{(j-m)}_{i-m} \right) e^{(j)}, & i \leq j-1 \\ \left(x_{i-j+1} - \sum_{m=1}^{j-1} e^{(j-m)}_{i-m} \right) e^{(j)}, & i \geq j \end{cases} \leq \begin{cases} e^{(j-1)} e^{(j)}, & i \leq j-1 \\ x_{i-j+1} e^{(j)}, & i \geq j. \end{cases}$$

We want to prove that $G^{e^{(j)}} \geq G^{e^{(0)}}$ for all j , which is equivalent to $D_j \geq 0$ for all these j . Indeed,

$$2 \sum_{i=1}^{N+j-1} i e_i^{(j)} \leq 2 \sum_{i=1}^{j-1} i e^{(j-1)} e^{(j)} + 2 \sum_{i=1}^N (i+j-1) X_i e^{(j)} \leq e^{(j-1)} e^{(j)} (j-1)j + e^{(j)} (NG + N + 2j - 1)$$

and

$$G^{e^{(j)}} = \frac{1 - 2 \sum_{i=1}^{N+j-1} i e_i^{(j)} + (N+j-1) G^{e^{(j-1)}}}{N+j} \leq \frac{NG+j}{N+j}.$$

Finally

$$D_k = \sum_{j=1}^k \frac{1 - e^{(j-1)} e^{(j)} (j-1)j - e^{(j)} (NG + N + 2j - 1) - G^{e^{(j-1)}}}{N+j} \geq \frac{k - e^{(1)} e^{(2)} \frac{(k-1)k(k+1)}{3} - k e^{(1)} (NG + N + k) - \sum_{j=1}^k G^{e^{(j-1)}}}{N+k}.$$

Since $D_k \geq 0$,

$$k - k e^{(1)} \left(e^{(2)} \frac{(k^2-1)}{3} + NG + N + k \right) - \sum_{j=1}^k \frac{NG+j-1}{N+j-1} \geq 0 \Rightarrow \frac{k - \sum_{j=1}^k \frac{NG+j-1}{N+j-1}}{k \left(e^{(2)} \frac{(k^2-1)}{3} + NG + N + k \right)} \geq e^{(1)}.$$

The case of Uniswap on 15 August 2021 presents the *worst case* scenario for this boundary, with $N = 227,946$, $k = 34,934$, $G = 0.99184$ and $e^{(2)} = 6.42 \cdot 10^{-11}$. This results in an upper bound for $e^{(1)}$ of $1.55 \cdot 10^{-8}$, which is greater than the largest $e^{(1)}$. The largest $e^{(1)}$ is of scale 10^{-9} . It was recorded for SushiSwap. After excluding all the addresses that hold less than \$1 worth of tokens, the Gini coefficient decreases in every case. This justifies our claim from Section 3.2, in which we argued that our analysis yields an upper boundary with respect to decentrality. Furthermore,

$$\frac{D_k}{G} = \frac{G^{e^{(k)}} - G}{G} \leq \frac{\frac{NG+k}{N+k} - G}{G} = \frac{k(1-G)}{G(N+k)} = 0.09305$$

for $G = 0.58816$, which is the lowest value recorded for the ZRX distribution. Consequently, the Gini coefficient would increase by at most 9.305% if we included these addresses in our analysis.

With similar arguments, one can prove that $1 - \text{NSE}$ increases when addresses with low-value holdings are added. This is also apparent from a continuity argument. When adding an address with 0 holdings, NSE is divided through $\frac{\log(N+1)}{\log(N)} > 1$ (for $\text{NSE} = 0$, δ_{i_0} the case is even simpler) and hence $1 - \text{NSE}$ increases; so by continuity of $1 - \text{NSE}$, there is a threshold such that for adding holdings smaller than the threshold, $1 - \text{NSE}$ will increase.

References

- [1] M. Atzori, Blockchain Technology and Decentralized Governance: Is the State still Necessary? J. Govern. Regul. 6 (2017) 45–62, https://doi.org/10.22495/jgr_v6_i1_p5.
- [2] L. Shin, The Cryptopians: Idealism, Greed, Lies, and the Making of the First Big Cryptocurrency Craze, PublicAffairs, 2022.
- [3] D. Tapscott, A. Tapscott, Blockchain Revolution: How the Technology behind Bitcoin and other Cryptocurrencies is Changing the World. Penguin, 2018.
- [4] J.K. Brekke, K. Beecroft, F. Pick, The Dissensus Protocol: Governing Differences in Online Peer Communities, Front. Hum. Dynam. 3 (2021), <https://doi.org/10.3389/fhumd.2021.641731>.
- [5] J.G. Allen, Bodies without Organs: Law, Economics, and Decentralised Governance, Stanford Journal of Blockchain Law and Policy, 2021. <https://stanford-jblp.pubpub.org/pub/law-econ-decentralised-governance>.
- [6] D.A. Zetzsche, D.W. Arner, R.P. Buckley, Decentralized Finance, J. Financ. Regul. 6 (2020) 172–203, <https://doi.org/10.1093/jfr/fjaa010>.
- [7] M. Brennecke, T. Guggenberger, B. Schellinger, N. Urbach, The De-Central Bank in Decentralized Finance: A Case Study of MakerDAO, in: Proceedings of the 55th Hawaii International Conference on System Sciences, AIS, 2022, pp. 6073–6082, <https://doi.org/10.24251/HICSS.2022.737>.
- [8] R. Ozcan, Decentralized Finance, in: Financial Ecosystem and Strategy in the Digital Era: Global Approaches and New Opportunities, Springer, 2021, pp. 57–75, https://doi.org/10.1007/978-3-030-72624-9_4.
- [9] K. Pandian, D. Pfeiffer, S. Qian, Decentralized Finance, in: The Emerald Handbook on Cryptoassets: Investment Opportunities and Challenges, Emerald, 2023, pp. 141–156, <https://doi.org/10.1108/978-1-80455-320-620221010>.
- [10] T. Barbereau, R. Smethurst, O. Papageorgiou, A. Rieger, G. Fridgen, DeFi, not so Decentralized: The Measured Distribution of Voting Rights, in: Proceedings of the 55th Hawaii International Conference on System Sciences, AIS, 2022, pp. 6043–6052, <https://doi.org/10.24251/HICSS.2022.734>.
- [11] X. Sun, C. Stasinakis, Decentralization Illusion in DeFi: Evidence from MakerDAO. <https://arxiv.org/abs/2203.16612>, 2022.
- [12] H. Amler, L. Eckey, S. Faust, M. Kaiser, P. Sandner, B. Schlosser, DeFi-ning DeFi: Challenges & Pathway, 2021. <http://arxiv.org/abs/2101.05589>.

- [13] T. Katona, Decentralized Finance: The Possibilities of a Blockchain “Money Lego” System, *Financ. Econ. Rev.* 20 (2021) 74–102, <https://doi.org/10.3389/FER.20.1.74102>.
- [14] M. Nadler, F. Schär, Decentralized Finance, Centralized Ownership? An Iterative Mapping Process to Measure Protocol Token Distribution, 2020. <https://arxiv.org/abs/2012.09306>, arXiv:2012.09306.
- [15] K. Qin, L. Zhou, Y. Afonin, L. Lazzaretti, A. Gervais, CeFi vs. DeFi — Comparing Centralized to Decentralized Finance, 2021. <https://arxiv.org/abs/2106.08157>.
- [16] S. Hassan, P. De Filippi, Decentralized Autonomous Organization, *Internet Policy Re.* 10 (2021), <https://doi.org/10.14763/2021.2.1556>.
- [17] E. Warren, Warren Releases Response from SEC Chair Gensler Affirming Need to Regulate Cryptocurrency Exchanges and Protect Investors and Our Financial System, 2021. <https://www.warren.senate.gov/newsroom/press-releases/warren-releases-response-from-sec-chair-gensler-affirming-need-to-regulate-cryptocurrency-exchanges-and-protect-investors-and-our-financial-system>.
- [18] M. DeCambre, Crypto Is Reshaping the World Economy, 50 Years after Nixon Ended the Dollar’s Peg to Gold. Here’s How Some Are Playing it, 2021. <https://www.marketwatch.com/story/nixon-reshaped-the-world-economy-50-years-a-go-is-crypto-on-the-brink-of-doing-the-same-now-11628893012>.
- [19] Financial Action Task Force, Virtual Assets: Update of FATF Guidance for a Risk-Based Approach to Virtual Assets and VASPs – Sixth Draft – Public Consultation, 2021.
- [20] G. Gensler, Remarks before the Aspen Security Forum, 2021. <https://www.sec.gov/news/public-statement/gensler-aspen-security-forum-2021-08-03>.
- [21] G. Gensler, Testimony before the United States Senate Committee on Banking, Housing, and Urban Affairs, 2021. <https://www.sec.gov/news/testimony/gensler-2021-09-14>.
- [22] S. Brown, P.J. Toomey, A. Walch, J. Brito, M. Belcher, Cryptocurrencies: What are they Good for?, 2021. <https://www.banking.senate.gov/hearings/cryptocurrencies-what-are-they-good-for>.
- [23] E. Warren, Warren Urges FSOC to Use Authority to Address Cryptocurrency Market, Protect Consumers, and Mitigate Financial Stability Risks, 2021. <https://www.warren.senate.gov/download/fsoc-crypto-letter-07262021>.
- [24] T. Emmer, Opening Address: Exploring the Risks of Central Bank Digital Currencies, 2023. <https://www.cato.org/events/central-bank-digital-currency>.
- [25] J.R. Jensen, V. von Wachter, O. Ross, How Decentralized is the Governance of Blockchain-Based Finance? Empirical Evidence from Four Governance Token Distributions, 2021. <https://arxiv.org/abs/2102.10096>, arXiv:2102.10096.
- [26] L. Anker-Sørensen, D.A. Zetsche, From Centralized to Decentralized Finance: The Issue of ‘Fake-DeFi’, 2021. <https://ebi-europa.eu/publications/working-paper-series/>.
- [27] S. Aramonte, W. Huang, A. Schrimpf, DeFi Risks and the Decentralisation Illusion, 2021. https://www.bis.org/publ/qtrpdf/r_qt2112b.pdf.
- [28] P. De Filippi, B. Loveluck, The Invisible Politics of Bitcoin: Governance Crisis of a Decentralised Infrastructure, *Internet Policy Re.* 5 (2016), <https://doi.org/10.14763/2016.3.427>.
- [29] P. De Filippi, M. Mannan, W. Reijers, Blockchain as a Confidence Machine: The Problem of Trust & Challenges of Governance, *Technol. Soc.* 62 (2020), 101284, <https://doi.org/10.1016/j.techsoc.2020.101284>.
- [30] Q. DuPont, Experiments in Algorithmic Governance — A History and Ethnography of “The DAO,” a Failed Decentralized Autonomous Organization, in: *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance*, Routledge, 2017, pp. 157–177, <https://doi.org/10.4324/9781315211909-8>.
- [31] W. Reijers, I. Wuisman, M. Mannan, P. De Filippi, C. Wray, V. Rae-Looi, A. Cubillos Vélez, L. Orgad, Now the Code Runs Itself: On-Chain and Off-Chain Governance of Blockchain Technologies, *Topoi* 40 (2018) 821–831, <https://doi.org/10.1007/s11245-018-9626-5>.
- [32] K.M. Eisenhardt, Building Theories from Case Study Research, *Acad. Manag. Rev.* 14 (1989) 20, <https://doi.org/10.5465/amr.1989.4308385>.
- [33] K.M. Eisenhardt, M.E. Graebner, Theory Building from Cases: Opportunities and Challenges, *Acad. Manag. J.* 50 (2007) 25–32, <https://doi.org/10.5465/amj.2007.24160888>.
- [34] R.K. Yin, *Case Study Research and Applications: Design and Methods*, sixth ed., Sage, 2018.
- [35] C. Brooke, Critical Research in Information Systems: Issue 1, *J. Inf. Technol.* 17 (2002) 45–47, <https://doi.org/10.1080/02683960210154139>.
- [36] D. Howcroft, E.M. Trauth, Choosing Critical IS Research, in: D. Howcroft, E. M. Trauth (Eds.), *Handbook of Critical Information Systems Research: Theory and Application*, Edward Elgar, 2005, pp. 1–15, <https://doi.org/10.4337/9781845426743.00007>.
- [37] B.C. Stahl, *Information Systems: Critical Perspectives*, Routledge, 2008.
- [38] K. Becker, Blockchain Matters – Lex Cryptographia and the Displacement of Legal Symbolics and Imaginaries. Law and Critique, 2022, <https://doi.org/10.1007/s10978-021-09317-8>.
- [39] B. Bodó, J.K. Brekke, J.H. Hoepman, Decentralisation: A Multidisciplinary Perspective, *Internet Policy Re.* 10 (2021) 1–21, <https://doi.org/10.14763/2021.2.1563>.
- [40] N.R. Hassan, J. Mingers, B. Stahl, Philosophy and Information Systems: Where are We and Where should we Go? *Eur. J. Inf. Syst.* 27 (2018) 263–277, <https://doi.org/10.1080/0960085X.2018.1470776>.
- [41] F. Rowe, Being Critical is Good, but Better with Philosophy! From Digital Transformation and Values to the Future of IS Research, *Eur. J. Inf. Syst.* 27 (2018) 380–393, <https://doi.org/10.1080/0960085X.2018.1471789>.
- [42] G.R.F. Ferrari (Ed.), *Plato: The Republic*, Cambridge University Press, 2000.
- [43] R. Morrison, N.C.H.L. Mazey, S.C. Wingreen, The DAO Controversy: The Case for a New Species of Corporate Governance? *Front. Blockchain* 3 (2020) <https://doi.org/10.3389/fbloc.2020.00025>.
- [44] S. Davidson, P. De Filippi, J. Potts, Blockchains and the Economic Institutions of Capitalism, *J. Inst. Econ.* 14 (2018) 639–658, <https://doi.org/10.1017/S1744137417000200>.
- [45] D. Rozas, A. Tenorio-Pornés, S. Díaz-Molina, S. Hassan, When Ostrom Meets Blockchain: Exploring the Potentials of Blockchain for Commons Governance, *Sage Open* 11 (2021), <https://doi.org/10.1177/2158244021100252>.
- [46] C. Santana, L. Albareda, Blockchain and the Emergence of Decentralized Autonomous Organizations (DAOs): An Integrative Model and Research Agenda, *Technol. Forecast. Soc. Change* 182 (2022), <https://doi.org/10.1016/j.techfore.2022.121806>.
- [47] J.P. Vergne, Decentralized vs. Distributed Organization: Blockchain, Machine Learning and the Future of the Digital Platform, *Organ. Theor.* 1 (2020), <https://doi.org/10.1177/2631787720977052>.
- [48] P. Dylan-Ennis, D. Kavanagh, L. Araujo, The Dynamic Imaginaries of the Ethereum Project, *Econ. Soc.* 52 (2023) 87–109, <https://doi.org/10.1080/03085147.2022.2131280>.
- [49] V. Buterin, The Meaning of Decentralization, 2017. <https://medium.com/@VitalikButerin/the-meaning-of-decentralization-a0c92b76a274>.
- [50] N. Schneider, Decentralization: An Incomplete Ambition, *J. Cult. Econ.* 12 (2019) 265–285, <https://doi.org/10.1080/17530350.2019.1589553>.
- [51] S.P. Gochhayat, S. Shetty, R. Mukkamala, P. Foytik, G.A. Kamhoua, L. Njilla, Measuring Decentrality in Blockchain Based Systems, *IEEE Access* 8 (2020) 178372–178390, <https://doi.org/10.1109/ACCESS.2020.3026577>.
- [52] V. Kostakis, C. Giotitis, The (a)political economy of Bitcoin, *tripleC* 12 (2014), <https://doi.org/10.31269/tripleC.v12i2.606>.
- [53] S. Nagarajan, Bitcoin hit the \$1 trillion market cap milestone twice as fast as Amazon and three times faster than Apple, data shows. <https://www.businessinsider.in/stock-market/news/bitcoin-hit-the-1-trillion-market-cap-milestone-twice-as-fast-as-amazon-and-three-times-faster-than-apple-data-shows/articleshow/82064390.cms>, 2021.
- [54] S. Nakamoto, Bitcoin: A Peer-To-Peer Electronic Cash System, 2008. <https://bitcoin.org/bitcoin.pdf>.
- [55] D. Drummer, D. Neumann, Is Code Law? Current Legal and Technical Adoption Issues and Remedies for Blockchain-Enabled Smart Contracts, *J. Inf. Technol.* 35 (2020) 337–360, <https://doi.org/10.1177/0268396220924669>.
- [56] H. Nabilou, How to Regulate Bitcoin? Decentralized Regulation for a Decentralized Cryptocurrency, *Int. J. Law Info Technol.* 27 (2019) 31, <https://doi.org/10.1093/ijlit/eaz008>.
- [57] H. Nabilou, Bitcoin Governance as a Decentralized Financial Market Infrastructure, Stanford J. Blockchain Law Policy (2021). <https://stanford-jblp.pubpub.org/pub/bitcoin-governance>.
- [58] M. Gupta, P. Gupta, Gini Coefficient based Wealth Distribution in the Bitcoin Network: A Case Study, in: R. Sharma, A. Mantri, S. Dua (Eds.), *Computing, Analytics and Networks*, vol. 805, Springer, 2018, pp. 192–202, https://doi.org/10.1007/978-981-13-0755-3_15.
- [59] I. Makarov, A. Schoar, Blockchain Analysis of the Bitcoin Market, 2021. <http://www.nber.org/papers/w29396>.
- [60] J. Parkin, The Senatorial Governance of Bitcoin: Making (de)centralized Money, *Econ. Soc.* 48 (2019) 463–487, <https://doi.org/10.1080/03085147.2019.1678262>.
- [61] L. Ante, Smart Contracts on the Blockchain — A Bibliometric Analysis and Review, *Telematics Inf.* 57 (2021), 101519, <https://doi.org/10.1016/j.tele.2020.101519>.
- [62] A. Groce, J. Feist, G. Grieco, M. Colburn, What are the Actual Flaws in Important Smart Contracts (and how can we Find Them)? in: J. Bonneau, N. Heninger (Eds.), *Financial Cryptography and Data Security* Springer, 2020, pp. 634–653, https://doi.org/10.1007/978-3-030-51280-4_34.
- [63] G.A. Oliva, A.E. Hassan, Z.M. Jiang, An Exploratory Study of Smart Contracts in the Ethereum Blockchain Platform, *Empir. Software Eng.* 25 (2020) 1864–1904, <https://doi.org/10.1007/s10664-019-09796-5>.
- [64] A. Pazaitis, P. De Filippi, V. Kostakis, Blockchain and Value Systems in the Sharing Economy: The Illustrative Case of Backfeed, *Technol. Forecast. Soc. Change* 125 (2017) 105–115, <https://doi.org/10.1016/j.techfore.2017.05.025>.
- [65] M. Casey, P. Vigna, In Blockchain We Trust, vol. 18, MIT Technology Review, 2018. <https://www.technologyreview.com/2018/04/09/3066/in-blockchain-we-trust/>.
- [66] G. Gensler, Response from SEC Chair Gensler Affirming Need to Regulate Cryptocurrency Exchanges and Protect Investors and Our Financial System, 2021. https://www.warren.senate.gov/imo/media/doc/gensler_response_to_warren_cryptocurrency_exchanges.pdf.
- [67] E. Warren, Warren Asks SEC Chair Gensler about Risks Posed by Cryptocurrency Exchanges, 2021. <https://www.warren.senate.gov/oversight/letters/warren-asks-sec-chair-gensler-about-risks-posed-by-cryptocurrency-exchanges>.
- [68] FinTech Innovation Hub, FinTech Innovation Hub Report, second ed., 2021. Technical Report 2. Financial Services Agency, https://www.fsa.go.jp/policy/bgin/FIH_Report_2nd_ja.pdf.
- [69] D.M. Berkovitz, Climate Change and Decentralized Finance: New Challenges for the CFTC, 2021. <https://www.cftc.gov/PressRoom/SpeechesTestimony/opaberkovitz7>.
- [70] L. Oliveira, L. Zavolokina, I. Bauer, G. Schwabe, To Token or not to Token: Tools for Understanding Blockchain Tokens, in: *Proceedings of the 39th International Conference of Information Systems*, 2018, pp. 1–17.

- [71] W.M. Klimon, *Beyond the Board: Alternatives in Nonprofit Corporate Governance*, Harvard Bus. Law Rev. 9 (2018).
- [72] N. Arnosti, S.M. Weinberg, Bitcoin: A Natural Oligopoly, *Manag. Sci.* (2022), <https://doi.org/10.1287/mnsc.2021.4095>.
- [73] Q. Lin, C. Li, X. Zhao, X. Chen, Measuring Decentralization in Bitcoin and Ethereum using Multiple Metrics and Granularities, in: 37th International Conference on Data Engineering Workshops, IEEE, 2021, pp. 80–87, <https://doi.org/10.1109/ICDEW53142.2021.00022>.
- [74] B. Kaplan, D. Duchon, Combining Qualitative and Quantitative Methods in Information Systems Research: A Case Study, *MIS Q.* 12 (1988) 571, <https://doi.org/10.2307/249133>.
- [75] M.D. Myers, M. Newman, The Qualitative Interview in IS Research: Examining the Craft, *Inf. Organ.* 17 (2007) 2–26, <https://doi.org/10.1016/j.infoandorg.2006.11.001>.
- [76] N.B. Schirmacher, J.R. Jensen, M. Avital, Token-Centric Work Practices in Fluid Organizations: The Cases of Yearn and MakerDAO, in: Proceedings of the 42nd International Conference on Information Systems, 2021, pp. 1–17.
- [77] G.G. Gable, Integrating Case Study and Survey Research Methods: An Example in Information Systems, *Eur. J. Inf. Syst.* 3 (1994) 15, <https://doi.org/10.1057/ejis.1994.12>.
- [78] J.E. de Azevedo Sousa, V. Oliveira, J. Valadares, G. Dias Gonçalves, S. Moraes Villela, H. Soares Bernardino, A. Borges Vieira, An Analysis of the Fees and Pending Time Correlation in Ethereum, *Int. J. Netw. Manag.* 31 (2021) e2113, <https://doi.org/10.1002/nem.2113>.
- [79] Y. Sovbetov, Factors Influencing Cryptocurrency Prices: Evidence from Bitcoin, Ethereum, Dash, Litecoin, and Monero, *J. Econ. Financ. Anal.* 2 (2018). <https://ojs.tripaledu.com/index.php/jefa/article/view/360>.
- [80] M. Spain, S. Foley, V. Gramoli, The Impact of Ethereum Throughput and Fees on Transaction Latency during ICOs, in: International Conference on Blockchain Economics, Security and Protocols, 2020, pp. 9:1–9:15, <https://doi.org/10.4230/OASICS.Tokenomics.2019.9>.
- [81] L. Nuzzi, Bitcoin Miners Are Escaping China, 2021. <https://medium.com/coinmetrics/bitcoin-miners-are-escaping-china-d3937e8f018c>.
- [82] K. Babel, P. Daian, M. Kelkar, A. Juels, Clockwork Finance: Automated Analysis of Economic Security in Smart Contracts, 2021. <http://arxiv.org/abs/2109.04347>.
- [83] W. Chan, A. Olmsted, Ethereum Transaction Graph Analysis, in: 12th International Conference for Internet Technology and Secured Transactions, 2017, pp. 498–500, <https://doi.org/10.23919/ICITST.2017.8356459>.
- [84] X. Xueshuo, W. Jiming, Y. Junyi, F. Yaozheng, L. Ye, L. Tao, W. Guiling, AWAP: Adaptive Weighted Attribute Propagation Enhanced Community Detection Model for Bitcoin De-anonymization, *Appl. Soft Comput.* 109 (2021), 107507, <https://doi.org/10.1016/j.asoc.2021.107507>.
- [85] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, A. Juels, Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability, in: Symposium on Security and Privacy, IEEE, 2020, pp. 910–927, <https://doi.org/10.1109/SP40000.2020.00040>.
- [86] C.F. Torres, R. Camino, R. State, Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain, in: 30th USENIX Security Symposium, 2021, pp. 1343–1359.
- [87] J. Liang, L. Li, D. Zeng, Evolutionary Dynamics of Cryptocurrency Transaction Networks: An Empirical Study, *PLoS One* 13 (2018), <https://doi.org/10.1371/journal.pone.0202202>.
- [88] V. Saraph, M. Herlihy, An Empirical Study of Speculative Concurrency in Ethereum Smart Contracts, in: International Conference on Blockchain Economics, Security and Protocols, 2020, <https://doi.org/10.4230/OASICS.Tokenomics.2019.4>.
- [89] H. Adams, Yesterday Binance Delegated 13M UNI from its Books, 2022. <https://twitter.com/haydenzadams/status/1582714218323017729>.
- [90] V. Buterin, Moving Beyond Coin Voting Governance, 2021. <https://vitalik.ca/general/2021/08/16/voting3.html>.
- [91] M. Miles, A. Huberman, J. Saldaña, *Qualitative Data Analysis: An Expanded Sourcebook*, second ed., Sage, 1994.
- [92] P. Mayring, *Qualitative Content Analysis: Theoretical Foundation, Basic Procedures and Software Solution*, Klagenfurt, 2014.
- [93] H.K. Klein, M.D. Myers, A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems, *MIS Q.* 23 (1999) 67–93, <https://doi.org/10.2307/249410>.
- [94] S.H. Cha, Comprehensive Survey on Distance/Similarity Measures between Probability Density, *Int. J. Math. Models Methods Appl. Sci.* 1 (2007).
- [95] J. Chung, P. Kannappan, C. Ng, P. Sahoo, Measures of Distance between Probability Distributions, *J. Math. Anal. Appl.* 138 (1989) 280–292, [https://doi.org/10.1016/0022-247X\(89\)90335-1](https://doi.org/10.1016/0022-247X(89)90335-1).
- [96] H. Dalton, The Measurement of the Inequality of Incomes, *Econ. J.* 30 (1920) 348, <https://doi.org/10.2307/2223525>.
- [97] L. Kostal, P. Lansky, O. Pokora, Measures of Statistical Dispersion based on Shannon and Fisher Information Concepts, *Inf. Sci.* 235 (2013) 214–223, <https://doi.org/10.1016/j.ins.2013.02.023>.
- [98] L. Kostal, O. Pokora, Nonparametric Estimation of Information-Based Measures of Statistical Dispersion, *Entropy* 14 (2012) 1221–1233, <https://doi.org/10.3390/e14071221>.
- [99] F. Liese, I. Vajda, On Divergences and Informations in Statistics and Information Theory, *IEEE Trans. Inf. Theor.* 52 (2006) 4394–4412, <https://doi.org/10.1109/TIT.2006.881731>.
- [100] I. Sason, S. Verdú, f -divergence inequalities, *IEEE Trans. Inf. Theor.* 62 (2016) 5973–6006, <https://doi.org/10.1109/TIT.2016.2603151>.
- [101] D. Freedman, R. Pisani, R. Purves, *Statistics*, fourth ed., Norton, 2007.
- [102] C. Gini, *Variabilità e mutabilità*, 1912.
- [103] P.P. Graczyk, Gini Coefficient: A New Way to Express Selectivity of Kinase Inhibitors against a Family of Kinases, *J. Med. Chem.* 50 (2007) 5773–5779, <https://doi.org/10.1021/jm070562u>.
- [104] W. Halffman, L. Leydesdorff, Is Inequality among Universities Increasing? Gini Coefficients and the Elusive Rise of Elite Universities, *Minerva* 48 (2010) 55–72, <https://doi.org/10.1007/s11024-010-9141-3>.
- [105] V. Buterin, Against Overuse of the Gini Coefficient, 2021. <https://vitalik.ca/general/2021/07/29/gini.html>.
- [106] C.E. Shannon, A Mathematical Theory of Communication, *Bell System Technical Journal* 27 (1948) 379–423, <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
- [107] J. Lin, Divergence Measures based on the Shannon Entropy, *IEEE Trans. Inf. Theor.* 37 (1991) 145–151, <https://doi.org/10.1109/18.61115>.
- [108] L.W. Cong, Y. Li, N. Wang, Token-Based Platform Finance, *J. Financ. Econ.* (2021), <https://doi.org/10.1016/j.jfineco.2021.10.002>.
- [109] R. Wandmacher, *Tokenomics*, in: *Cryptofinance and Mechanisms of Exchange: The Making of Virtual Currency*, Springer, 2019, pp. 113–123.
- [110] C. Wan, In a first, MakerDAO protocol to auction MKR tokens to cover \$4M bad debt, 2020. <https://www.theblock.co/post/58606/in-a-first-makerdao-protocol-to-auction-mkr-tokens-to-cover-4m-bad-debt>.
- [111] adrienbe, What Was MakerDAO's MKR Initial Supply?, 2019. https://www.reddit.com/r/MakerDAO/comments/apuu9e/what_was_makerdaos_mkr_initial_supply/.
- [112] E. Buzzio, M. Clarendon, Democracy Decentralised: Voting, Governance and Transparency, dGen, 2020.
- [113] X. Sun, X. Chen, C. Stasinakis, G. Serpinis, Voter Coalitions and Democracy in Decentralized Finance: Evidence from MakerDAO, arXiv, 2023, <https://doi.org/10.48550/ARXIV.2210.11203>.
- [114] V. Grover, K. Lyytinen, New State of Play in Information Systems Research, *MIS Q.* 39 (2015) 271–296, <https://doi.org/10.25300/MISQ/2015/39.2.01>.
- [115] OECD, Ownership and Control of Ships, OECD Directorate for Science, Technology, and Industry, 2003. Technical Report, [https://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DSTI/DOT/MTC\(2003\)61/R/REV1&docLanguage=En](https://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DSTI/DOT/MTC(2003)61/R/REV1&docLanguage=En).
- [116] C. Pacini, N.F. Stowell, Panama Papers and the Abuse of Shell Entities, in: H. K. Baker, L. Purda-Heeler, S. Saadi (Eds.), *Corporate Fraud Exposed*, Emerald, 2020, pp. 361–382, <https://doi.org/10.1108/978-1-78973-417-120201023>.
- [117] K. Stastna, Anonymous Paper Shares Helped the Wealthy Hide Billions, Panama Papers Show, 2016. <https://www.cbc.ca/news/business/panama-papers-be-arer-shares-1.3525434>.
- [118] E.v.d. Does de Willebois, J. Sharman, R. Harrison, J.W. Park, E. Halter, The Puppet Masters: How the Corrupt Use Legal Structures to Hide Stolen Assets and what to Do about it, World Bank Publications, 2011. <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/784961468152973030/The-puppet-masters-how-the-corrupt-use-legal-structures-to-hide-stolen-assets-and-what-to-do-about-it>.
- [119] K. Wiehe, Oh, Snap! Do Multi-Class Offerings Signal the Decline of Shareholder Democracy and the Normalization of Founder Primacy? *Va Law Bus. Rev.* 12 (2017).
- [120] G. Hayden, M.T. Bodie, Shareholder Democracy and the Curious Turn toward Poop Primacy, *William Mary Law Rev.* 51 (2009) 2071–2121.
- [121] G.M. Hayden, M.T. Bodie, The Corporation Reborn: From Shareholder Primacy to Shared Governance, *Boston Coll. Law Rev.* 61 (2020) 2419–2485.
- [122] E. Kamonjoh, Controlled Companies in the Standard & Poor's 1500: A Follow-Up Review of Performance and Risk, 2016. <https://bpb-us-w2.wpmucdn.com/sites.udel.edu/dist/8/12944/files/2022/08/Controlled-Companies-IRRCI-2015-FINAL-3-16-16.pdf>.
- [123] V. Pareto, *A Treatise on General Sociology* vols. 3–4, Dover, 1963 [1916].
- [124] V. Pareto, *The Rise and Fall of Elites: An Application of Theoretical Sociology*, Routledge, 2017 [1901].
- [125] A. Tocqueville, *The Ancien Régime and the French Revolution*, Penguin, 2008 [1856].
- [126] C. Fordahl, Lions and Foxes: Revisiting Pareto's Bestiary for the Age of Late Pluto Democracy, *Distinktion: J. Soc. Theor.* 21 (2020) 316–333, <https://doi.org/10.1080/1600910X.2020.1752274>.
- [127] A. Wiedand, L. Heppding, Centralized and Decentralized Finance: Coexistence or Convergence? in: T. Walker, E. Nikbakht, M. Kooli (Eds.), *The Fintech Disruption: How Financial Innovation is Transforming the Banking Industry* Springer, 2023, pp. 11–51, https://doi.org/10.1007/978-3-031-23069-1_2.
- [128] A. Cronje, Time to Retire “Decentralized Finance”, 2021. <https://twitter.com/AndreCronjeTech/status/1468273605502980106>.
- [129] R. Selkis, Crypto Theses for 2022, 2021. <https://messari.io/pdf/messari-report-crypto-theses-for-2022.pdf>.
- [130] K. Warwick, S. Kulechov, T. Chitra, W. Villanueva, DeFi: The Future of Finance, 2021. <https://www.youtube.com/watch?v=0GoACn1RPGU>.
- [131] Rekt, Curve Wars, 2021. <https://rekt.news/curve-wars/>.
- [132] Rekt, SushiSwap Scandal, 2021. <https://rekt.news/sushiswap-scandal>.
- [133] B. McGing (Ed.), *Polybius: The Histories*, Oxford University Press, 2010.
- [134] D. Kerr, M. Jennings, A Legal Framework for Decentralized Autonomous Organizations, 2021. <https://a16z.com/wp-content/uploads/2021/10/DAO-Legal-Framework-Jennings-Kerr10.19.21-Final.pdf>.
- [135] M. Boiron, Legal Framework: Non-U.S. Trusts in Decentralized Autonomous Organizations, 2022. <https://dydx.foundation/blog/en/legal-framework-no-n-us-trusts-in-daos>.

- [136] R.B. Lamb, Utah Passes Innovative DAO Legislation, 2023. The National Law Review 13, <https://www.natlawreview.com/article/utah-passes-innovative-dao-legislation>.
- [137] C. Brummer, R. Seira, Legal Wrappers and DAOs, SSRN, 2022, <https://doi.org/10.2139/ssrn.4123737>.
- [138] S. Inoua, Beware the Gini Index! A New Inequality Measure, 2021. <https://arxiv.org/ftp/arxiv/papers/2110/2110.01741.pdf>.
- [139] 0x, 0x Governance Update #1, 2021. <https://blog.0x.org/0x-governance-update-1/>.

3.3. Research Paper 3: *DeFi, not so Decentralized: The Measured Distribution of Voting Rights*

DeFi, Not So Decentralized: The Measured Distribution of Voting Rights

Barbereau, Tom
University of Luxembourg
tom.barbereau@uni.lu
ORCID: 0000-0002-8554-0991

Smethurst, Reilly
University of Luxembourg
reilly.smethurst@uni.lu
ORCID: 0000-0002-6191-475X

Papageorgiou, Orestis
University of Luxembourg
orestis.papageorgiou@uni.lu
ORCID: 0000-0002-2245-5079

Rieger, Alexander
University of Luxembourg
alexander.rieger@uni.lu
ORCID: 0000-0001-7996-4678

Fridgen, Gilbert
University of Luxembourg
gilbert.fridgen@uni.lu
ORCID: 0000-0001-7037-4807

Abstract

Bitcoin and Ethereum are frequently promoted as decentralized, but developers and academics question their actual decentralization. This motivates further experiments with public permissionless blockchains to achieve decentralization along technical, economic, and political lines. The distribution of tokenized voting rights aims for political decentralization. Tokenized voting rights achieved notoriety within the nascent field of decentralized finance (DeFi) in 2020. As an alternative to centralized crypto-asset exchanges and lending platforms (owned by companies like Coinbase and Celsius), DeFi developers typically create non-custodial projects that are not majority-owned or managed by legal entities. Holders of tokenized voting rights can instead govern DeFi projects. To scrutinize DeFi's distributed governance strategies, we conducted a multiple-case study of non-custodial, Ethereum-based DeFi projects: Uniswap, Maker, SushiSwap, Yearn Finance, and UMA. Our findings are novel and surprising: quantitative evaluations of DeFi's distributed governance strategies reveal a failure to achieve political decentralization.

1. Introduction

In January 2021, Uniswap became the first non-custodial crypto-asset exchange to reach USD \$100 billion trading volume – a significant event for the nascent field of *decentralized finance* (DeFi). DeFi projects are typically governed by online communities of anonymous stakeholders rather than companies or public institutions. Instead of registered shareholders' voting rights, many DeFi projects involve tokenized

voting rights, and they offer unregulated financial services that are based on crypto-assets instead of sovereign currencies. For DeFi, the meaning of 'decentralized' thus entails independence from regulated companies, investor registration processes, and central banks [1, 2, 3]. The descriptor 'decentralized' does not necessarily denote a uniform distribution of capital or voting power [4]. Political decentralization is anything but guaranteed.

In 1856, the historian Alexis de Tocqueville [5] investigated the fall of the French aristocracy and the subsequent dissolution of power. Contrary to the revolutionaries' expectations, the abolition of feudalism in France at the end of the 18th century did not effectively limit the power of all central authorities. Authority putatively transferred from the aristocracy to the people, yet there soon emerged a powerful bureaucratic system known as the *Première République*. "When a nation abolishes aristocracy," De Tocqueville concluded, "centralization follows as a matter of course: everything tends toward unity of power, and it requires no small contrivance to maintain divisions of authority."

De Tocqueville's maxim retains an air of truth in the era of putatively decentralized networks like Bitcoin and Ethereum. The ideals of decentralization that inspired the invention of Bitcoin and Ethereum are undermined by tendencies towards socio-political centralization [6, 7, 8, 9, 10] and technical centralization [11, 12]. Developers and active members of crypto-asset communities are aware of these tendencies; hence they propose novel, experimental strategies to achieve decentralization. DeFi projects are the most prominent decentralization experiments from late 2020 [2]. Many entail strategies to distribute tradable *voting rights tokens* (also known as *governance tokens*) for the sake of political decentralization [1].

DeFi's voting rights tokens are speculative crypto-assets, not backed by anything physical [13]. They enable token-holders to iteratively revise a project's rules, parameters, and features. A broad range of DeFi projects issue voting rights tokens, from non-custodial crypto-asset exchanges to derivatives platforms [1]. Maker – a non-custodial crypto-asset lending platform – initiated the trend with its MKR token.

Uniswap is the largest non-custodial, Ethereum-based exchange. Uniswap issued its voting rights token (UNI) in September 2020 via a retrospective *airdrop*. The airdrop rewarded anonymous wallet owners that interacted with the Uniswap exchange in the past (by swapping tokens or providing liquidity). UNI token-holders today govern the Uniswap exchange. They have the right to vote on proposals that pertain to treasury funds and changes to the exchange's rules; and, if they wish, they can sell their UNI tokens on various crypto-asset exchanges, or they can purchase more UNI tokens to increase their voting power. As of 9 June 2021, the UNI token is in the top 10 list of crypto-assets (sorted by market capitalization), and it sells for more than four times its initial price of \$5.

While academic research has addressed crypto-assets, Bitcoin, Ethereum, algorithmic governance, and the various meanings of 'decentralization,' non-custodial DeFi projects are comparatively uncharted territory, and tokenized voting rights issued by DeFi projects are an under-researched topic. It is not yet known if De Tocqueville's pessimism about decentralization remains valid for DeFi's distributed voting rights; hence we formulated the following research questions. How did major, Ethereum-based DeFi projects distribute voting rights? What are the outcomes of the different distribution strategies?

We conducted an exploratory, multiple-case study in response to these questions [14]. We selected as cases five non-custodial, Ethereum-based DeFi projects that include tradable voting rights tokens: Uniswap, Maker, SushiSwap, Yearn Finance, and UMA. Our case selection covers four major DeFi categories [1]: lending (Maker), exchanges (Uniswap and SushiSwap), assets (Yearn Finance), and derivatives (UMA). We selected two exchanges instead of just one exchange, because SushiSwap essentially began as a clone of Uniswap.

We followed a *mixed methods approach* [15], which is recognized for its ability to discover the unexpected. First, we qualitatively analyzed project documentation that details the properties and the distribution strategy for each case's tokenized voting rights. We then quantitatively examined the "physical" artifact [14], based on methods proposed by Gochhayat et al. [11].

We specifically used metrics that quantify the degree of uniformity for each case's distribution of voting rights tokens.

2. Background: The promise of 'decentralization'

Bitcoin is the fastest asset to reach a \$1 trillion market capitalization [16]. It achieved this feat with no legal entity attached to it as a manager or a majority owner. Bitcoin's pseudonymous creator, Satoshi Nakamoto, established a finite supply (21 million units of BTC) that cannot be altered by a central bank or any other financial authority [17]. Bitcoin transactions are not registered on servers owned by a single bank, financial institution, or consortium; the Bitcoin ledger is instead distributed across a global computing network that consists of voluntary participants [18]. In comparison with sovereign currencies, physical commodities, and other tangible assets, Bitcoin is difficult to regulate or subject to capital controls [19, 20].

Social scientists often interpret Bitcoin as a symptom of widespread distrust in both financial and political institutions following the 2007-2008 global financial crisis [6, 8]. The raw hex data from Bitcoin's *genesis block*, mined by Nakamoto on 3 January 2009, contains within it The Times' daily headline: "Chancellor on brink of second bailout for banks." Sociologists class Bitcoin's early adopters as fundamentalist libertarians, many of whom conceive 'decentralization' as a set of economic and political ideals [10, 21]. De Filippi & Loveluck [6] traced the Bitcoin community's conception of 'decentralization' to "Friedrich Hayek's and Milton Friedman's ambition to end the monopoly of nation-states."

The ideals of decentralization enjoy a broad appeal among libertarians. For the libertarian Right, decentralization's principal mission is to end the oligopoly of banks and tax authorities, and to create a market without costly intermediaries and powerful regulators. For the libertarian Left, decentralization's principal mission is to abolish centralized monetary policy and State capital controls [6, 22]. There is no necessary link, however, between the Bitcoin community's ideals of decentralization and an actualized, uniform distribution of capital or voting power within the Bitcoin network [9, 23]. As Kostakis & Giotitsas [4] put it, "in theory you have equipotential individuals (that is, everyone can potentially participate in a project), but in practice what one gets is concentrated capital and centralized governance."

2.1. Bitcoin's re-centralization and unregulated store of value

On the one hand, Bitcoin exhibits non-uniform distributions of capital, governance authority, and computing power, which together imply a failure to achieve decentralization. On the other hand, the Bitcoin network poses a challenge to centralized legal authorities and regulators, which means it is indeed 'decentralized' in the anti-establishment, libertarian sense.

According to the Gini index, capital (BTC) is distributed less evenly in the Bitcoin network than it is in economies like South Africa or Brazil [23, 24]. The 2015 *block-size* debate and subsequent revision to the Bitcoin protocol rules demonstrated the concentrated governance power held by the Bitcoin Core Developers and the Lead Developer [6, 9]. In addition, costly mining hardware (*application-specific integrated circuits*, or ASICs), owned by a sophisticated minority of Bitcoin miners, offers evidence of concentrated computing power [11, 12, 25]. Today, one private company in China, Bitmain, dominates the global ASIC production industry [26], and it owns the Bitcoin mining pools AntPool and BTC.com. These pools together comprise over 20% of the Bitcoin network's hash rate (Figure 1 – data retrieved from btc.com on 7 June 2021).




Figure 1. Mining pools dominate the hash rate

The Bitcoin network does, however, successfully manifest this libertarian ideal of decentralization: it is uncontrollable or unregulatable. Since the global Bitcoin network is diffuse and not wholly grounded in any specific jurisdiction, it is a challenge for regulators [27, 28, 29]. The Bitcoin network is also a challenge to regulate because there is no legal entity attached to it as a manager or a majority owner. In legal terms, Bitcoin can thus be classified as an "unincorporated distributed

ledger system" [3]. Lawmakers also struggle to define the legal status of Bitcoin's store of value (BTC) [19]; hence regulatory responses to BTC are incoherent and autarkic [27].

2.2. Ethereum's unregulated contracts, non-labile organizations, and centralized traits

Ethereum introduced its own, unregulated store of value (ETH) in 2015; but, more importantly for the present study, Ethereum also introduced unregulated chain code, referred to as *smart contracts* [30]. Smart contracts do not require an intermediary authority to execute them – another variation on the theme of 'decentralization' [10]. Smart contracts can bind together members of *decentralized autonomous organizations* (DAOs), and they can also assign voting rights [31].

A DAO is a protocol that 'lives' on the Internet and consists of formalized rules [31]. If truly decentralized, no single entity (including a DAO's creator) can alter, reverse, or control a DAO once it is deployed on the Internet. A DAO is virtually unregulatable and non-labile: it simply executes its pre-determined set of 'if-then' rules [2, 32]. If a DeFi project delegates management duties to an online community instead of a legal entity, then it is effectively a DAO [1, 2].

Like Bitcoin, the Ethereum network is sometimes criticized for its compromised or failed decentralization [11, 25]. Researchers cast doubt upon Ethereum's ability to abolish corporate intermediaries and usher in a "decentralized world" [25, 33]. In 2016, Ethereum's core developers exercised autocratic powers [8, 34, 35] and established a new *fork* of the blockchain, in response to a disastrous hack. De Tocqueville's maxim thus applies to Ethereum's governance as well: "centralization follows as a matter of course" [5].

2.3. DeFi projects not bound to regulated companies: centralized, once again?

The year 2020 witnessed multi-billion-dollar trading volumes for Ethereum-based financial services projects that do not rely on regulated, custodial intermediaries such as banks, brokerages, or company-governed exchanges [1, 3, 36]. These projects – commonly categorized as DeFi – include non-custodial crypto-asset exchanges, lending platforms, and derivatives platforms [1]. DeFi projects are often governed by a non-labile DAO rather than a legal entity [2, 3]; hence they are distinct from centralized crypto-asset

exchanges and lending platforms owned by regulated companies like Coinbase and Celsius.

DeFi exchanges like Uniswap and SushiSwap allow participants to vote for proposed changes to rules, parameters, and features, so long as the participants purchase or else earn voting rights tokens (via contributions to the exchange) [13]. This phenomenon can be interpreted in at least two distinct ways: (a) governance is decentralized, because decision-making power is distributed among a community of token-holders (and therefore not held by a centralized company or institution), and (b) governance is timocratic, since voting rights are exclusively attributed to ‘property owners’ (in this case, token-holders). The second interpretation is especially pertinent to SushiSwap. On 7 September 2020, a small group of individuals held the majority of voting rights tokens, and they successfully seized control of the exchange [37]. De Tocqueville’s maxim, once again, rings true.

A DeFi project named Yearn Finance pre-empted outcomes like this, hence it deployed a refined *fair launch* strategy. Fair launch strategies vary from case to case; but in general, a fair launch does not allocate a large portion of *pre-mined* (spontaneously generated) tokens to developers or early investors and advisors [38]. A fair launch is premised on the following assumption: if a significant portion of voting rights tokens are distributed to privileged insiders, then it is difficult to describe a project’s governance as decentralized. The fair launch is a sub-topic of Ethereum *tokenomics* and token distribution strategies [13, 39, 40].

Yearn Finance’s fair launch strategy, deployed in July 2020, distributed none of the project’s voting rights tokens (YFI) to developers and early investors. Yearn Finance’s fair launch did not abolish timocracy, since only YFI token-holders possess voting rights. What its fair launch did, specifically, is remove the distinction between ‘insiders’ (that are allocated voting rights by default) and ‘outsiders’ (that must either purchase voting rights on exchanges, provide liquidity, or contribute work to the project). Other DeFi projects, such as Beefy Finance, Harvest Finance, and YAM, later deployed fair launch strategies as well.

DeFi’s various fair launch strategies indicate new avenues for research into algorithmic governance and political decentralization, set against a backdrop of timocracy, re-centralization, and concentrated powers.

3. Methods

To understand DeFi’s tokenized voting rights and to evaluate the success of various distribution strategies, we conducted an exploratory, multiple-case study [14] with a mixed methods approach [15]. For the sake of consistent quantitative data extraction, we only chose DeFi projects that utilize Ethereum, and we only used Etherscan. We limited our study to Ethereum, because it has more daily transactions than rival ledgers like the Binance Smart Chain, Solana, and Terra. We selected distributions of voting rights tokens that have a market capitalization of more than \$1 billion. Four of our cases cover distinct DeFi categories [1]: lending (Maker), exchanges (Uniswap), assets (Yearn Finance), and derivatives (UMA). We added a fifth case, SushiSwap, because of its close relation to Uniswap plus the fact that SushiSwap issued voting rights tokens before Uniswap did. The initial version of SushiSwap, dated 26 August 2020, derived most of its smart contract code from Uniswap Version 2 [41]. At the time, there was one major distinction between the two exchanges: holders of voting rights tokens could govern SushiSwap, whereas developers governed Uniswap. This distinction proved short-lived. On 16 September 2020, Uniswap’s developers altered their governance plans and launched a voting rights token (UNI) [42]. Table 1 provides an overview of our selected cases.

Table 1. Selected projects for a multiple-case study

Project	Market cap*	Categorization
Uniswap	\$21,058,088,732	Exchange
Maker	\$3,272,720,932	Lending
SushiSwap	\$2,086,130,878	Exchange
Yearn Finance	\$1,888,059,030	Assets
UMA	\$1,489,888,679	Derivatives

* Data source: CoinMarketCap (14 May 2021).

Case studies are generally used to explore contemporary phenomena that are not easily distinguishable from their context [14]. For phenomena that involve a high degree of “explorability”, case studies allow researchers to gather insights about motivations, to discover causal links, and to explain why singular events occurred [43]. Within the nascent and under-researched field of DeFi [1, 2, 3], we assume that the case study method is valid.

Case studies build on a combination of six evidence sources: interviews, documentation, direct observations, participant observations, archival records, and physical

artifacts [14]. Since each DeFi case that we selected is open source, we could draw on several of these sources (Table 2). We relied on project documentation to identify differences among the cases' voting rights tokens and distribution strategies, and we used Etherscan to analyze the Ethereum blockchain and the distribution of voting rights tokens over time.

Table 2. Sources of case study data

Type	Description
(1) Documentation	(1) 145 pages of blog entries (2) 94 pages of whitepapers*
(2) "Physical" artifact	(1) ERC-20 tokens on the Ethereum blockchain

* Maker and Uniswap each have a formal whitepaper. Our three other cases do not.

We adopted the two-stage process proposed by Miles et al. [44] to scrutinize each case's documentation. Firstly, we considered the documentation individually and assigned initial codes. As a team of researchers, we convened regularly to review the emerging concepts and to ensure consistency in the coding system [45]. For the second stage, we clustered codes and assigned them to higher-level themes that emerged contingently via data collection (inductive coding) or else pertained to an existing hypothesis (deductive coding). Table 3 gives a brief example of our study's deductive and inductive coding.

Table 3. Exemplary coding scheme

Quote	1 st cycle coding	2 nd cycle coding
"UNI is a tradable asset and functions like most other standard ERC-20 tokens, except it has a deeper power as a voting mechanism."	UNI is an ERC-20 token with voting parameters.	<u>Deductive</u> Uniswap has voting rights tokens (UNI)
"Delegating UNI binds the voting power of your tokens to an address so it may be used to vote."	Delegating UNI to an address changes its status.	<u>Inductive</u> Voting with UNI occurs on-chain

For the second data source, using Etherscan, we retrieved the addresses that hold any of the tokens that correspond to voting units at two separate time points: t (24 May 2021) and $(t - 6)$ months. This allowed us to evaluate the case's decentrality at the present state and the change of decentrality over time. From the yielded data, we excluded automata's smart contract addresses, because automata typically do not exercise voting rights

or participate in governance processes. For the cases in which voting either takes place on-chain or voting rights are *staked* on-chain, we excluded the addresses with balances lower than the transaction fee (gas), since they cannot afford to cast a vote. We subtracted the balance of these excluded addresses from the overall token supply. We then bundled together the balances of addresses that have the same owner (for example, exchanges like Binance, Huobi, or KuCoin) and assumed that addresses with unknown owners belong to different individuals. As a consequence of this final assumption, our analysis presents the 'best case' scenario regarding the decentrality of the voting rights tokens.

Table 4. Quantitative data collection

Case	Holders (excl. contracts and low balances)
Uniswap	245,162
Maker	81,254
SushiSwap	77,442
Yearn Finance	35,471
UMA	13,891

We based our analysis on the number of accounts specified in Table 4 and the relevant account balances. Our conception of 'political decentralization' is based on the decentrality of voting rights tokens. To evaluate the decentrality of each case's voting rights tokens, we relied on seven metrics. For details about the seven metrics and their respective equations, please refer to Gochhayat et al. [11]. The seven metrics are:

1. **Normalized Fairness (NF)** [46]: A decentralized system would have higher fairness, because every participant would have equal voting power.
2. **Normalized Shannon's Entropy (NSE)** [47]: A centralized token distribution would have low entropy, since a few participants would heavily affect every decision.
3. **Gini coefficient (G)** [48]: A decentralized token distribution would have a low Gini coefficient, since most address holders would have equal voting power.
4. **Cosine Similarity (CS)**: A decentralized token distribution would have high similarity with a system in which every address has equal voting power.
5. **Jensen-Shannon Divergence (JSD)** [49]: A decentralized system would have a high degree of divergence from a system in which few participants have the majority of the voting power.
6. **Normalized Euclidean Distance (NED)**: A decentralized token distribution would have a small

degree of distance from a system in which every address owner has equal voting power.

7. **Soergel Distance (SD):** The same holds as for NED.

For the sake of a casual comparison or analogy, we cited findings from Gochhayat et al. [11] about the decentrality of the “governance layer” for both Bitcoin and Ethereum. We normalized our computed metrics (or replaced the non-normalized metrics with normalized versions) to assess the values in the [0,1] interval.

4. Findings

4.1. Properties of voting rights tokens

The voting rights tokens we studied are case-specific (Table 5). They have only three things in common: the tokens are all fungible and Ethereum-based (ERC-20), they grant holders the right to vote on community proposals, and they can be traded on crypto-asset exchanges. There are not many entitlements that are pre-determined by developers. In a figurative sense, voting rights tokens grant holders permission to enter the Senate Floor and cast votes on the measures of the day, which are contingent and thus not pre-determined.

Table 5. Voting rights tokens’ properties by case (as specified in project documentation)

Case	Cast votes	Choose voting proposals	Manage treasury	Manage collateral	Staking reward	Dispute resolution
UNI	X	X	X			
MKR	X	X		X		
SUSHI	X	X			X	
YFI	X	X				
UMA	X					X

The projects’ documentation, in total, specifies just four unique entitlements for holders of voting rights tokens: treasury management power (UNI), collateral selection power (MKR), a staking reward (xSUSHI for SUSHI holders), and dispute resolution power (UMA).

Uniswap’s documentation mentions a *community treasury*. Uniswap’s developers created the community treasury and allocated to it 43% of the total token supply (430 million UNI). The developers ceded control of the treasury to UNI token-holders on 18 October 2020. UNI

token-holders could then “vote to allocate UNI towards grants, strategic partnerships, governance initiatives, [...] and other programs” [42]. The documentation’s language is purposefully unspecific, because it is not possible to predict what decisions the UNI token-holders will make over time about the treasury funds or about particular proposals for changes to the project. Uniswap’s forum participants can pose questions about the project’s governance; but only UNI token-holders (and their *delegates*) can determine what proposals are formalized as executable code, then they can vote to implement or reject the selected proposals.

Maker’s documentation specifies the ability of MKR token-holders to vote on what Ethereum-based assets are used as collateral by the lending platform. The MKR token-holders also vote on the collateral’s corresponding risk parameters. (Volatile Ethereum-based assets are likely attributed strict risk parameters, whereas stable Ethereum-based assets are likely attributed lenient risk parameters.) Maker’s forum participants can formulate proposals for votes, not just MKR token-holders. MKR token-holders can elect *Active Proposals*, then they can exercise their right to vote on the Active Proposals. The first selection process for MKR token-holders is called *proposal polling*; the second process is called *Executive Voting*.

SushiSwap’s SUSHI token can be staked on the exchange. Holders then receive a tokenized derivative (xSUSHI) as a reward. The SushiSwap exchange charges participants a 0.3% trading fee. A small portion (0.05%) of this trading fee is used to generate the xSUSHI tokens, which are then allocated to the holders of staked SUSHI tokens. At any time on the exchange, holders can convert the xSUSHI derivative to SUSHI. (The xSUSHI/SUSHI exchange rate is variable.) Only SUSHI token-holders can select what forum proposals are put up for a vote, then they can exercise their right to vote on the selected proposals.

UMA token-holders can likewise exercise their right to vote on proposals that emerge from the UMA community forum. In addition to this, UMA token-holders have an exclusive right to cast dispute resolution votes via the *UMA Voter dApp*. The work required to resolve disputes generates new UMA tokens (as payment for the work). The UMA token supply is thus inflationary. There are two grounds for disputes: incorrect crowd-sourced (off-chain) pricing information, and contracts liquidated for an improper reason. UMA token-holders can assess the disputes and verify the pricing data and liquidation data. UMA employs this system, named the *Data Verification Mechanism*, instead of an oracle’s data feed.

Yearn Finance’s project documentation does not specify any unique entitlements for YFI token-holders. Forum participants can formulate proposals – any proposal whatsoever – without needing to hold YFI. YFI token-holders can then choose what proposals are put up for a vote, and they can exercise their right to vote on the selected proposals.

4.2. Token distribution strategies

Distribution strategies are especially important for voting rights tokens, since the token distribution determines how many individuals can exercise control over a project and how much voting power individuals possess. Our five DeFi cases all aimed for distributed governance, yet each case deployed a different token distribution strategy.

MKR tokens were pre-mined (and thus not produced by the work or stake of the lending platform’s participants) and sold via an exchange. Some UNI tokens were pre-mined; others were produced by the stake of the exchange’s participants. UMA’s initial token supply was pre-mined, but as noted, UMA’s inflationary supply is generated by dispute resolution work. Neither SUSHI nor YFI tokens were pre-mined. It follows that SushiSwap and Yearn Finance are our only two cases to deploy a fair launch strategy (which precludes pre-mined tokens).

Yearn Finance’s YFI tokens must be earned via contributions to the project (such as liquidity provision)

or else purchased via crypto-asset exchanges. Since the YFI tokens were released in accordance with a strict fair launch policy, no tokens were distributed to developers. Yearn Finance’s developers consequently lacked capital and could not afford to pay a third party to audit the YFI token’s smart contract code. Yearn Finance’s fair launch strategy thus increased the risk placed upon the community of YFI holders and offered a lack of traditional auditing guarantees. The outcome is ambiguous: the distribution of voting power incurs a distribution of risk.

SushiSwap’s fair launch is also ambiguous, but its details differ notably from those of Yearn Finance’s fair launch. No SUSHI tokens were pre-mined and allocated to privileged insiders like venture capitalists, but 10% of all SUSHI tokens generated by the exchange’s liquidity providers automatically transferred to a treasury accessible only by developers. On 5 September 2020, a pseudonymous core developer, Chef Nomi, sold over \$13 million worth of SUSHI tokens from the SushiSwap treasury. In response to community protests and public criticism, Chef Nomi apologized on Twitter and returned the funds to the SushiSwap treasury six days later; but nothing guaranteed this outcome.

4.3. Decentrality of the token distributions

The decentrality of our five cases’ voting rights tokens is very low. We therefore deduce that, for each of our cases, the voting power is highly concentrated.

Table 6. Measured governance decentrality of cases

		1 - NF	1 - NSE	G	1 - CS	JSD	NED	SD
Bitcoin [11]	t = 2.4.2020	0.850	0.315	0.821	-	-	-	-
Ethereum [11]	t = 7.4.2020	0.962	0.438	0.904	-	-	-	-
Uniswap	t – 6 months	0.99974	0.61173	0.98810	0.98362	0.88687	0.99178	0.99980
	t = 24.5.2021	0.99980	0.59213	0.99132	0.98568	0.89963	0.99281	0.99955
Maker	t – 6 months	0.99941	0.57847	0.99491	0.97534	0.91246	0.98759	0.98757
	t = 24.5.2021	0.99938	0.55255	0.99475	0.97489	0.91257	0.98737	0.98557
SushiSwap	t – 6 months	0.99870	0.46674	0.98724	0.96364	0.85819	0.98165	0.99999
	t = 24.5.2021	0.99908	0.46918	0.99098	0.96952	0.88238	0.98464	0.99999
Yearn Finance	t – 6 months	0.99786	0.38747	0.91224	0.95725	0.64872	0.96148	0.91483
	t = 24.5.2021	0.99892	0.46993	0.98290	0.96671	0.83568	0.98322	0.96033
UMA	t – 6 months	0.99886	0.67603	0.99635	0.96421	0.93759	0.98194	0.99985
	t = 24.5.2021	0.99921	0.65222	0.99651	0.97056	0.93797	0.98517	0.99974

Table 6 presents our quantitative measurements together with the analogous measurements by Gochhayat et al. [11]. For the sake of a clearer overview, we changed the direction of the metrics NF, NSE, and CS in the table by subtracting their normalized values from 1. As a result, high values indicate centralization.

The two cases that adopted a fair launch strategy, Yearn Finance and SushiSwap, have token distributions that are relatively less centralized than the other cases' token distributions. This aside, the distribution of YFI tokens tended towards centralization after the tokens' inception in July 2020. Likewise, the distribution of SUSHI tokens tended towards centralization after their inception in August 2020. The distribution of YFI tokens started with a greater degree of decentrality, and it remains less centralized than the distribution of SUSHI tokens. This might be caused by the 10% allocation of SUSHI tokens to the SushiSwap treasury.

UMA tokens are the most centralized. Two other cases that did not deploy a fair launch, Maker and Uniswap, are almost identical. The former's token distribution is slightly less centralized. The centralization of UNI tokens is largely attributable to the launch strategy that allocated 40% of UNI tokens to Uniswap's developers, early investors, and advisors. As of 24 May 2021, if automata's smart contract addresses are excluded, then just 23 addresses control more than 50% of the active token supply. (The active supply excludes the tokens that are held by smart contract addresses.)

5. Discussion and outlook

De Tocqueville's maxim about re-centralization can be rehashed for DeFi. Our multiple-case study yields a significant finding: decentralization is not actualized in the sense of evenly distributed voting power, even though decentralization is achieved in the sense of independence from centralized legal authorities and regulated companies. In political terms, hopes for evenly distributed, democratic governance are presently unfounded in our five selected cases, including the two cases that deployed a fair launch strategy. The measured distribution outcomes reveal concentrated voting power for all five cases (in comparison with a hypothetical, fully decentralized case in which each address would have equal voting rights).

In response to our findings, one cannot overlook the power of large token-holders, commonly nicknamed *whales*. Whales thwart the ambitions of Yearn Finance, along with other DAOs, to evenly distribute tokens and network resources [35]. Granular analysis of our data

revealed that for Yearn Finance, 50% of the active token supply is controlled by just 21 accounts. These 21 accounts can potentially collude and exercise massive voting power. If they act together as a 'whale pod,' they can also manipulate token value via 'pump-and-dump' schemes [50].

Regulation limits the power of whales in traditional financial markets. Whales can swim freely by comparison in DeFi's murky waters. The unique value proposition of DeFi projects – that is, the type of decentralization that DeFi projects actually achieve – is the ability to bypass regulation and 'live' on the Ethereum network [2, 3, 36]. Ethereum's "unincorporated distributed ledger technology" has the potential to "undermine traditional forms of accountability" as well as erode the "effectiveness of traditional financial regulation and enforcement" [3]. In addition, DeFi's anonymous transactions and lack of KYC/AML checks are predictable concerns for regulators [51, 3, 52, 29, 20].

DeFi thus rehashes and intensifies the Bitcoin and Ethereum networks' history of regulatory challenges. In response, researchers and developers could turn to what Zetzsche et al. [3] describe as *embedded regulation* – regulators' objectives and limits hard-coded into a financial services platform. Embedded regulation could hypothetically limit the power of whales, for the sake of political decentralization. There is no guarantee, however, that embedded regulation will find success among DeFi communities and crypto-asset investors. A successful or prominent case of embedded regulation does not yet exist.

6. Conclusion

Non-custodial DeFi projects typically do not involve sovereign currencies, investor registration processes, or a corporate office located within a particular jurisdiction; and they are not managed or majority-owned by legal entities [3]. This motivated our exploratory, multiple-case study of five Ethereum-based DeFi projects and their tradable voting rights tokens: Uniswap, Maker, SushiSwap, Yearn Finance, and UMA. Our core finding is that each case's token distribution strategy failed to achieve measurable decentrality. In theory, voting rights tokens enable decentralized, community-based governance; but according to our measurements, the control of major DeFi projects is highly centralized.

We acknowledge the limitations of metrics used to assess decentrality as well as the fact that different metrics can yield different outcomes. We are confident,

however, that our measured outcomes reflect reality. We employed a variety of metrics from different categories. Each metric indicated that, for all our cases, the distribution of voting rights tokens began centralized, and as time progressed, the distribution became even more centralized (Table 6). Each of the metrics we computed for our cases exhibit centralization. This means that the fair launch strategies deployed by Yearn Finance and SushiSwap achieved little success.

The nascent field of DeFi offers multiple avenues for future research. The concentrated voting power, general motives, and trading activities of DeFi whales are under-researched topics. DeFi forum posts about whales, DeFi communities' voting proposals to limit the power of whales, and DeFi commentaries by social media influencers represent valuable data sources for future social-scientific research. Legal scholars can investigate why non-custodial DeFi projects are especially challenging for regulators, and they can study developers' embedded regulation experiments.

DeFi's genuine novelty consists in the relationship between its distributed governance and regulators. What regulators face is both ironic and serious – a politically centralized yet globally diffuse form of control over multi-billion-dollar financial services. This form of control cannot be seized, limited, or broken up under current regimes.

7. Acknowledgements

Supported by PayPal and the Luxembourg National Research Fund FNR, Luxembourg (P17/IS/13342933/PayPal-FNR/Chair in DFS/Gilbert Fridgen).

8. References

- [1] Amler, H., L. Ekey, S. Faust, M. Kaiser, P. Sandner, and B. Schlosser, "DeFi-ning DeFi: Challenges & Pathway", *arXiv: Cryptography & Security*, 2021.
- [2] Katona, T., "Decentralized Finance: The Possibilities of a Blockchain 'Money Lego' System", *Financial and Economic Review* 20(1), 2021, pp. 74–102.
- [3] Zetzsche, D.A., D.W. Arner, and R.P. Buckley, "Decentralized Finance", *Journal of Financial Regulation* 6(2), 2020, pp. 172–203.
- [4] Kostakis, V., and C. Giotitsas, "The (A)Political Economy of Bitcoin", *tripleC* 12(2), 2014.
- [5] De Tocqueville, A., *L'Ancien Régime et la Révolution*, Michel-Lévy Frères, Paris, 1856.
- [6] De Filippi, P., and B. Loveluck, "The invisible politics of Bitcoin: governance crisis of a decentralised infrastructure", *Internet Policy Review* 5(3), 2016.
- [7] Hsieh, Y.-Y., J.-P. Vergne, and S. Wang, "Chapter 3: The internal and external governance of blockchain-based organizations", In M. Campbell-Verduyn, ed., *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global governance*. Routledge, Taylor & Francis Group, London; New York, 2017.
- [8] Reijers, W., and M. Coeckelbergh, "The Blockchain as a Narrative Technology: Investigating the Social Ontology and Normative Configurations of Cryptocurrencies", *Philosophy & Technology* 31(1), 2018, pp. 103–130.
- [9] Parkin, J., "The senatorial governance of Bitcoin: making (de)centralized money", *Economy and Society* 48(4), 2019, pp. 463–487.
- [10] Schneider, N., "Decentralization: an incomplete ambition", *Journal of Cultural Economy* 12(4), 2019, pp. 265–285.
- [11] Gochhayat, S.P., S. Shetty, R. Mukkamala, P. Foytik, G.A. Kamhoua, and L. Njilla, "Measuring Decentrality in Blockchain Based Systems", *IEEE Access* 8, 2020, pp. 178372–178390.
- [12] Gervais, A., G.O. Karame, V. Capkun, and S. Capkun, "Is Bitcoin a Decentralized Currency?", *IEEE Security & Privacy* 12(3), 2014, pp. 54–60.
- [13] Oliveira, L., L. Zavolokina, I. Bauer, and G. Schwabe, "To Token or not to Token: Tools for Understanding Blockchain Tokens", *International Conference of Information Systems*, s.n. (2018).
- [14] Yin, R.K., *Case study research and applications: design and methods*, SAGE, Los Angeles London New Delhi Singapore Washington DC Melbourne, 2018.
- [15] Kaplan, B., and D. Duchon, "Combining Qualitative and Quantitative Methods in Information Systems Research: A Case Study", *MIS Quarterly* 12(4), 1988, pp. 571.
- [16] Nagarajan, S., "Bitcoin hit the \$1 trillion market cap milestone twice as fast as Amazon and three times faster than Apple, data shows", *Business Insider*, 2021.
- [17] Nakamoto, S., "Bitcoin: A Peer-to-Peer Electronic Cash System", 2008.
- [18] Böhme, R., N. Christin, B. Edelman, and T. Moore, "Bitcoin: Economics, Technology, and Governance", *Journal of Economic Perspectives* 29(2), 2015, pp. 213–238.
- [19] De Filippi, P., "Bitcoin: a regulatory nightmare to a libertarian dream", *Internet Policy Review* 3(2), 2014, pp. 11.
- [20] Foley, S., J.R. Karlsen, and T.J. Putniņš, "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?", *The Review of Financial Studies* 32(5), 2019, pp. 1798–1853.
- [21] Vigna, P., and M. Casey, *The age of cryptocurrency: how bitcoin and the blockchain are challenging the global economic order*, Picador, St. Martin's Press, New York, 2016.
- [22] Pazaitis, A., P. De Filippi, and V. Kostakis, "Blockchain and value systems in the sharing economy: The illustrative case of Backfeed", *Technological Forecasting and Social Change* 125, 2017, pp. 105–115.

- [23] Gupta, M., and P. Gupta, "Gini Coefficient Based Wealth Distribution in the Bitcoin Network: A Case Study", In R. Sharma, A. Mantri and S. Dua, eds., *Computing, Analytics and Networks*. Springer Singapore, Singapore, 2018, 192–202.
- [24] World Bank, "Gini Index (World Bank estimate)", *The World Bank Data*, 2019. <https://data.worldbank.org/indicator/SI.POV.GINI>
- [25] Lin, Q., C. Li, X. Zhao, and X. Chen, "Measuring Decentralization in Bitcoin and Ethereum using Multiple Metrics and Granularities", *2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW)*, (2021), 80–87.
- [26] de Vries, A., "Bitcoin's Growing Energy Problem", *Joule* 2(5), 2018, pp. 801–805.
- [27] Drummer, D., and D. Neumann, "Is code law? Current legal and technical adoption issues and remedies for blockchain-enabled smart contracts", *Journal of Information Technology* 35(4), 2020, pp. 337–360.
- [28] Lehmann, M., "Who Owns Bitcoin? Private Law Facing the Blockchain", *European Banking Institute*, 2019.
- [29] Möser, M., R. Böhme, and D. Breuker, "Towards Risk Scoring of Bitcoin Transactions", In R. Böhme, M. Brenner, T. Moore and M. Smith, eds., *Financial Cryptography and Data Security*. Springer Berlin Heidelberg, 2014, 16–32.
- [30] Buterin, V., "Ethereum: A next-generation smart contract and decentralized application platform", 2014.
- [31] Buterin, V., "DAOs, DACs, DAs and More: An Incomplete Terminology Guide", *Ethereum Blog*, 2014.
- [32] Wright, A., and P. De Filippi, "Decentralized Blockchain Technology and the Rise of Lex Cryptographia", *SSRN Electronic Journal*, 2015.
- [33] Sunyaev, A., N. Kannengießer, R. Beck, et al., "Token Economy", *Business & Information Systems Engineering*, 2021.
- [34] Andersen, J.V., and C. Ingram Bogusz, "Self-Organizing in Blockchain Infrastructures: Generativity Through Shifting Objectives and Forking", *Journal of the Association for Information Systems*, 2019, pp. 1242–1273.
- [35] DuPont, Q., "Chapter 8: Experiments in algorithmic governance – A history and ethnography of 'The DAO,' a failed decentralized autonomous organization", In *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global governance*. Routledge, Taylor & Francis Group, London; New York, 2017.
- [36] Chen, Y., and C. Bellavitis, "Blockchain disruption and decentralized finance: The rise of decentralized business models", *Journal of Business Venturing Insights* 13, 2020, pp. e00151.
- [37] Young, M., "A Guide and Short History of SushiSwap", *BeInCrypto*, 2020.
- [38] Chen, Y., "Blockchain tokens and the potential democratization of entrepreneurship and innovation", *Business Horizons* 61(4), 2018, pp. 567–575.
- [39] Catalini, C., and J.S. Gans, "Initial Coin Offerings and the Value of Crypto Tokens", *MIT Sloan Research Paper* 5347(18), 2018, pp. 37.
- [40] Fridgen, G., F. Regner, A. Schweizer, and N. Urbach, "Don't Slip on the Initial Coin Offering - A Taxonomy for a Blockchain enabled form of Crowdfunding", *Proceedings of the 26th European Conference on Information Systems*, (2018).
- [41] SushiSwap, "The SushiSwap Project: An evolution of Uniswap with SUSHI tokenomics", *SushiSwap Blog*, 2020. <https://web.archive.org/web/20200827023308/https://medium.com/sushiswap/the-sushiswap-project-c4049ea9941e>
- [42] Uniswap, "Introducing UNI", 2020. <https://uniswap.org/blog/uni/>
- [43] Gable, G.G., "Integrating case study and survey research methods: an example in information systems", *European Journal of Information Systems* 3(2), 1994, pp. 15.
- [44] Miles, M.B., A.M. Huberman, and J. Saldaña, *Qualitative Data Analysis: an expanded sourcebook*, Sage Publications, Inc, 1994.
- [45] Klein, H.K., and M.D. Myers, "A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems", *MIS Quarterly* 23(1), 1999, pp. 28.
- [46] Jain, R.K., D.-M.W. Chiu, and W.R. Hawe, "A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems", *ACM Transaction on Computer Systems*, 1984, pp. 39.
- [47] Shannon, C.E., "A Mathematical Theory of Communication", *Bell System Technical Journal* 27(3), 1948, pp. 379–423.
- [48] Gini, C., "On the Measure of Concentration with Special Reference to Income and Statistics", *Colorado College Publication* 208, 1936, pp. 73–79.
- [49] Lin, J., "Divergence measures based on the Shannon entropy", *IEEE Transactions on Information Theory* 37(1), 1991, pp. 145–151.
- [50] Kamps, J., and B. Kleinberg, "To the moon: defining and detecting cryptocurrency pump-and-dumps", *Crime Science* 7(1), 2018, pp. 18.
- [51] Anderson, R., I. Shumailov, and M. Ahmed, "Making Bitcoin Legal", In V. Matyáš, P. Švenda, F. Stajano, B. Christianson and J. Anderson, eds., *Security Protocols XXVI*. Springer International Publishing, Cham, 2018, 243–253.
- [52] Renwick, R., and R. Gleasure, "Those who control the code control the rules: How different perspectives of privacy are being written into the code of blockchain systems", *Journal of Information Technology*, 2020, pp. 026839622094440.

3.4. Research Paper 4: *Pricing dynamics and herding behaviour of NFTs*

Pricing dynamics and herding behaviour of NFTs

Gilbert Fridgen¹  | Roman Kräussl²  |
Orestis Papageorgiou¹  | Alessandro Tugnetti³ 

¹SnT - Interdisciplinary Center for Security, Reliability and Trust, University of Luxembourg, Luxembourg

²Bayes Business School (formerly Cass), Hoover Institution at Stanford University and CEPR, London

³Department of Finance, University of Luxembourg, Luxembourg

Correspondence

Orestis Papageorgiou, SnT -
Interdisciplinary Center for Security,
Reliability and Trust, University of
Luxembourg, 29JFK, Blvd JF Kennedy,
L-1855 Luxembourg.
Email: orestis.papageorgiou@uni.lu

Funding information

Luxembourg National Research Fund
(FNR), Grant/Award Numbers: 16326754,
NCER22/IS/16570468/NCER-FT; PayPal,
PEARL, Grant/Award Number: 13342933/
GilbertFridgen; Luxembourg Institute for
Advanced Studies (IAS), Young Academics
scheme, DATART project

Abstract

This paper analyzes the sales of 875,389 art non-fungible tokens (NFTs) on the Ethereum blockchain to identify the key determinants influencing NFT pricing and market dynamics. We find that market liquidity and trade volume are strong predictors of NFT prices. Contrarily, social media activity negatively correlates with prices. Introducing an artist ranking system, our study reveals a “superstar effect”, with a few artists dominating sales, and herding behaviour within the NFT market.

KEYWORDS

art market, cryptocurrency, Ethereum blockchain, herding, liquidity, network, nonfungible tokens, price index, social media, speculation

This research was funded in part by the Luxembourg National Research Fund (FNR), grant reference 16326754 and NCER22/IS/16570468/NCER-FT, by PayPal, PEARL grant reference 13342933/Gilbert Fridgen and by the Luxembourg Institute for Advanced Studies (IAS), Young Academics scheme, DATART project. For the purpose of open access and in fulfillment of the obligations arising from the grant agreement, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *European Financial Management* published by John Wiley & Sons Ltd.

JEL CLASSIFICATION

C55, G11, Z11

1 | INTRODUCTION

Nonfungible tokens (NFTs) rose in prominence in 2021, amassing more than \$44 billion in traded volume and attracting the attention of both institutional and retail investors (Chainalysis, 2022). In less than a year, NFTs transitioned from being recognized by only a small community of blockchain enthusiasts to becoming the subject of full-length articles in major news outlets, being described as a revolution not only in the art industry but also in various other sectors (Wilson et al., 2021). This prominence even led to Collins Dictionary choosing “NFT” as the 2021 word of the year (Collins, 2021). Such popularity inevitably prompted a considerable amount of scrutiny regarding the utility of NFTs and whether they can truly reform the digital asset sector. One might argue that NFTs are not merely a fleeting trend but rather part of a broader technological evolution in the blockchain space. Some perceive the crash in 2022 as a correction through the underlying value of NFTs, aligning with a more nuanced understanding of market dynamics. Others may contend that the drop in value does not diminish the inherent properties that make NFTs unique; instead, it could potentially pave the way for a more sustainable growth pattern that aligns with long-term financial models.

Despite the hype subsiding and the trading volume falling, NFTs have recently found applications ranging from empowering business models innovation (Hartwich et al., 2023; Li & Chen, 2023) to serving as the primary component in several initial coin offerings used by firms to raise capital (Holden & Malani, 2022). Besides these applications, digital collectibles and artworks represent the most prevalent category within the NFT marketplace. Probably the most influential art NFT sale came in March 2021 with Christie's auction house first-ever sale of a digital artwork: Beeple's “Everydays: The First 5000 Days” was sold for \$69.3 million, positioning the artist among the most valuable living creators. This, combined with other notable sales such as Pak's “The Merge”, the most expensive NFT sold to date, and certain CryptoPunks selling for more than \$10 million, propelled NFTs to the forefront of the art world in 2021, dominating the digital art space (Financial Times, 2021a) and generating almost as much volume as the traditional art space (Financial Times, 2021b). Even in 2022, when every other NFT category experienced a dramatic decrease in the number and value of sales, art NFTs have seen a slight increase in the average sale value (McAndrew, 2022). Despite facing market volatility, NFTs with actual use-cases, such as those in gaming, cultural preservation and fractional ownership, are emerging as particularly promising investment avenues, signalling a market evolution that increasingly prioritizes utility-driven assets over speculative collectibles and thereby suggesting a maturing landscape where authentic value will likely catalyze future growth, thus providing opportunities for both creators and investors to strategically position themselves in sectors of tangible utility and enduring value (Hategan, 2023).

Regardless of growth, relatively little is understood regarding the attributes that lend value to art NFTs and influence their selling probability. Although scarcity is widely acknowledged, both by academics (Alizadeh et al., 2023; Bamakan et al., 2022; Bao & Roubaud, 2022; Mekacher et al., 2022) and by practitioners (Sotheby's, 2021; Valeonti et al., 2021), as an essential trait, other factors, such as the impact of the artist, and the collection characteristics, remain vague, while academic research articles tend to focus on a few collections and attributes (Cong et al., 2022; Huang & Goetzmann, 2023; Sockin & Xiong, 2022).

Dowling (2022a) provides a pioneering exploration of Decentraland LAND NFTs, drawing a mental association between virtual and physical land, and identifying possible inefficiencies in pricing behaviour akin to early cryptocurrency pricing. This notion of inefficiency is further examined from a broader perspective by Cheah and Fry (2015) and Urquhart (2016), underlining the nascent nature of NFT markets. Schaar and Kampakis (2022) take a quantitative approach to investigate the CryptoPunks collection, highlighting an impressive average monthly return of 34.19% over 3 years and the significant role that rarity plays in determining prices. This emphasis on rarity is also central to the work of Mekacher et al. (2022), who delve into 410 collections to demonstrate that rarer NFTs command higher prices and are less risky. Horky et al. (2022) employ a combination of econometric tools and machine learning in their study of the digital art market through SuperRare, finding that hedonic pricing models furnish valuable insights into NFT prices, independent of cryptocurrencies.

Broadening the scope to encompass the entire market, Dowling (2022b) work stands out by interrogating the connection between NFT pricing and cryptocurrencies, uncovering limited volatility transmission effects but a notable comovement between the two. Ante (2022) adds nuance to this relationship by revealing how Bitcoin price shocks trigger an increase in NFT sales, while Ether price shocks have a converse effect. Borri et al. (2023) take an ambitious step by constructing a comprehensive data set for the overall NFT market, unearthing the nuanced interplay between creator fees, rarity, visual characteristics and prices. Kong and Lin (2021) enrich this perspective by emphasizing the role of well-connected investors in the pricing dynamics of NFTs. Mazur (2021) shifts the focus to the risk and return characteristics of NFT-based start-ups, illustrating a staggering investment multiple of 40 over the long term. In their seminal paper, Nadini et al. (2021) map the structure and evolution of the entire NFT market, identifying sale history and visual features as key price predictors. Hemenway Falk et al. (2022) introduce a novel dimension by exploring the value of creator royalties, uncovering their potential to reshape the NFT market through risk-sharing, dynamic pricing and price discrimination.

While the existing literature on NFT pricing and determinants has provided a comprehensive view of the market, the exploration of herding behaviour in the context of art NFTs presents a unique and uncharted territory. Herding behaviour, as observed in traditional financial markets (Banerjee, 1992; Bikhchandani et al., 2021; Shiller, 1995), refers to the inclination of investors to follow the majority's actions rather than making independent decisions based on intrinsic value. Art NFTs offer a perfect primer to study this phenomenon, primarily due to their parallel nature with the physical art market where the phenomenon of herding behaviour has been well-documented. Azarmi and Menny (2013) study of the fine arts market explores the tendency of investors to gravitate towards well-established artists, leaving the quality of the work secondary to the historical auction performance of the artist. They found that a small fraction of artists dominates financial activity and that contemporary and less-established artists are subject to more herding behaviour. This observation reflects a pattern where investors follow trends and the actions of others, often leading to inflation in the value of specific artists or styles. Art NFTs amplify these dynamics by adding layers of complexity. Unlike physical art, where authenticity, provenance, and physical condition play a role in valuation, art NFTs operate on a digital platform where these factors are replaced by digital scarcity, network effects and integration with cryptocurrencies. The rapid rise of art NFTs, by significant price volatility and intense media coverage, has many resemblances to the herding behaviour observed in the traditional art market. Investors and collectors rush towards certain NFTs, often driven more by hype and the actions of others than a deep understanding of the digital art's inherent value (Nadini et al., 2021).

Moreover, the volatility and correlation with the broader cryptocurrency market further amplify the complexity of art NFTs. The fluctuations in their prices, often driven by external factors such as changes in Ethereum or Bitcoin values, create a turbulent environment where herding behaviour could thrive. The allure of quick profits, media attention and the novelty of owning a unique digital asset can lead to a cascade effect where one investment decision triggers another, often without a rational analysis of the asset's underlying value. The work of Bao et al. (2023) provides the first empirical evidence on herding in the NFT market. They identify three waves of herding in the NFT market, with daily market returns becoming more volatile during these periods. Their findings also reveal that herding is more likely to emerge as the proportion of newcomers increases and that media exposure drives investors' attention when herding arises. They further find a connection with the return on Ethereum but a diminishing effect with the return on Bitcoin. Yousaf and Yarovaya (2022) examine herding behaviour across three cryptocurrency classes, including NFTs. Their time-varying analysis identified herding in conventional cryptocurrencies during the most recent bubble of 2021 but failed to demonstrate evidence of herding in NFTs during various market conditions. This contrast presents an intricate picture of the NFT market, suggesting that herding behaviour may manifest differently across various aspects of the broader cryptocurrency space. The speculative nature of NFTs and DeFi markets, often perceived as bubble behaviours, has been explored by Wang et al. (2022). They document that both NFT and DeFi markets exhibit speculative bubbles, with NFT bubbles being more recurrent and having higher explosive magnitudes. The bubbles in these markets are correlated with market hype and broader cryptocurrency market uncertainty, though they also recognize periods where bubbles are not detected, indicating intrinsic value in these markets.

Our study's approach merges financial, descriptive and social network features of art NFTs. We aim to explore how market frictions operate in a digital asset environment, and how they may affect pricing, liquidity and investor behaviour. Specifically, this paper investigates the relationship between art NFT pricing and market herding behaviour. The primary objective is to construct a comprehensive price index for art NFTs, integrating multifaceted attributes, such as financial, artistic and social network influences. This setting aims to challenge conventional asset pricing models, acknowledging the unique, nonfundamental value characteristics of art NFTs (Taleb, 2021). The methodology employs hedonic and repeat-sales regression (RSR) models, analyzing differentiated market trends and isolating intrinsic item qualities to understand price dynamics over time.

Second, we examine investor herding in the art NFT market. We assess whether collective investor behaviour, influenced more by group dynamics than individual rationale, plays a significant role in art NFT valuation cycles. Our research uses Cross-Sectional Absolute Deviation (CSAD) to analyze price comovements, particularly in the context of low liquidity and asymmetric information typical of the art NFT domain. This analysis aims to uncover deviations from fundamental valuation driven by investor herding, providing insights into price formation and market behaviour in the evolving NFT landscape. We focus on the entire art NFT market, analyzing every art collection and transaction performed on Ethereum, the blockchain with the most NFT transactions and home to many successful token offerings from firms. Our sample comprises 875,389 art NFTs that have been deployed on the Ethereum blockchain and have been involved in 385,884 sales.

Our findings confirm the positive correlation between the number of trades and the floor price of NFTs with their average price, elucidating critical elements of the NFT market's pricing structure. We find an unexpected inverse relationship between social media activity and NFT

prices, presenting a challenge to popular conceptions of the NFT market. Remarkably, our exploration of herding behaviour revealed a small group of artists who command most of the market's activity, mirroring trends seen in traditional art markets, while also spotlighting the market's susceptibility to potential speculative bubbles. Our results also demonstrate that in instances of high market volatility, investors diverge from herding, adopting more risk-averse strategies. New entrants initially exert a stabilizing effect on market prices, but eventually contribute to increased financial volatility. Notably, a bullish broader cryptocurrency market acts as a moderating variable, diverting capital away from NFTs, thus mitigating herd-driven price inflation in this asset class.

These insights into the art NFT market's pricing dynamics and liquidity constraints contribute to the existing financial literature, providing a rich context for understanding a new asset class within a blockchain-enabled environment. The link between herding behaviour, speculative bubbles and the inherent value of NFTs opens up avenues for understanding the art market's financial dynamics and the broader economic landscape. The integration of insights from traditional herding behaviour theories (e.g., Banerjee, 1992; Shiller, 1995) with contemporary studies on NFTs and art markets provides a nuanced understanding of market dynamics. This interdisciplinary approach is instrumental in drawing parallels between classical assets, like, securities, physical art and other collectibles, and emerging digital assets, like, NFTs.

What sets NFTs apart is the unprecedented access to a wealth of real-time data about transactions and financial characteristics of these digital assets. Unlike traditional markets where information might be fragmented, delayed or obscured by various market frictions, the blockchain technology underpinning NFTs ensures that every transaction is transparent, timestamped and publicly accessible. The ability to track and analyze these transactions in real-time opens up new avenues for research, allowing for a more nuanced understanding of market dynamics, pricing mechanisms and investor behaviour. As an example, researchers can study market frictions in a way that was previously unattainable. Market frictions, such as transaction costs, information asymmetry and liquidity constraints, play a critical role in asset pricing and investment strategies. They have been studied extensively in traditional financial markets by authors, like, Amihud and Mendelson (1986), Stiglitz (1989) and Vayanos and Wang (2012). In the context of NFTs, these frictions might manifest differently, and their impact on market dynamics could be distinct from what is observed in traditional markets.

The remainder of this paper is structured as follows. Section 2 discusses the data extraction techniques. Section 3 describes the econometric models used to construct the price indexes and to study the herd bias. Section 4 discusses our findings and Section 5 concludes.

2 | DATA

Blockchain-related analyses are often classified into two broad categories: on-chain and off-chain. On-chain research involves data retrieved directly from the blockchain's public ledger, whereas off-chain analysis utilizes data sources outside the blockchain, such as price-tracking websites. For our analysis, we rely on both categories to obtain as much data as possible on NFTs and their pricing. Specifically, by leveraging the granular data available through NFT transactions, we strive to provide a more comprehensive understanding of the art NFT market. Please, refer to Appendix A for a detailed explanation of data extraction, cleaning and preparation processes.

Our sample comprises 875,389 art NFTs that have been deployed on the Ethereum blockchain and have been involved in 385,884 sales between 15 July 2018 and 10th 10 February 2022. To facilitate a deeper understanding of this extensive data set, we have organized the information into two distinct tables.

Table 1 presents the descriptive statistics, summarizing key features such as median, standard deviation and other relevant statistical measures that provide an overall picture of the art NFT market. These statistics offer insights into the general trends, distributions and characteristics that define the market landscape. Furthermore, they are instrumental in constructing a comprehensive price index for the art NFT market. The table contains financial variables such as the returns of the Bloomberg Galaxy Crypto Index (*BGCI_return*) since past research has indicated the valuation of NFTs can move in tandem with broader cryptocurrency markets, suggesting that investor sentiment in cryptocurrencies could spill over into NFT valuations. To understand the market's valuation floor, the table includes the lowest selling price of an NFT in a collection (*floor_price*), which sets the baseline for the collection's value. Complementing this, the most recent selling price (*last_price*) offers a snapshot of the current market demand. The number of times an NFT has been sold (*num_trades*) is tracked to assess liquidity and investor interest, while the time span between the first and last sales (*timediff*) provides insight into the NFT's market presence and activity over time. Additionally, it contains visual analytics, quantifying the aesthetic attributes of NFTs, such as the dominance of specific colours within the image (*white, black and so forth*), the total number of colours in each image (*num_colours*) and the complexity of the image expressed by the normalized Shannon's entropy (*norm_shannon_entropy*). These visual analytics reflect the NFT's aesthetic uniqueness, which can influence its market value (Nadini et al., 2021). Lastly, the table captures the presence of NFT collections on Twitter, with metrics like the number of accounts each collection follows (*following_count*) and the monthly average of quotes received (*quote_count_month*). These metrics serve as indicators of the collection's community engagement and outreach, which are essential for understanding the social standing and potential influence of NFT collections. Primarily, the rationale behind leveraging Twitter data, as opposed to other social media, such as Reddit, Google and Facebook, hinges on its relative accessibility and the ease of an expedited data aggregation and analysis process, consistent with the current literature on digital asset pricing (Kapoor et al., 2022). Furthermore, Twitter stands out as a primary source of insights about the NFT market. As noted by Nadini et al. (2021), a predominant proportion of NFT transactions, happens on platforms like *OpenSea*, where artists predominantly link their Twitter accounts. This unique aspect of the NFT market makes Twitter a more relevant and comprehensive source for gauging market sentiment and trends, compared with other social media platforms. This extensive inclusion is aimed at enhancing the index's scope and accuracy, drawing on methodologies from the hedonic and repeat-sales literature (Bailey et al., 1963; Borri et al., 2023).

Table 1 shows that the distribution of most variables is negatively skewed; for many of them, this remains around 0 up to the second quartile and then increases exponentially towards the third and fourth quartiles. The application of the Jarque–Bera (JB) and Kwiatkowski–Phillips–Schmidt–Shin tests to our data uncovers intricate statistical properties. Notable deviations from normality and the existence of unit roots in continuous variables such as *BGCI_return* emphasize the nonlinear characteristics of the art NFT market. The JB statistics reveal pronounced skewness and kurtosis, confirming the fragmented nature of the NFT market (Caporale et al., 2021).

The variables *floor_price* and *last_price* highlight the extensive range of pricing that characterizes the art NFT market. This dynamic pricing behaviour correlates with observations in

TABLE 1 Descriptive statistics art nonfungible tokens (NFTs).

This table reports a detailed statistical overview of art NFTs. Panel A focuses on the continuous variables, providing a comprehensive statistical summary of the data extracted. This includes the median, standard deviation, minimum, maximum, Jarque–Bera (JB) test statistic for normality and the Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test for stationarity across both the aggregate and transaction data sets. The aggregate data set encompasses augmented data per NFT, including data from Discord and Twitter, as well as color and image complexity data. The transaction data set contains detailed information pertaining to every sale of the extracted art NFTs. Panel B details the discrete variables in both data sets, listing the unique values and the top frequent value with its corresponding frequency. The information compiled in this table stems from a rigorous data extraction process, leveraging on-chain and off-chain sources as described in Appendix A, providing a rich and diversified analysis of the NFTs, encompassing both their visual characteristics and market behaviour. As a final note, following data extraction, the variables measuring total and monthly messages and unique users in both the announcement and general channels of Discord servers have a number of missing values greater than 70%. For this reason, they will not be reported or used in the analysis. * $p \leq 0.1$; ** $p \leq 0.05$; *** $p \leq 0.01$.

Panel A: Continuous variables						
Variable	Median	σ	Minimum	Maximum	JB Statistic	KPSS test
Aggregate data set						
<i>BGCI return</i>	0	0.03	−0.26	0.22	7.91e + 04***	0.94**
<i>black</i>	0.04	0.27	0	1	1.15e + 05***	4.91***
<i>blue</i>	0.04	0.23	0	1	2.23e + 05***	14.38***
<i>floor price</i>	616.01	2.38e + 04	1	6.25e + 06	1.13e + 13***	2.49***
<i>gray</i>	0.01	0.2	0	1	6.19e + 05***	10.62***
<i>green</i>	0	0.17	0	1	1.44e + 06***	7.56***
<i>last price</i>	728.68	2.82e + 04	1	6.25e + 06	3.59e + 12***	4.29***
<i>norm shannon entropy</i>	0.29	0.12	0	0.53	1.69e + 04***	5.38***
<i>orange</i>	0.03	0.2	0	1	4.22e + 05***	17.01***
<i>purple</i>	0	0.13	0	1	6.34e + 06***	10.13***
<i>red</i>	0.05	0.21	0	1	3.96e + 05***	4.16***
<i>timediff</i>	48.94	242.59	0.01	1307.19	5.50e + 05***	16.89***
<i>white</i>	0.01	0.25	0	1	2.10e + 05***	2.66***
<i>yellow</i>	0	0.1	0	1	1.23e + 07***	1.64***
Transaction data set						
<i>average price</i>	773.42	2.23e + 04	1	6.25e + 06	1.35e + 13***	0.73***
<i>floor price</i>	479.99	2.00e + 04	1	6.25e + 06	3.08e + 13***	0.6**
<i>last price</i>	775.45	2.85e + 04	1	6.25e + 06	3.96e + 12***	0.77***

(Continues)

TABLE 1 (Continued)

Panel A: Continuous variables						
Variable	Median	σ	Minimum	Maximum	JB Statistic	KPSS test
max price	972.04	3.17e + 04	1	6.25e + 06	1.25e + 12***	1.06***
usd amount	701.84	2.41e + 04	1	6.25e + 06	8.60e + 12***	0.7***
Panel B: Discrete variables						
Variable	Unique values		Top frequent value		Frequency of top value	
Aggregate data set						
collection name	506		Foundation (FND)		21.96	
deployer creator generalities	8		Company		40.03	
deployer creator name	287		Pak		3.49	
discord account	2		1		61.85	
discord server	237		SuperRare		7.67	
followers count	408		334,093		21.96	
following count	349		854		21.96	
like count month	404		10,452		21.96	
listed count	209		1		23.35	
marketplace collection	2		0		70.09	
nft type	2		erc721		97.76	
num of colors	9		9		33.01	
num tweets	377		2397		21.96	
num of owners	156		2		77.91	
num of trades	149		1		77.87	
platform of last sale	6		OpenSea		69.49	
quote count month	380		300		21.96	
reply count month	392		1274		21.96	
retweet count month	401		1284		21.96	
twitter account	2		1		97.58	
twitter handle	417		FND		21.96	
verified	2		0		64.11	
Transaction data set						
buyer	93,184		0 × 8888888888e9997e...		0.63	
collection name	506		FND		15.64	
nft	238,420		0x73da73ef3a6982109c4.../8		0.69	
num of owners	156		2		53.82	

TABLE 1 (Continued)

Panel B: Discrete variables			
Variable	Unique values	Top frequent value	Frequency of top value
<i>num of trades</i>	149	1	53.73
<i>platform of last sale</i>	6	OpenSea	75.11
<i>seller</i>	72,618	0x8c9f364bf7a56ed058...	1.11

the financial literature, and it frequently results from complex factors, such as rarity and artist recognition. Rarity acts as a recognized determinant of value in both conventional and digital art markets, creating a perception of exclusivity and uniqueness (Renneboog & Spaenjers, 2013; Schaar & Kampakis, 2022). In the realm of NFTs, rarity's definition further refines and authenticates via the blockchain technology, thereby augmenting its attraction. Conversely, artist recognition plays an instrumental role in the valuation of art pieces, and eminent artists demand elevated prices that reflect their consolidated reputation and brand (Mandel, 2009).

Examining the discrete variables panel provides insights into the structural composition of the NFT ecosystem. From the six NFT marketplaces included in our analysis, namely, *OpenSea*, *Rarible*, *NFTX*, *LooksRare*, *Foundation* and *SuperRare*, the dominance of *OpenSea* is prevalent, capturing 69.49% of all transactions. The platform's dominance is congruent with academic research on network effects, where platforms that captivate more users gain advantages in terms of augmented liquidity and information dissemination, thereby fortifying their market dominance (Parker & Van Alstyne, 2005; Wilson et al., 2021). We observe that most NFTs are created through companies, accounting for 40.03% of total NFT creations. Notably, *Foundation* stands out by capturing 21.96% of the total number of created NFTs. This could be attributed to the user-friendly interface provided by *Foundation* and similar marketplaces, which allow NFT creation without the need for programming skills. *Foundation's* influence extends to the art NFT sales, where it accounts for 15.29% of the market. This is probably because it curates pieces from distinguished artists, which aligns with scholarly research concerning consumer behaviour, where branded collections often evoke greater trust and thus attract a larger number of buyers (Rojas-Lamoren et al., 2022).

Analogously, Table 2 illustrates the dynamics of herding behaviour in the art NFT market, providing a detailed analysis of the patterns and tendencies among investors. This table provides descriptive statistics for the continuous and discrete variables used in the study of herd bias. It examines how investor decisions are influenced by the actions of others, leading to collective market trends. Finally, Table A1 in Appendix A provides a brief description of the quantitative variables presented in Tables 1 and 2.

Following the data extraction, we rely on exploratory data analysis to identify key characteristics of our data set. By analyzing the number of sales and sales volume in USD per year, we observe that art NFTs followed the overall market trend and skyrocketed in 2021. Specifically, 2018 and 2019 account for just 0.72% of the number of sales and 0.02% of the USD volume, while 2021 accounts for 93.88% of sales and 99.21% of the volume. Focusing on 2020 and 2021, we see a similar trend when examining the average monthly volume and sales price. The volume exhibits a strong uptrend for both years, however, the average price, despite initially following a similar pattern, starts plateauing after February 2021. This indicates that the volume increase in the art market in 2021 was mostly due to an increase in the number of sales and not in price.

TABLE 2 Herding behaviour descriptive statistics.

This table reports descriptive statistics for the continuous and discrete variables used in the study of herd bias. The continuous variables are split into two data sets: the Artists data set and the Transactions data set. Each variable's median value is displayed, along with its standard deviation, minimum and maximum values. The Jarque–Bera statistic (JB Stat.), which tests for normality in the data distribution, is also provided. A triple asterisk denotes a significance level of $p \leq 0.01$. The Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test results for stationarity are also shown, with a single asterisk denoting $p \leq 0.1$. For instance, in the Artist Data Set, *sold ratio*, which shows that the average artist sells only 2% of their work, reveals a great disparity as some artists manage to sell all their work. The wide-ranging *average price* highlights the diversity in the pricing strategy across artists, with the median being 556, while some nonfungible tokens (NFTs) are priced, on average, as high as $1.35e + 06$. Trade activity, represented by the *num of trades*, is generally limited, given the fact that 25% of NFTs are sold twice. The *buyer seller pair* variable confirms this trend, as despite the top frequent value being just 1 unique pair, some artists attract thousands of unique pairs. NFT selling times, captured by *timediff*, typically do not exceed three months, but some artworks may require over three years to find a buyer. The *Volume* variable, a measure of an artist's monthly market activity, further supports this imbalance: while the median market activity is 4492.12, the most popular artists can stir activity levels up to $1.87e + 07$. *Share* suggests that the NFT market can be dominated by a few artists, with some accounting for up to 82% of the total monthly market capitalization, despite the median share being zero. Finally, the *Demand* variable summarizes the shifting nature of consumer preferences in the NFT market, with a wide range from 0.97 to 3570.16. In the Transactions data set, discrete variables such as *discord account* and *deployer artist* provide additional insights. Most artists appear to have a Discord account, while independent artists constitute the majority. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Panel A: Continuous variables						
Variable	Median	σ	Minimum	Maximum	JB Stat.	KPSS test
Artists data set						
average price	556	4.05e + 04	2.06	1.35e + 06	3.44e + 07***	0.15*
cum score	1.77e + 07	1.95e + 07	788.19	7.13e + 07	259.43***	7.9***
cum score year rank	443.5	467.63	1	1576	228.6***	7.94***
market cap	4492.12	1.01e + 06	2.41	1.87e + 07	2.08e + 06***	0.28***
demand	0.97	114.36	0	3570.16	1.85e + 07***	0.42*
score	1170.45	1.71e + 05	62.28	3.12e + 06	1.98e + 06***	0.28*
share	0	0.05	0	0.82	1.16e + 06***	0.11*
sold ratio	0.02	0.16	0	1	2.45e + 04*	0.14*
timediff	88.6	186.77	1.04	1231.58	7891.99***	0.32*
volume	4492.12	1.01e + 06	2.41	1.87e + 07	2.08e + 06***	0.28*
Transactions data set						
CSAD	0.06	0	0.03	0.1	16.29***	3.03***
market return	0.02	0.01	0	0.04	19.72***	2.92***
newcomer proportion	0.24	0.1	0.04	0.67	26.51***	1.25***
Panel B: Discrete variables						
Variable	Unique values		Top frequent value		Frequency of top value	
Artists data set						
buyer seller pair	344		1		13.09	

TABLE 2 (Continued)

Panel B: Discrete variables			
Variable	Unique values	Top frequent value	Frequency of top value
<i>deployer creator generalities</i>	2	Individual	87.08
<i>discord account</i>	2	Yes	54.07
<i>num of trades</i>	842	2	25.38

We note that 71.9% of the examined NFTs have never been sold. Of those sold, only 22.1% were sold more than once. However, NFTs that are sold more than once have a higher chance of being sold a third time, with the probability increasing to 27.1%, which further increases to 30.71% for the fourth sale.

The observed illiquidity and fragmentation in the art NFT market resonate with characteristics commonly associated with emerging markets (Amihud & Mendelson, 1986; Kirilenko et al., 2017; So & Wang, 2014) where high volatility and uncertainty prevail.

These trends suggest a dynamic market shaped by a blend of intrinsic and extrinsic factors. In particular, extrinsic factors, such as market hype and overall crypto market trends, add further complexity to pricing dynamics. Similar to phenomena observed during technology bubbles (Ofek & Richardson, 2003), the influence of social media and broader market sentiment can create rapid price swings and speculative behaviours.

3 | APPROACH

The methodology section of our analysis is split into two distinct parts, each addressing a unique aspect of the art NFT market. In the first part, we undertake an evaluation of the art NFT pricing structures. Indeed, the evolving landscape of financial markets, particularly with the emergence of new asset classes like art NFTs, has brought to the forefront a critical reassessment of traditional valuation models. In financial literature, there's a burgeoning consensus that these novel assets defy conventional categorization and valuation metrics (Chanel et al., 1996; Nadini et al., 2021; Rosen, 1974), primarily due to their distinct blend of financial, artistic and social attributes. Art NFTs, unlike traditional assets, encapsulate a unique intersection of digital scarcity, artistic expression and community-driven value, challenging the traditional paradigms of intrinsic value assessment. This divergence from established valuation norms necessitates a more adaptive and multifaceted approach to understand their market behaviour:

Hypothesis 1. A holistic and robust price index for art NFTs can be developed by integrating financial, artistic and social network features, providing a more accurate reflection of their value than traditional financial metrics alone.

This hypothesis arises from recognizing that such assets lack intrinsic fundamental value (Taleb, 2021). Our approach seeks to establish a pricing model for art NFTs that goes beyond traditional financial valuation, encapsulating the complex interplay of artistic merit and social

influence. To this end, we adopt a two-sided approach that employs both hedonic and RSR models. The hedonic pricing model is a revealed preference method used to estimate the influence of various features on the price of a good, and it is particularly suited for studying markets where the goods being transacted are highly differentiated (Rosen, 1974), like, the art NFT market. The RSR model, on the other hand, focuses on items that have been sold more than once, thus allowing us to control for the intrinsic qualities of each item and isolate the pure price effect over time. This method is designed to comprehensively capture the intricate dynamics among financial, artistic and social aspects that jointly determine the valuation of art NFTs.

The second part of this methodology section analyzes the herd bias within the art NFT market. In this setting, it becomes imperative to consider the possibility that investor rationality, a cornerstone assumption in traditional financial markets (Amihud & Mendelson, 1986; Campbell & Shiller, 1988), may not hold true in this novel and evolving landscape. The art NFT market, characterized by its nascent stage, digital nature and unique asset properties, presents a scenario where traditional notions of market behaviour and investor rationality may be challenged or redefined.

The concept of herding behaviour, which implies investors making decisions based on the actions of others rather than based on independent analysis or changes in intrinsic asset values, becomes particularly important in this context. This is compounded by the unique attributes of the NFT market, such as its relatively low liquidity and high information asymmetry (Li & Chen, 2023). These factors can create an environment conducive to herding, where investors are more likely to follow market trends and collective sentiment rather than basing their investment decisions on fundamental analysis:

Hypothesis 2. Herding behaviour significantly influences the valuation cycles of art NFTs, with investor actions conforming to group dynamics, especially in scenarios of high market volatility and information asymmetry.

We utilize the CSAD as a key analytical tool to analyze herding behaviour and artist-driven market trends within the art NFT market. CSAD is instrumental in revealing how market participants react to volatility and conform to prevalent market trends, often prioritizing these over fundamental analysis. This trend is especially noticeable in markets fraught with uncertainty and information asymmetry, characteristics intrinsic to the art NFT sector (Bao et al., 2023). The low liquidity of the NFT market can actually amplify herding effects because a few investors' actions can disproportionately influence prices, prompting others to follow (Azarmi & Menny, 2013; Banerjee, 1992; Demirer & Kutan, 2006). As such, a small number of transactions can excessively influence the market, leading to a scenario where prices comove more due to investor behaviour rather than underlying fundamentals. The strength of the CSAD model lies in its ability to identify herding by assessing the comovement of NFT prices, observing when prices move more uniformly than fundamentals would suggest. The model's efficacy lies in its ability to detect this uniformity in price movements. In a normal market scenario, where decisions are made independently, price movements of different assets would be more varied and less correlated. However, in situations where herding is present, there is a notable convergence in the way prices of different assets move, indicating that investors are likely following similar cues or strategies, rather than basing their decisions on changes in intrinsic asset values. The CSAD model's applicability in low liquidity markets like NFTs is underscored by studies that have examined herding behaviour in similar contexts, such as small-cap stocks or emerging market equities (Chang et al., 2000; Park & Sabourian, 2011; Yousaf & Yarovaya, 2022).

3.1 | Pricing dynamics

The hedonic regression approach applied to the study of art NFTs price dynamics is rooted in the consumer theory of utility maximization, where NFTs are considered bundles of characteristics, and the price reflects the consumer's willingness to pay for these characteristics (Lancaster, 1966; Rosen, 1974). This approach has been widely used to study price determinants in various markets, such as real estate (Muehlenbachs et al., 2015), traditional art (Adams et al., 2021; Aubry et al., 2023; Chanel et al., 1996; Renneboog & Spaenjers, 2013) and more recently, in the analysis of digital assets like NFTs (Borri et al., 2023; Horky et al., 2022).

For the analysis of art NFTs in this study, we commenced by employing a hedonic regression model, described by the following equation:

$$\ln(y_i + 1) = \beta_0 + \beta \mathbf{X}_i + \gamma D + \delta T + \epsilon_i, \quad (1)$$

where y_i is the average price of the NFT, \mathbf{X} represent a vector of continuous variables, D denotes control variables, T stands for time fixed effects and ϵ is the error term. Here, the continuous variables include features such as colour proportions, Shannon's entropy and floor price, reflecting both aesthetic and complexity aspects of the NFTs. The control variables represent a set of dummy variables accounting for intrinsic characteristics (e.g., the number of colours in each image) of the digital assets and presence on social media. Time coefficients are incorporated to construct a price index, accounting for temporal variations and market trends.

To correct the potential selection bias in observed price data, we performed a Heckman two-stage regression. Selection bias arises when the sample selection for observation is not random, and there is a correlation between the observable selection and the unobservable factors affecting the dependent variable (Korteweg et al., 2016). In the context of art NFTs, this bias can occur when only the prices of traded NFTs are observed, while unsold NFTs remain unobserved. This nonrandom selection can create a bias in estimating the relationship between price and characteristics of the NFTs, as the unobserved factors affecting the decision to trade might also influence the price. To correct for this, the inverse Mills ratio (IMR) is derived from the Probit estimation and encapsulates the likelihood of the NFT being traded. This IMR is incorporated into the hedonic regression as a new variable, λ , ensuring unbiased estimates:

$$\lambda = \frac{\phi(\gamma \mathbf{X})}{\Phi(\gamma \mathbf{X})}. \quad (2)$$

The model corrects for the selection bias by capturing the unobserved factors affecting both the selection process and the dependent variable:

$$\ln(y_i + 1) = \beta_0 + \beta \mathbf{X}_i + \gamma D + \delta T + \theta \lambda_i + \epsilon_i. \quad (3)$$

The two-stage approach ensures that the relationship between the price and characteristics of NFTs is estimated without the bias induced by nonrandom selection.

We addressed multicollinearity using the Variance Inflation Factor method, ensuring that the explanatory variables were not highly correlated (O'Brien, 2007). Furthermore, we assume that our data respect all the assumptions of the hedonic and Heckman models, such as linearity, homoscedasticity and the normality of error terms and residuals.

In the context of hedonic modelling for art NFT prices, the time-varying and invariant variables are crucial in determining the underlying price index. However, it is vital to ensure that the hedonic regression parameters are stable over time. If not, this instability can lead to biased estimates of time dummy coefficients and consequent price indexes (Kuminoff et al., 2010).

To mitigate this issue, we employ the chained Fisher index, derived from sectional hedonic regressions, following the methodology used in classical financial literature (Diewert, 1976; Triplett, 2004). Specifically, we construct the Fisher index by first calculating the Laspeyres and Paasche indexes and then obtaining the geometric mean of these two indexes. Given the estimated β coefficients from the hedonic regression and the normalized weights q for each characteristic, the Laspeyres index for a given period t is calculated as

$$Laspeyres_t = \frac{\sum_{j=1}^J \beta_{j,t+1} q_{j,t}}{\sum_{j=1}^J \beta_{j,t} q_{j,t}}. \quad (4)$$

Similarly, the Paasche index for the same period is determined as

$$Paasche_t = \frac{\sum_{j=1}^J \beta_{j,t+1} q_{j,t+1}}{\sum_{j=1}^J \beta_{j,t} q_{j,t+1}}. \quad (5)$$

The Fisher index is then the geometric mean of these two indexes:

$$F_{t+1} = \sqrt{Laspeyres_t \times Paasche_t}. \quad (6)$$

The Fisher index can be chained across time to generate a consistent price index for art NFTs.

In the final step, we employ RSRs. This technique focuses specifically on the items that have been sold multiple times, thereby controlling for the unobservable heterogeneity of those assets (Bailey et al., 1963; Korteweg et al., 2016; Mei & Moses, 2002). This approach offers a complementary perspective to hedonic regression by capturing price changes for identical NFTs, thus eliminating the need to control for detailed characteristics. The model is estimated with the regression:

$$\ln(y_{i,t}) = \beta_0 + \sum_{t=1}^T \beta_t D_{i,t} + \varepsilon_{i,t}, \quad (7)$$

where $y_{i,t}$ represents the ratio between the prices of two consecutive sales, β_t the coefficients for the dummy variable $D_{i,t}$ with a value of 1 in the period when the resale occurs, -1 in the period of the previous sale, and 0 otherwise and $\varepsilon_{i,t}$ is the error term.

However, a significant concern in RSRs is, again, the potential selection bias and the heteroskedasticity of the error term. Indeed, as specified before, in RSR, the sample only consists of NFTs that have been sold multiple times. This specific selection can introduce bias since NFTs that are sold more than once may not be a random subset of all NFTs. They might have unique characteristics that make them more or less likely to be resold, leading to a

systematic deviation from the overall population. To correct for this bias and account for the varying holding periods between repeat sales, we utilize the Case Shiller three-stage regression method (Case & Shiller, 1989; Korteweg et al., 2016).

We begin by conducting the RSR as previously described, extracting the residuals. This initial step lays the foundation for modelling the heteroskedasticity in the error term. Such nonconstant variability could be related to factors like the holding period, which might have varying effects on the error term across different observations. We regress the squared residuals against a constant and the specific holding periods of art NFTs:

$$u_t^2 = \gamma_0 + \gamma_1 \times \text{holding_period} + \zeta_t. \quad (8)$$

By modelling the squared residuals in this manner, we capture the pattern of heteroskedasticity and its relationship with the holding period. Finally, we fit the original RSR using Generalized Linear Models with weights derived from the reciprocal of the square root of the fitted values from the second stage.

3.2 | Herd bias

The notion of herding behaviour refers to the propensity of market participants to conform to the investment decisions of a larger group, often at the expense of their own available information. In the context of the art NFT market, this can manifest in the form of preferences aligning with recognized or leading artists, thereby creating self-reinforcing patterns of demand. Such mechanisms can lead to a superstar phenomenon where a small number of artists dominate the market. This phenomenon has been empirically investigated in traditional art markets (Rosen, 1981), but the exploration in digital markets is new.

Our analytical framework encompasses three primary dimensions: (1) artistic quality as reflected by market metrics, (2) herding behaviour rooted in classical financial theories and (3) the intricate relationship between historical performance, popularity and influence factors like social media engagement. This last dimension, particularly, provides insights into how external elements, such as an artist's Discord account presence or the artist status, interact with traditional market dynamics to shape an artist's ranking and demand. By leveraging the OLS estimation approach of the multinomial logit model of Azarmi and Menny (2013), we collectively analyze consumer choices based on the individual market shares of artists rather than focusing on separate purchasing decisions.

Our analytical representation takes the following form:

$$\log(\Pi_{it}) = \alpha_i + \alpha_t + \sum_{j=1}^J \beta_j (X_{ijt} - \bar{X}_{jt}) + \epsilon_{it}, \quad (9)$$

where the intrinsic value of an artwork created by artist i is symbolized by a_i , the j -th explanatory variable is denoted by X_{ijt} and ϵ_{it} refers to the error term of the model. The corresponding arithmetic means of X_{ijt} are represented by \bar{X}_{jt} . Π_{it} symbolizes the selection made by art enthusiasts and is rooted in the relative demand attributed to artist i at time instance t . We define this relative demand as

$$\Pi_{it} = \frac{S_{it}}{\tilde{S}_t} = \frac{\frac{V_{it}}{V_t}}{\sqrt[n]{\prod_{j=1}^n \frac{V_{jt}}{V_t}}}. \quad (10)$$

In this model, the relative demand Π_{it} is derived by dividing an artist's market share S_{it} by the geometric mean of the individual market shares across all artists, represented by \tilde{S}_t . The term S_t is ascertained by dividing the monthly market capitalization V_{it} specific to an artist by the cumulative market capitalization V_t for all n artists at time t . We define the market capitalization for each artist similarly to how it's calculated in the stock market, by multiplying the number of artworks sold by the average price of these artworks for each month and summing the total for each artist.

We assume that consumers do not have predetermined preferences for a specific artist or art style within a given sample category. For instance, an individual evaluating works from a particular NFT collection would exhibit equal interest in diverse artistic expressions, irrespective of the creator.

As an explanatory variable, the cumulative score of the annual ranking lagged by one period captures the inertia in an artist's popularity, reflecting how past performance continues to affect current standings. The yearly rank of the artist and the previous monthly rank provide insights into the temporal dynamics of an artist's market position. The interaction terms in the regression model serve specific purposes in understanding the dynamics of the art NFT market. The complexity of the art NFT market is further dissected through specific interaction terms in the model. The term involving yearly rank and Discord account (*year * discord_account*) captures the influence of social media engagement on an artist's rank, reflecting how an active presence on platforms like Discord can amplify the artist's visibility and market appeal. The interaction between yearly rank and the artist status (*year * deployer_artist*) distinguishes between artists who are independent creators and those associated with companies. In the art NFT market, the interaction effects help explain how information cascades might form. For instance, a high yearly rank, coupled with active engagement on Discord, might create a positive feedback loop where popularity leads to more visibility, further enhancing demand. Conversely, the artist's status might affect how information about the artist disseminates, with individual artists potentially benefiting from a more personalized connection with collectors.

We complement and enrich the study of herding behaviour and artist stardom by examining the CSAD. This method serves as an analytical tool for discerning how market participants adapt to fluctuations within the market. By evaluating the cumulative sum of deviations from a certain trend or mean, we identify instances of herding behaviour, where investors are inclined to follow prevailing trends rather than base their decisions on fundamental analysis. Such herding behaviour forms a common thread between the stardom effect and cascade theory, as well as the study of CSAD. These interconnected areas pertain to the manner in which investors or market participants might succumb to trends, overlook private information or make determinations grounded in the conduct of others, rather than a rational evaluation of value. In the context of NFTs, being relatively nascent and heavily influenced by social dynamics, celebrity endorsements and market hype, often leads to behaviours where individuals align with the crowd instead of pursuing independent analysis. We recognize this alignment as a phenomenon where a few artists dominate the market.

Initially, we calculate daily returns of individual art NFTs as in Bao et al. (2023) and we adapt the CSAD measure, as formulated by Chang et al. (2000):

$$CSAD_t = \frac{1}{N} \sum_{i=1}^N |R_{i,t} - R_{m,t}|. \quad (11)$$

In the NFT market, which is characterized by low liquidity, the price movements are often more pronounced due to the significant impact of trades in the market. When herding occurs, investors tend to move collectively, leading to a decrease in the dispersion of returns across different assets, that is, they start moving more in sync with each other. In such scenarios, the CSAD value tends to decrease as the absolute deviations of individual asset returns from the mean return diminish. This uniformity in price movements, despite the low liquidity, is a key indicator of herding.

The effectiveness of CSAD in capturing herding in low liquidity environments lies in its sensitivity to the convergence of asset returns. When investors herd, the correlation between individual asset returns and the market average increases, reducing the CSAD value. This contrasts with a market dominated by independent decision-making, where one would expect a higher level of dispersion in returns, and hence, a higher CSAD value. Further, to enhance the robustness of the findings and capture the dynamic interactions in the market, we extend the methodology by introducing a lagged dependent variable, as advised by Fu and Wu (2021):

$$CSAD_t = \alpha + \gamma_1 |R_{m,t}| + \gamma_2 R_{m,t}^2 + \gamma_3 CSAD_{t-1} + \varepsilon_t. \quad (12)$$

The convergence between individual and market returns in this model reflects the manifestation of artist stardom, where market preferences tend to converge towards specific artists.

To delve deeper into the dynamics of herding, we employ the fixed transition probabilities Markov-switching model (MSM) (Diebold & Rudebusch, 1999), integrating the methodology developed by Bao et al. (2023). This approach enables the distinction between different market regimes, revealing how herding behaviour fluctuates over time:

$$CSAD_t = \begin{cases} \alpha + \gamma_{1,1} |R_{m,t}| + \gamma_{2,1} R_{m,t}^2 + \gamma_3 CSAD_{t-1} + \varepsilon_{1,t} & \text{if } s_t = 1, \\ \alpha + \gamma_{1,2} |R_{m,t}| + \gamma_{2,2} R_{m,t}^2 + \gamma_3 CSAD_{t-1} + \varepsilon_{2,t} & \text{if } s_t = 2, \end{cases} \quad (13)$$

with transition probabilities expressed as

$$P(s_{t+1} = i | s_t = j) = p_{ij}, \quad i, j \in \{1, 2\}.$$

The transition matrix for this first-order Markov-switching process takes the form

$$P = \begin{pmatrix} P_{1,1} & 1 - P_{2,2} \\ 1 - P_{1,1} & P_{2,2} \end{pmatrix},$$

where $0 < P_{i,i} < 1$ and $0 < P_{j,j} < 1$. In this context, Regime 1 represents the state with herding, characterized by a negative $\gamma_{2,1}$, while Regime 2 denotes the state without herding, indicated by positive or insignificant $\gamma_{2,2}$.

The constant term α serves as a nonswitching parameter, symbolizing the long-term average of $CSAD_t$, while the transition probabilities, $P(s_{t+1} = i | s_t = j)$, govern the progression between these two states. The probability $1 - P_{2,2}$ specifically elucidates the likelihood of the emergence of herding behaviour.

TABLE 3 Hedonic and Heckman two-stage model estimations.

This table reports the estimation results of hedonic models (Models 1, 2 and 3) and Heckman two-stage models (Models 1', 2' and 3') applied to the average price of art nonfungible tokens (NFTs) collected from February 1, 2020, to February 10, 2022. The dependent variable is the logarithm of the average price plus one. The models progressively include additional controls: Models 1 and 1' do not include social network controls (*i.e.*, *verified*, *twitter account*, *discord account*, *deployer creator generalities*) or market dummy variables (*i.e.*, *marketplace collection*, *platform of last sale*, *nft type*). Models 2 and 2' introduce social network controls and Models 3 and 3' further include market dummy variables. Month fixed effects (FE) are included in all models. We acknowledge that the variables *Mar-21*, *Apr-21*, *Jul-21*, *Aug-21*, *Sep-21*, *Oct-21*, *Nov-21*, *Dec-21*, *Jan-22*, *Feb-22*, *marketplace collection*, *black* and *white* exhibited VIF values greater than 10. Upon reevaluation with these time FEs and other variables removed, the magnitudes of the regression coefficients for the vast majority of variables remained practically unchanged. This suggests that the high VIFs mainly influenced the redistribution of effects among the time FEs. Furthermore, the coefficients on the colors *gray*, *green*, *blue*, *yellow*, *purple*, *orange* and *red* are not reported in the table. These have coefficients of magnitude and sign similar to those obtained for *black* and *white*. Each entry in the table is the coefficient estimate from the respective model, with the standard error reported in parentheses. Asterisks denote significance levels: *denotes significance at the 10% level, **at the 5% level and ***at the 1% level. The number of observations for all models is 238,420. The adjusted R^2 , residual standard error and F statistics are reported at the bottom of the table. The *mills ratio* row represents the Inverse Mills Ratio, which is included only in the Heckman models (1', 2' and 3'). * $p \leq 0.1$; ** $p \leq 0.05$; *** $p \leq 0.01$.

Dependent variable: $\log(\text{average price} + 1)$	(1)	(1')	(2)	(2')	(3)	(3')
<i>const</i>	−0.495*** (0.021)	−0.204*** (0.013)	−0.494*** (0.021)	−0.181*** (0.014)	−0.738** (0.022)	−0.400*** (0.016)
<i>black</i>	−0.266*** (0.011)	−0.264*** (0.011)	−0.259*** (0.011)	−0.257** (0.011)	−0.264*** (0.011)	−0.261*** (0.011)
<i>white</i>	−0.257*** (0.011)	−0.257*** (0.011)	−0.252*** (0.011)	−0.250*** (0.011)	−0.248** (0.011)	−0.247*** (0.011)
<i>Norm_Shannon_entropy</i>	−0.006 (0.009)	−0.007 (0.009)	0.025*** (0.009)	0.020** (0.009)	0.047** (0.009)	0.042*** (0.009)
<i>Num_of_colors</i>	0.001 (0.000)	0.001 (0.000)	0.001 (0.000)	0.000 (0.000)	−0.001* (0.000)	−0.001 (0.000)
<i>BGCI_return</i>	0.086*** (0.023)	0.252*** (0.052)	0.077*** (0.023)	−0.027 (0.053)	0.070** (0.023)	−0.001 (0.053)

TABLE 3 (Continued)

Dependent variable: $\log(\text{average price} + 1)$	(1)	(1')	(2)	(2')	(3)	(3')
$\log(\text{number_of_trades} + 1)$	1.065*** (0.003)	1.067*** (0.003)	1.068** (0.003)	1.071*** (0.003)	1.094*** (0.003)	1.097*** (0.003)
$\log(\text{floor_price} + 1)$	0.935*** (0.001)	0.936*** (0.001)	0.931*** (0.001)	0.932*** (0.001)	0.925** (0.001)	0.927*** (0.001)
$\log(\text{following_count} + 1)$	−0.003*** (0.000)	−0.003*** (0.000)	−0.007** (0.000)	−0.007*** (0.000)	−0.003*** (0.000)	−0.004*** (0.000)
$\log(\text{listed_count} + 1)$	0.024*** (0.000)	0.023*** (0.000)	0.025** (0.001)	0.025*** (0.000)	0.018** (0.000)	0.017*** (0.001)
$\log(\text{quote_count_month} + 1)$	0.001 (0.001)	−0.000 (0.001)	−0.015*** (0.001)	−0.015*** (0.001)	−0.008*** (0.001)	−0.008*** (0.001)
$\log(\text{timediff} + 1)$	0.030** (0.001)	0.030** (0.001)	0.023*** (0.001)	0.023*** (0.001)	0.024*** (0.001)	0.024*** (0.001)
<i>Mills_ratio</i>		−0.044*** (0.012)		0.023* (0.012)		0.015 (0.012)
Social network controls		×		✓		✓
Market FE		×		×		✓
Month FE		✓		✓		✓
Other colors controls		✓		✓		✓
Observations				238,420		
Adjusted R^2	0.933	0.933	0.933	0.933	0.934	0.934

4 | DISCUSSION OF RESULTS

4.1 | Price indices design

For the construction of hedonic indices, we select sold non-wash traded NFTs starting from January 1, 2020, due to the rarity of the transactions before that date, and we are left with 238,420 unique NFTs. We identify several significant pricing determinants for art NFTs. Table 3 illustrates our hedonic regression results. A pivotal outcome, observed consistently across all our models, is the positive and significant coefficient on the log transformed variables *number_of_trades* and *floor_price*, with values around 1.07 and 0.93, respectively. This suggests that NFTs with a higher number of trades and a higher floor price are likely to command a higher average price. This finding aligns with existing literature on asset pricing, such as Campbell and Shiller (1988) and Fama and French (1993), highlighting the relevance of these factors in financial markets. In particular, the variable *floor_price* may be capturing the price floor effect, suggesting that the minimum acceptable price for the NFT could be a significant determinant of its average price. As Kahneman and Tversky (1979) illustrate, this minimum price often sets a psychological benchmark or reference point for transactions, influencing final prices. In a similar vein, the role of a base price or minimum price as a determinant of the final price of an asset has been emphasized by Campbell and Shiller (1988). In the context of our study, a higher floor price could convey a higher perceived value of the NFT, resulting in a higher average price. Further, the floor price might encapsulate the price floor effect. In auction theory, Milgrom and Weber (1982) underscore the importance of the reserve price (the lowest price at which the seller is willing to sell the item) in determining the final price. This concept could apply to NFTs as well, where the floor price acts as a reserve price and impacts the final price of the NFT.

On the other hand, the negative coefficients on the log transformed variables *following_count* and *quote_count_month* are indeed unexpected, given that one would typically expect higher levels of social media activity to positively influence the price of NFTs. However, our regression results suggest that this is not the case for the art NFT market, at least within the context of the observed data. One possible explanation for these counterintuitive results is the over-saturation hypothesis: When an NFT or a collection of NFTs is excessively quoted or posted about on social media, it could lead to an over-saturation of information. This phenomenon could reduce the perceived rarity associated to the NFT, thereby driving down its price. This idea aligns with the findings of Dellarocas (2003) and BabićRosario et al. (2020), which discuss the potential downsides of excessive online exposure.

The application of Heckman's two-step model, as exhibited in models (1'), (2') and (3'), reveals insights about the potential selection bias in the study of art NFT pricing dynamics. The IMR (λ) for the model (1'), derived from the Probit model in the first step of the Heckman process, is significant and negative in magnitude, confirming the presence of selection bias in the sample of traded NFTs. The negative sign of the IMR suggests that the unobserved factors influencing the decision to trade NFTs are inversely correlated with the error term in the average price equation. This implies that those NFTs less likely to be traded (due to unobserved factors) tend to have higher prices than what could be predicted based solely on their observable characteristics. Higher-priced NFTs, which are less likely to be traded, might contain superior artistic or unique features not captured by the observable characteristics.

In terms of the robustness of the results, the consistency of most variable coefficients across both the hedonic and Heckman models indicates that our findings are resilient to adjustments

for selection bias. However, some variables, such as the return on the Bloomberg Galaxy Crypto Index (*BGCI_return*), demonstrate significant changes in their coefficients when transitioning from the hedonic models to the Heckman models. The coefficients' substantial increase in models (1) and (1') suggests that the relationship between these variables and NFT prices is subject to the influence of selection bias. As such, the return on the Galaxy Crypto Index could be capturing some of the unobserved factors affecting the decision to trade NFTs. Previous studies have found a significant impact of selection bias on the relationship between asset prices and market-wide indicators (Ang et al., 2020; Bianchi & Babiak, 2022). The pronounced positive coefficient in our model suggests a correlation between the performance of cryptocurrencies, as measured by the Bloomberg Galaxy Crypto Index (BGCI), and the pricing of NFTs.

The application of the RSR model drastically reduced the number of observations from 238,420 to 107,101. This is expected as the RSR model only considers assets that have been traded more than once, which is a crucial characteristic of the RSR model that leverages repeated sales data to control for the unobserved time-invariant characteristics of the assets. Figure 1 draws the estimated art NFT price indices.¹ We observe that all four indices—the hedonic, hedonic-Heckman, repeat sales and repeat sales–Case Shiller—display a similar trend of price growth over the period from August 2020 to February 2022. There is a particularly noticeable spike in price growth around February 2021. This surge in price indices coincides with several significant events in the NFT market. For instance, in March 2021, a digital artwork by the artist Beeple was sold as an NFT at Christie's auction for a staggering \$69 million. This sale not only marked a historic moment for digital art but also significantly boosted the perceived value of NFTs. Furthermore, in the same month, the first-ever tweet by Twitter CEO Jack Dorsey was sold as an NFT for \$2.9 million. These high-profile sales events likely contributed to the increased attention and perceived value of NFTs, which is reflected in the sharp increase in the price indices around this time.

In the subsequent period, from April to June 2021, the indices had a sudden drop. This could be linked to several notable developments in the NFT market. Between February 2021 and June 2021, the NFT market, experienced a significant decline due to an unsustainable surge in prices driven by speculative investments, which led to a market correction as initial excitement waned. This period also saw increasing regulatory scrutiny and uncertainty, with the rapid advancements in NFT technology outpacing regulatory frameworks, creating instability and eroding investor confidence. Concurrently, the broader cryptocurrency market experienced a downturn, with major cryptocurrencies like Bitcoin and Ethereum seeing significant price drops, which further impacted NFT transactions. Additionally, the market was flooded with low-quality projects and a growing realization that many NFTs had little intrinsic value, leading to oversupply and decreased demand, causing many NFTs to become virtually worthless. The sharp spike in the indices around July 2021 coincides with significant investments in the NFT space. At that time, OpenSea, one of the largest NFT marketplaces, announced that it had raised \$100 million in a Series B funding round led by AH Capital Management. This

¹Since the RSR model does not effectively handle sales occurring within the same time period—relying instead on time variation between sales to identify price changes and control for time-invariant characteristics—we acknowledge this limitation in our data. To mitigate this, we constructed repeat sales indices on monthly, weekly, and daily bases. We opted for the daily index because it minimized the number of discarded transactions (35.50%). In Figure 1, we aggregated the repeat sales indices to a monthly level by averaging the parameters within each timeframe.




FIGURE 1 Art nonfungible token (NFT) price indices. This figure presents the four different art NFT price indices: the hedonic, hedonic-Heckman, repeat sales and repeat sales-Case Shiller. These indices trace the evolution of art NFT prices from February 2020 to February 2022. The hedonic model, for instance, evaluates the prices of art NFTs based on their observable characteristics, while the repeat sales model focuses on price changes for identical art NFTs traded over time. The hedonic-Heckman model augments the hedonic model by correcting for potential selection bias. The repeat sales-Case Shiller model extends the repeat sales model by accounting for both heteroskedasticity and selection bias. From August 2020 to February 2022, the hedonic, hedonic-Heckman, repeat sales and repeat sales-Case Shiller indices all displayed similar price growth trends, with a noticeable spike around February 2021. This spike appears to coincide with significant events in the NFT market, such as Beeple's \$69 million NFT sale at Christie's and the \$2.9 million sale of Jack Dorsey's first tweet, potentially boosting NFT values. From April to June 2021, the indices dropped, likely resulting from a market correction driven by speculative investments, regulatory scrutiny and a broader cryptocurrency downturn. In July 2021, a spike in indices may be attributed to OpenSea's \$100 million funding round. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/efm.12566)]

investment valued the company at \$1.5 billion. This massive investment likely boosted market confidence and contributed to the surge in NFT prices reflected in the price indices.

After correcting for selection bias, Figure 1 shows that the repeat-sales-Case Shiller model, which corrects for both heteroskedasticity and selection bias, is slightly more volatile, particularly after February 2021. This could be attributed to the increased sensitivity of this model to changes in repeat-sales data. This unique sensitivity stems from the model's design, which leverages repeat-sales data to capture pure price changes, while simultaneously accounting for potential heteroskedasticity and selection bias. As a result, any substantial variations in repeat-sales data are likely to be more pronounced in this model, leading to increased volatility in the derived price index. Overall, from June 2021 to October 2021, the Repeat Sales index was generally above the Hedonic index, possibly due to speculative trading where investors resold NFTs at inflated prices driven by enthusiasm and potential speculative bubbles. In contrast, the Hedonic index remained relatively stable, suggesting that the intrinsic characteristics of NFTs didn't change much in value.

The hedonic and RSR offer robust analytical frameworks for understanding the pricing dynamics of the digital assets in our sample. However, they are not exempt from potential issues. A significant concern in these models is the time instability of the parameters, which can lead to biased estimates of the price indices (Kuminoff et al., 2010). This instability can stem from changing market dynamics, evolving consumer preferences or shifts in the mix of characteristics present in the traded NFTs over time.

The Laspeyres and Paasche indices, as shown in Figure 2a, provide a clear picture of the monthly price dynamics of the NFT market. Despite observable fluctuations in the market, these indices remain relatively stable and mirror each other closely throughout the sample period. This stability suggests that the cost of purchasing a specific bundle of NFT characteristics has remained consistent over time, a finding in line with the assumptions of the hedonic model (Rosen, 1974). However, a significant boom in the NFT market was observed starting from September 2021. Despite this surge, the Laspeyres and Paasche indices remained stable, prompting us to conduct a more detailed investigation into the underlying causes.

We plot the changes in the β coefficients and quantities (intended as the median of characteristics in our data set) over time in Figure 2b. The β coefficients, indicative of the implicit prices of NFT characteristics, show substantial variability throughout the period under consideration. Conversely, the quantities of the NFT characteristics, denoted by q , follow a

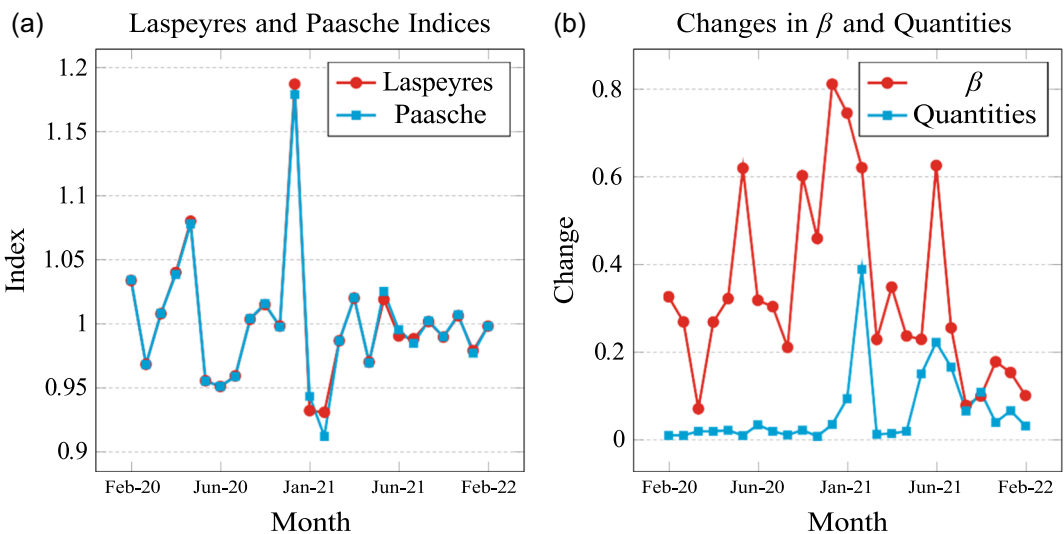


FIGURE 2 Fisher index. (a) The Laspeyres and Paasche indices (components of the overall Fisher price index) over the examined period from February 2020 to February 2022. These indices reflect the monthly price dynamics of the nonfungible token (NFT) market. Despite observable fluctuations in the market, particularly a significant surge in September 2021, these indices remain relatively stable. This stability suggests a consistent cost for a specific bundle of NFT characteristics, underpinning the assumptions of the hedonic model. (b) The changes in the β coefficients and quantities over the same time period. The β coefficients, which represent the implicit prices of NFT characteristics, show substantial variability, capturing the evolving nature of the NFT market and changes in consumer preferences. Conversely, the quantities, denoted by q , remain relatively constant, indicating a stable mix of NFT characteristics sold over time. This contrast between the variability of β coefficients and the stability of q suggests that the changes in prices are generally balanced by changes in quantities. This balance further highlights the stability of the Laspeyres and Paasche indices, and consequently the Fisher index, even during the market boom in September 2021. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

smoother trajectory, suggesting that the mix of characteristics in the sold NFTs has remained relatively constant over time. This stability in quantities, coupled with the variability in β coefficients, causes the steady nature of the Laspeyres and Paasche indices during the observed market boom. Essentially, the stability of the Fisher index, which is a geometric mean of the Laspeyres and Paasche indices, implies that changes in prices (β) are generally balanced by changes in quantities (q).

The graphical representation of the number of trades and floor price trends in Figure 3 further underscores the dynamic nature of the NFT market. We observe a significant increase in the number of trades, or market activity, towards the end of the period, which coincides with a rise in the floor price. This pattern suggests a potential correlation between market activity, prices and the quantities of characteristics traded. In other words, as the volume of trade and average prices increased (evidenced by the rise in the floor price), the quantities of NFT characteristics sold also increased, effectively stabilizing the price indices.

4.2 | Herding behaviour

Following classical financial theories, we believe that consumers tend to follow the crowd, aligning their preferences and demands with popular artists, thereby creating a cascade effect. This herding behaviour is manifested in the NFT market, leading to a superstar phenomenon where a handful of artists dominate the market, as proposed by Rosen (1981).




FIGURE 3 Number of trades and floor price trends. This graph presents the evolution of the number of trades and floor price trends in the art nonfungible token (NFT) market from February 2020 to February 2022. The $\log(\text{num of trades} + 1)$ (depicted in red) and $\log(\text{floor price} + 1)$ (depicted in blue) are represented on a normalized, monthly averaged scale. An observable upswing in both metrics is evident from around October 2021, suggesting a significant market boom during this period. This upward trend likely reflects evolving consumer preferences, technological advancements, increased public awareness and broader market acceptance of NFTs as digital assets. Importantly, as trading volume and average prices increased, so did the quantities of NFT characteristics sold, effectively balancing the price indices. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/efm.12566)]

We employ a comprehensive ranking system for artists. This system works by considering a diverse range of market metrics and social factors which encompass various aspects of an artist's presence in the NFT market, allowing us to evaluate the artists' market positions and popularity. The ranking process is performed monthly to reflect the dynamic nature of the NFT market. Specifically, we leverage the previous analysis database and select all deployers based on their Twitter handles who are not companies or marketplaces. We group the NFTs by artist, leading to a total of 413 unique NFT artists based on Twitter handles.

We factor in the ratio of an artist's sold to her unsold artworks, which gives us an understanding of the market demand for her works. The higher the ratio, the greater the demand for the artist's work. Similarly, we account for the average price of sold NFTs. A higher average price typically indicates a higher perceived value or desirability of an artist's work. The market capitalization for each artist, calculated on a monthly basis, is also considered. This measure, defined as the total value of her sold artworks within a given month, provides an insight into the artist's overall market presence and financial impact. We also take into account the average number of trades for each sold NFT. This measure provides an indication of the liquidity and secondary market activity of an artist's work. More trades generally suggest greater interest and engagement from collectors and investors. Additionally, we look at the number of unique buyer–seller pairs, which gives us an idea of the diversity of interactions associated with an artist's works. A greater number of unique pairs usually indicates a broader collector base and more varied market activity. Lastly, we consider the selling time of each sold NFT. This measure indicates the average time it takes for an artist's work to sell. Shorter selling times often suggest higher demand and faster market turnover. By considering all these equally weighted factors, we compute a composite score for each artist. This score, where each factor has an equal weight, serves as the basis for our monthly ranking of artists.

To contextualize the variables employed in our study, we start with an examination of the descriptive statistics illustrated in Table 2. The substantial variation in *volume*—with a median value of 4492.12—exposes the disparity between artists, highlighting an uneven playing field where a few artists garner significant market activity, while the majority struggle to gain traction. As this variable reflects the market activity for individual artists in a given month, this could potentially signal illiquidity issues, where the market is concentrated around a small number of artists, akin to the “superstar” phenomenon. Moreover, *share*, which represents the proportion of the total monthly market capitalization attributable to an individual artist, further accentuates this fragmentation. The median value of 0 indicates that the market is disproportionately dominated by a few artists, suggesting a significant skew in market power. Lastly, the *demand* variable, representing an artist's popularity relative to other artists, underscores the fluctuating and potentially illiquid nature of the art NFT market. With a median value of 0.97 and a wide range of values, it demonstrates how rapidly consumer preferences can shift in the art NFT market. This suggests a market prone to sudden swings and potentially fragile liquidity, as demand can quickly pivot from one artist to another.

Table 4 presents the findings from our regression models, dividing the artists into three distinct categories based on their monthly rankings. Artists are grouped by the month they are active, and their rankings are evaluated within these monthly groups. They are then divided into three quantile-based categories based on their rank, ensuring that each category has roughly the same number of artists where possible. Since rankings can vary from month to month due to changes in the metrics that determine an artist's rank, it is common for an artist to appear in different categories across different months. This method allows for each group to contain about 780 artists, even though there are only 413 unique artists, because the

TABLE 4 Least square estimation herd bias.

This table reports the findings from a comprehensive regression analysis based on a least square estimation of the model of Azarmi and Menny (2013), investigating the determinants of demand for artworks in the nonfungible token market. We divide artists into three categories according to their monthly rankings: the highest-ranked artists (Model 1), midtier artists (Model 2) and lower-tier artists (Model 3). The dependent variable in each model is the logarithm of demand for artworks by each artist category. A diverse array of explanatory variables is included, reflecting historical performance metrics, artist rankings and social media engagement. Yearly rankings are consistently negative and significant, indicating that higher rankings correlate with lower demand, suggesting a preference for niche or emerging artists. Furthermore, lagged monthly rankings show that previous high ranks reduce current demand for top and bottom artists but increase it for midtier artists, likely due to a “momentum effect”. Nevertheless, engagement on Discord does not boost demand for low-tier artists, possibly due to a smaller follower base, but enhances visibility and demand for midtier artists. Finally, independently deployed works by top-ranked artists see higher demand, reflecting a preference for authenticity. However, this strategy is less beneficial for midtier artists, likely due to the lack of support systems. Each entry in the table delineates the coefficient estimates, with standard errors provided in parentheses. Artist and month fixed effects (FE) are included to control for unobserved time-invariant artist-specific characteristics and common temporal shocks. The total number of observations, the adjusted R^2 and the F statistic are reported for each artist category. Significance levels are denoted as follows: * $p \leq 0.1$; ** $p \leq 0.05$; *** $p \leq 0.01$

	Dependent variable: $\log(demand_{it})$		
	(1)	(2)	(3)
<i>const</i>	2.9905*** (0.0189)	−0.0320*** (0.0124)	−2.9255*** (0.0233)
<i>cum score year rank_{t−1}</i>	−0.0002* (0.0001)	−0.0001 (0.0001)	0.0001 (0.0001)
<i>year rank_t</i>	−0.0087*** (0.0006)	−0.0058*** (0.0002)	−0.0073*** (0.0003)
<i>month rank_{t−1}</i>	−0.0017* (0.0012)	0.0003 (0.0006)	−0.0008 (0.0009)
<i>year * discord account_t</i>	−0.0001 (0.0003)	0.0002* (0.0001)	−0.0001 (0.0001)
<i>year * deployer artist_t</i>	0.0013** (0.0005)	−0.0005*** (0.0002)	0.0000 (0.0002)
Artist FE	✓	✓	✓
Month FE	✓	✓	✓
Observations	787	777	796
R^2	0.7346 297.30***	0.7855 329.64***	0.6280 172.53***
F statistic (robust)	(df = 538)	(df = 450)	(df = 511)

same artists can appear in multiple categories across various months. The categories are as follows: the highest-ranked artists (Model 1), the midtier artists (Model 2) and the lower-tier artists (Model 3).

We find that $year_rank_t$ is negative and statistically significant across all models, suggesting that artists with higher yearly rankings have lower demand, potentially indicating art enthusiasts' preference for niche artists or their inclination to discover new talent rather than following the crowd.

The variable $month_rank_{t-1}$ further underscores the complexity of the art choice dynamics. It is negative for the highest- and lowest-ranked artists (Models 1 and 3), suggesting that a higher rank in the previous month decreases the current demand. However, the relationship is positive for midtier artists (Model 2), which could be attributed to a "momentum effect" where midranking artists who perform well in 1 month continue to attract attention in the subsequent month.

The interaction term $year_t * discord_account$ although not significant for the lowest-ranked artists (Model 3), suggests that active engagement on Discord does not necessarily boost demand for these artists' works. This observation might seem counterintuitive, especially considering the common belief that social media engagement generally increase artist visibility and market appeal, as highlighted by Bao et al., 2023. However, it could be that for low-tier artists, the benefit of additional exposure on platforms like Discord are limited by factors such as a smaller existing follower base or less recognition in the broader art market. In contrast, for mid-tier artists, where the interaction term is positive and significant, engagement on Discord likely plays a crucial role in enhancing visibility and demand. This suggests that mid-tier artists, who may already have some market presence but have not yet reached saturation, can significantly leverage social media to broaden their audience and increase market appeal.

Turning our attention to the interaction term $year_t * deployer_artist$, we find that it is positive for the highest-ranked artists (Model 1). Artists who independently deploy their work, as opposed to those associated with or acting on behalf of a company, are met with higher demand. This finding could be indicative of art enthusiasts' preference for the authenticity, creativity and individuality that independent artists often embody. However, the dynamics shift when we look at midtier artists (Model 2). For these artists, the term is negative, implying that independent deployment might not be as beneficial. This could be due to the additional challenges and responsibilities that come with operating independently. Lesser-known artists might lack the resources, network or reputation that companies provide, thereby facing hurdles in establishing themselves in the competitive art NFT market. Furthermore, art enthusiasts may perceive company-associated artists as more reliable or credible.

Our findings support the hypothesis of herding behaviour. The yearly score shows a clear and consistent picture: The auction performance of artists (as represented by their rank) has a profound influence on consumer choice. This is particularly pronounced for the high and mid-tier artist samples (Model 1 and Model 2), where herding has a stronger impact on consumer behaviour, and the effect of current performance is even more apparent.

The clear result obtained from our regression analysis is the conspicuous inequalities in the art NFT market: the choices of artworks are highly skewed towards the highest and mid-ranked artists. This result implies a heavy concentration of market influence amongst a small number of artists. Indeed, the presence of stardom in art markets is characterized by the fact that a select group of artists disproportionately commands both attention and revenue. These "star" artists dominate the marketplace, often overshadowing their less well-known peers. Their works are sought after by collectors and command high prices, contributing to a skewed

distribution of wealth within the artist community. To visually illustrate this concentration, we turn to the 2021 market capitalization Lorenz curve for the top 200 artists in Figure 4a. The Lorenz curve's substantial deviation from the equality line—indicative of a scenario where every artist shares an equal proportion of the market cap—provides a clear visualization of the pronounced inequality in the distribution of market capitalization among artists. The curve's steep incline towards the end of the distribution highlights that a small fraction of artists command a disproportionately large share of the market cap, reinforcing our regression findings.

Figure 4b shows the Gini Index, a single measure derived from the Lorenz curve that quantifies the degree of inequality over time. The relatively high and fluctuating Gini Index values across all three samples underscore a persistent and varying level of inequality in the

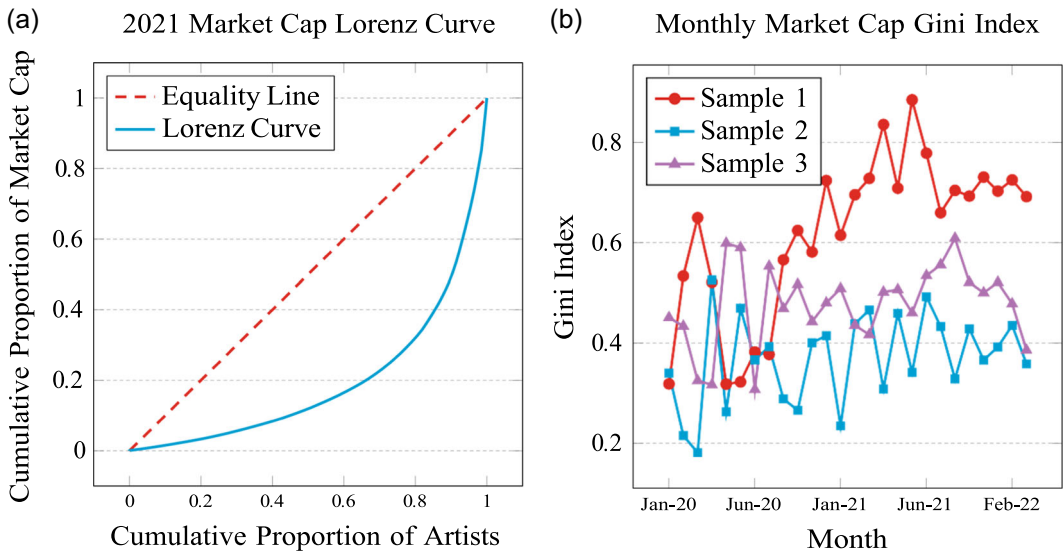


FIGURE 4 Lorenz curve and Gini index for art nonfungible tokens (NFTs). In the analysis of the 2021 Market Cap Lorenz Curve for the top 200 artists, a striking deviation from the line of equality is evident. This line of equality serves as a baseline, symbolizing an idealistic scenario where each artist enjoys an equal proportion of market capitalization—an equitable distribution of wealth and influence. The Lorenz curve, however, tells a different story, one marked by stark disparities and profound inequality. The Lorenz curve's x-axis measures the cumulative proportion of artists, a metric that ranges from the least to the most successful. The y-axis quantifies the cumulative proportion of the market capitalization, thereby reflecting the accumulated wealth and influence within the artist community. The substantial bowing of the Lorenz curve away from the line of equality indicates the pronounced inequality in the distribution of market capitalization among artists. The curve takes a steep ascent towards the end of the distribution, suggesting that a minor fraction of artists wield a disproportionately large share of the market cap. The monthly market capitalization Gini Index investigates the degree and temporal dynamics of this inequality. The Gini Index, a derivative measure from the Lorenz curve, encapsulates the inequality within a single, powerful metric. Plotted over time, it offers a chronological snapshot of inequality, tracing its ebb and flow through different periods. Three different samples were analyzed across a series of months (top, mid and low tier artists), each revealing its own tale of inequality. However, the common thread across all three is the high Gini Index values that fluctuate over time, underscoring a persistent and variable level of inequality in the market capitalization distribution among artists. Among these, Sample 1 exhibits the most pronounced inequality, reflecting the dominance of the highest-ranked artists. (a) Lorenz curve and (b) Gini index. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/efm.12566)]

market capitalization distribution among artists. The most pronounced inequality is observed in the top-tier artists sample, which aligns with our regression analysis indicating the dominance of the highest-ranked artists.

Building on our prior analysis, we aim to identify if and when herding behaviour is likely to occur in the art NFT market of high and mid-tier artists, especially during periods of significant market fluctuations. For a reliable study of CSAD, it is crucial to have a sufficient number of active artists to enhance liquidity and ensure robust results. Therefore, we restrict our sample to artists who have performed at least 10 sales. Additionally, we focus on data starting from January 1, 2021, as this was the year with the majority of NFT transactions. This approach leads to a total of 405 artists, ensuring that our findings are based on a well-defined and active segment of the market.

The ordinary least squares (OLS) regression results in Table 5 suggest that market volatility, represented by the absolute market return, has a positive influence on the CSAD. The significant positive coefficient indicates that increased market volatility is associated with increased herding behaviour: when the market experiences significant ups and downs, investors are more likely to follow the crowd, potentially leading to inflated or deflated asset prices. The lagged CSAD also has a significant positive coefficient, suggesting a persistence of herding behaviour over time. This implies that if herding behaviour was present in the past, it is likely to continue into the future, potentially creating self-reinforcing trends in the NFT market. Most interestingly, however, is the non-significant coefficient of the squared market return, which is unexpected since we typically expect a negative coefficient here. This anomaly suggests that our initial model might not fully capture the underlying dynamics, warranting further analysis. To investigate this further, we use the Markov Switching Model. Essentially, during times of intense market fluctuations, there is a tendency for the dispersion among individual NFT returns to decrease. This suggests that in high-volatility periods, investors are likely following a common narrative or sentiment, causing their trading behaviours to align. Thus, the individual returns of these art NFTs begin to converge around a central tendency, reducing the CSAD (Drehmann et al., 2005; Economou et al., 2011; Park & Sabourian, 2011).

In the MSM, the coefficients under Regime 1 (where herding is present) and Regime 2 (where herding is absent) provide more specific insights into the dynamic nature of herding behaviour in the NFT market. We notice under Regime 1 that the squared market return exhibits a negative coefficient compared with its OLS counterpart. This finding is in alignment with the hypothesis that herding is amplified under extreme market conditions. More specifically, in periods of elevated market returns, the prevalence and intensity of herding seem to be exacerbated. This reinforces the notion that market participants are especially prone to follow the crowd when the stakes are high, subsequently reducing the CSAD of individual asset returns. In contrast, Regime 2 features a positive coefficient for the squared market return, although it is poorly statistically significant. This suggests that extreme market conditions could serve as a deterrent to the formation of herding behaviour when such a behaviour is not already manifest. The model's transition probabilities indicate that once herding is initiated (Regime 1), it tends to persist over time. In this case, herding can be self-reinforcing due to factors, such as information cascades or social influence.

The absolute market return parameter in the MSM serve as complementary layers of interpretation. The significant positive coefficient on absolute market return in Regime 1 suggests that herding is more prominent in periods of higher returns. This can be interpreted as a form of "return-chasing" behaviour, where traders and investors flock to trending assets, thereby enhancing herding.

TABLE 5 Ordinary least squares (OLS) and Markov switching models (MSMs) for herd bias.

This table reports the coefficients of the OLS regression and the MSM to investigate the presence and dynamics of herding behaviour in the art nonfungible token market. The dependent variable is the Cross-Sectional Absolute Deviation (CSAD), a measure of herding. The first column shows the OLS regression results, illustrating the overall dynamics of herding behaviour in the market. Both the absolute market return ($|R_{m,t}|$) and lagged CSAD ($CSAD_{t-1}$) have positive coefficients, indicating that market volatility and past herding behaviour increase the tendency for herding. The next two columns display the coefficients under two different regimes in the MSM: Regime 1, which signifies the presence of herding, and Regime 2, which signifies the absence of herding. In Regime 1, the negative coefficient of the squared market return confirms and reinforces the relationship between market returns and herding. In Regime 2, the squared market return displays a positive, though poorly significant, coefficient, suggesting the difficult formation of herding behaviour when it is not already present. The last column shows the transition probabilities ($p[1 \rightarrow 1]$ and $p[2 \rightarrow 1]$) in the MSM, providing insights into the persistence and transition dynamics of herding behaviour. The table contains 405 observations. The model's goodness of fit is represented by the log-likelihood, Bayesian Information Criterion (BIC) and Akaike Information Criterion (AIC). Significance levels are denoted as follows: *** $p \leq 0.01$; ** $p \leq 0.05$; * $p \leq 0.1$.

	Dependent Variable: $CSAD_t$			Nonswitching Parameters
	OLS	Markov switching model		
		Regime 1	Regime 2	
$const$	0.0006*** (0.000)	0.0012*** (0.000)	0.0003* (0.000)	
$ R_{m,t} $	0.1704*** (0.045)	1.6929*** (0.080)	0.0702* (0.037)	
$R^2_{m,t}$	0.0830 (0.898)	-7.1029*** (0.017)	0.2868 (0.759)	
$CSAD_{t-1}$	0.9739*** (0.007)	0.4864*** (0.029)	0.9507*** (0.011)	
$p[1 \rightarrow 1]$				0.8115*** (0.081)
$p[2 \rightarrow 1]$				0.0185*** (0.008)
Observations		405		
Log likelihood	2064.3		2148.75	
BIC	-4105		-4231.45	
AIC	-4121		-4257.49	

Figure 5 effectively demonstrates the temporal oscillation between regimes of herding behaviour and independent decision-making within the NFT market, as postulated by our model. Specifically, the shaded areas highlight periods when investors, driven by market momentum, decide to follow the crowd, leading to an increase in herding behaviour. These instances are closely aligned with our MSM findings, which reveal a negative relationship between squared market returns and CSAD, hence indicating increased herding behaviour during market volatility.

The shaded regions in Figure 5 align with notable events in the NFT space that likely incited increased market activity and investor interest. For example, Beeple's \$69 million




FIGURE 5 Dynamics of market return and regime switching. This graph provides a temporal illustration of the nonfungible token (NFT) market returns and the corresponding prevalence of herding behaviour. The shaded areas represent periods of herding behaviour and correlate with significant market returns. This alignment indicates a positive relationship between absolute market returns and Cross-Sectional Absolute Deviation, offering a nuanced visual representation of herding behaviour during periods of market volatility. The graph showcases oscillation between herding and independent decision-making, reflecting the complex, nonlinear relationship between market returns and herding behaviour. For instance, high but not extreme market returns still induce herding behaviour, as indicated by the shaded areas. These shaded periods, driven by market momentum, correspond to key events in the NFT space that amplified market activity. In March 2021, the \$69 million auction of Beeple's artwork led to increased trading activity. Similarly, the "Loot (for Adventurers)" project in September 2021 ignited a surge in market returns and herding behaviour. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/efm.12566)]

auction watershed event in March 2021 attracted significant media attention and public interest, leading to increased trading activity and a surge in market returns. Similarly, the shaded region in September 2021 coincides with the period when the "Loot (for Adventurers)" project gained significant traction. This project led to a surge in trading activity, manifesting as increased market returns and herding behaviour.

To further enhance our understanding of the mechanisms that drive herding behaviour in the art NFT market, we will investigate in the following how the coefficient on the squared market return is modulated by the presence of investor newcomers and fluctuations in the BGCI. Specifically, we seek to assess whether, and to what extent, the percentage of newcomers in the market and variations in BGCI returns impact the strength and direction of the coefficient on squared market returns, which serves as our proxy for herding behaviour. To do so, we use a weighted least squares (WLS) regression model. We first calculate the proportion of newcomers in the market for each date, defining a newcomer as an investor who has traded only once. This proportion is then merged with the existing dataset containing market returns and CSAD values. We perform weighted regressions across a grid of newcomer proportions, using Silverman's Rule of Thumb to determine the bandwidth for our kernel weights. This allows us to assess how the coefficient on squared market returns varies with different levels of newcomer

presence, providing a nuanced view of herding behavior under varying market conditions. The WLS model is well-suited for this analysis as it accounts for heteroskedasticity and provides more reliable estimates by giving different weights to observations based on the newcomer proportion. By plotting the coefficients and their confidence intervals across the newcomer proportion grid, we can visualize and interpret the impact of newcomers and BGCI fluctuations on herding behavior in the NFT art market.

Drawing on insights from Bao et al. (2023), we know that experienced investors tend to secure higher returns per unit of cryptocurrency invested compared with their inexperienced counterparts. This disparity may be attributed to inexperienced investors, often newcomers, purchasing NFTs at higher average prices, possibly due to their limited familiarity with market valuations and dynamics. Given that the majority of these newcomers engage in trading activities only once, their lack of market understanding makes them susceptible to following market consensus. When there is an influx of such newcomers into the art NFT market, the propensity for herding behaviour to manifest is likely heightened.

Simultaneously, the rationale for incorporating BGCI returns stems from the observed high correlation between cryptocurrency markets and NFT markets. Fluctuations in the BGCI can serve as a proxy for general market sentiment within the broader cryptoasset ecosystem, which invariably influences the art NFT space. Thus, BGCI returns can either amplify or attenuate herding tendencies, depending on the prevailing market conditions.

Figure 6a illustrates the relationship between newcomer proportion and the coefficient γ_2 of Equation (13). Initially, when the share of newcomer is lower than 0.30, γ_2 touches 0, suggesting minimal herding behavior. This phase indicates that a smaller influx of newcomers does not significantly drive collective behavior, and market participants tend to act more independently. As the percentage of newcomers increases to the range of 0.35, γ_2 becomes more substantially negative, reflecting a pronounced increase in herding behavior. In this intermediate phase, the actions of newcomers start to align more closely, leading to a stronger collective movement in market behavior. This alignment is indicative of newcomers following similar trends and signals, amplifying the herding effect.

However, as the share of newcomers continues to rise beyond 0.55, γ_2 begins to increase, although it remains negative. This trend might mark the onset of a saturation effect (Bikhchandani et al., 2021; Cong et al., 2021). At this point, the market might experience collective uncertainty, as the proportion of inexperienced participants could become too large. This might lead to a shift towards more individualistic behaviour. Such a shift would be characterized by a less pronounced negative γ_2 , indicating reduced herding.

The relationship between γ_2 and BGCI return in Figure 6b presents a non-linear and fluctuating path. Our findings reveal a counterintuitive result. Contrary to expectations that higher returns on the BGCI would amplify herding behaviour in the NFT market, we observe that stronger performance in the cryptocurrency market actually attenuates the prevalence of herding tendencies within the NFT landscape.

According to Demirer and Kutun (2006) and Huo et al. (2023), assets that are the focus of institutional herding initially experience positive abnormal returns. In other words, when institutional investors move en masse to buy a particular asset, its price surges above what would be its intrinsic value under normal market conditions. However, this is a transitory phenomenon; prices tend to revert to their mean or intrinsic value in the long run, nullifying the temporary gains accrued due to herding.

Corbet et al. (2022) elaborate on this by offering a theoretical framework. Their model suggests that the inflation of asset prices is not a random or inexplicable event but a direct




FIGURE 6 γ_2 Dynamics, newcomers and Bloomberg Galaxy Crypto Index (BGCi) return. This graph presents the relationships between the γ_2 coefficient, the percentage of newcomer investors in the nonfungible token (NFT) market and the BGCi returns. The γ_2 coefficient represents the sensitivity of the cross-sectional asset deviation to the squared market returns, a proxy for herding behaviour. The functions in the graphs have been smoothed with a moving average with a window of 30 to reduce noise and highlight underlying trends. Initially, when the share of newcomers is lower than 0.30, γ_2 touches 0, suggesting minimal herding behaviour. This phase indicates that a smaller influx of newcomers does not significantly drive collective behaviour, and market participants tend to act more independently. As the percentage of newcomers increases to the range of 0.35, γ_2 becomes more substantially negative, reflecting a pronounced increase in herding behaviour. In this intermediate phase, the actions of newcomers start to align more closely, amplifying the herding effect. However, as the share of newcomers rises beyond 0.55, γ_2 begins to increase, although it remains negative. This trend marks the onset of a saturation effect. At this point, the market might experience collective uncertainty, leading to a shift towards more individualistic behaviour, characterized by a less pronounced negative γ_2 , indicating reduced herding. The relationship between γ_2 and BGCi return is nonlinear and fluctuating. Interestingly, instead of increasing herding behaviour, a strong performance in the cryptocurrency market seems to reduce it in the art NFT market. Typically, assets subject to institutional herding experience initial return spikes, but these gains are short lived as prices correct to intrinsic values. Herding inflates prices temporarily, which then revert over time. A strong BGCi, indicating a healthy cryptocurrency market, leads investors to shift capital from NFTs to cryptocurrencies, thereby reducing herding in the art NFT market. Thus, strong BGCi performance stabilizes investor behaviour and reduces herding-induced distortions in art NFTs. (a) Newcomer proportion and γ_2 and (b) BGCi return and γ_2 . [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

outcome of herding behaviour among institutional investors. Once the herd dissipates or shifts focus, the inflated prices naturally correct themselves, aligning more closely with their true, intrinsic value. This implies a sort of pull effect that intrinsic value exerts on asset prices, drawing them back to equilibrium once the distorting influence of herding has waned.

We consider in the following how these theories interface with our observations regarding the BGCi and the art NFT market. Figure 6b illustrates that the relationship between BGCi return and γ_2 is characterized by significant fluctuations rather than a straightforward trend. This non-linear relationship suggests that a strong performance in the BGCi may moderate

herding tendencies in the NFT market but this effect is complex and variable. In practice, when the BGCI is performing well, it signifies strength in the broader cryptocurrency market. As herding in the NFT market ceases and leads to negative returns, investors in the NFT market may seek alternative avenues for investment, often transitioning their capital into more stable or promising assets, like, cryptocurrencies. This inflow of capital positively impacts the BGCI. A strong BGCI performance serves as a signal or even a magnet for capital reallocation away from the NFT market, thereby moderating herding behaviour in NFTs. This mechanism acts as a balancing force across the two correlated but distinct asset classes, art NFTs and cryptocurrencies represented by the BGCI. This balancing act ensures a more stable market environment, reducing the distortions typically caused by herding and allowing for a more rational asset pricing across both markets. Hence, the performance of the BGCI serves not only as a reflection of market sentiment but also as a moderating variable that indirectly influences investor behaviour in the NFT market.

5 | CONCLUSIONS

We have analyzed the value determinants and herding behaviour in the art NFT market, establishing a ranking system for artists and examining the factors that drive NFT prices. Our analysis reveals a nuanced picture of the NFT ecosystem. We find a positive relationship between the number of trades and floor prices with the average price of NFTs, underscoring the influence of market activity and perceived value on pricing. In contrast to prevailing assumptions, our research indicates that social media activity may inversely affect NFT prices, potentially due to market oversaturation and reduced perceived rarity. This challenges the conventional wisdom that increased visibility invariably boosts demand.

We observe a pronounced herding behaviour, particularly among top and mid-tier artists, manifesting in a market heavily skewed towards a select few, creating pronounced inequalities. This finding is critical as it highlights the concentration of market influence and wealth within a small segment of artists, reflecting a broader trend of inequality that transcends the art world into the digital space. Additionally, herding behaviour is influenced by the entrance of new investors and fluctuations in the cryptocurrency market. While newcomers tend to reinforce herding, an increase in their numbers could lead to a saturation effect, ultimately introducing more volatility. Conversely, a strong cryptocurrency market appears to moderate herding within the NFT market, likely due to a reallocation of capital towards more “stable” crypto investments. Furthermore, our study provides insights into the complex dynamics of consumer choices, showing that higher artist rankings do not guarantee continued future demand. This suggests a market that values novelty and the discovery of new talents over established reputations.

Our findings validate a prevalent belief in the crypto space as a whole, that due to the unregulated nature of the industry, cryptocurrency-related endeavours attract swindlers that do more harm than good both in terms of reputation and finances. However, this does not imply that art NFTs have no artistic worth. As speculative activities diminish and the market stabilizes, there is an opportunity for the art NFT sector to foster a more creative and sustainable environment that benefits both artists and collectors. While the art NFT market is subject to the whims of speculation and herding, it also appears to possess the potential for innovation and value creation. The key lies in understanding and navigating the market's complex dynamics,

from social media influences to broader economic factors, to harness the potential of NFTs as a novel artistic medium.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

DATA AVAILABILITY STATEMENT

The authors elect not to share the data supporting the findings of this study.

ORCID

Gilbert Fridgen  <http://orcid.org/0000-0001-7037-4807>

Roman Kräussl  <http://orcid.org/0000-0001-8933-9278>

Orestis Papageorgiou  <http://orcid.org/0000-0002-2245-5079>

Alessandro Tugnetti  <http://orcid.org/0000-0003-3811-2828>

REFERENCES

- Adams, R. B., Kräussl, R., Navone, M., & Verwijmeren, P. (2021). Gendered prices. *The Review of Financial Studies*, 34(8), 3789–3839.
- Alizadeh, S., Setayesh, A., Mohamadpour, A., & Bahrak, B. (2023). A network analysis of the non-fungible token (NFT) market: Structural characteristics, evolution, and interactions. *Applied Network Science*, 8(1), 1–18.
- Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid–ask spread. *Journal of Financial Economics*, 17(2), 223–249.
- Ang, A., Liu, J., & Schwarz, K. (2020). Using stocks or portfolios in tests of factor models. *Journal of Financial and Quantitative Analysis*, 55(3), 709–750.
- Ante, L. (2022). The non-fungible token (NFT) market and its relationship with Bitcoin and Ethereum. *FinTech*, 1(3), 216–224.
- Aubry, M., Kräussl, R., Manso, G., & Spaenjers, C. (2023). Biased auctioneers. *The Journal of Finance*, 78(2), 795–833.
- Azarmi, T. F., & Menny, P. R. P. (2013). Star artists and herding in fine arts' market: Theory and empirical evidence. Available at SSRN: <https://ssrn.com/abstract=2232307>
- Babić Rosario, A., de Valck, K., & Sotgiu, F. (2020). Conceptualizing the electronic word-of-mouth process: What we know and need to know about eWOM creation, exposure, and evaluation. *Journal of the Academy of Marketing Science*, 48(3), 422–448.
- Bailey, M. J., Muth, R. F., & Nourse, H. O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58(304), 933–942.
- Bamakan, S. M. H., Nezhadsistani, N., Bodaghi, O., & Qu, Q. (2022). Patents and intellectual property assets as non-fungible tokens; Key technologies and challenges. *Scientific Reports*, 12(1), 2178.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.
- Bao, T., Ma, M., & Wen, Y. (2023). Herding in the non-fungible token (NFT) market. *Journal of Behavioral and Experimental Finance*, 39, 100837.
- Bao, H., & Roubaud, D. (2022). Non-fungible token: A systematic review and research agenda. *Journal of Risk and Financial Management*, 15(5), 215.
- Bianchi, D., & Babiak, M. (2022). On the performance of cryptocurrency funds. *Journal of Banking & Finance*, 138, 106467.
- Bikhchandani, S., Hirshleifer, D., Tamuz, O., & Welch, I. (2021). *Information cascades and social learning* [NBER Working Paper, ID 28887].
- Borri, N., Liu, Y., & Tsyvinski, A. (2023). The economics of non-fungible tokens. Available at SSRN: <https://ssrn.com/abstract=4052045>
- Campbell, J. Y., & Shiller, R. J. (1988). The dividend–price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3), 195–228.

- Caporale, G. M., Kang, W.-Y., Spagnolo, F., & Spagnolo, N. (2021). Cyber-attacks, spillovers and contagion in the cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money*, 74, 101298.
- Case, K. E., & Shiller, R. J. (1989). The efficiency of the market for single-family homes. *The American Economic Review*, 79(1), 125–137.
- Chainalysis. (2022). *The Chainalysis 2021 NFT market report*.
- Chanel, O., Chanel, O., Gérard-Varet, L.-A., & Ginsburgh, V. (1996). The relevance of hedonic price indices: The case of paintings. *Journal of Cultural Economics*, 20(1), 1–24.
- Chang, E. C., Cheng, J. W., & Khorana, A. (2000). An examination of herd behavior in equity markets: An international perspective. *Journal of Banking & Finance*, 24(10), 1651–1679.
- Cheah, E.-T., & Fry, J. (2015). Speculative bubbles in bitcoin markets? An empirical investigation in the fundamental value of Bitcoin. *Economics Letters*, 130, 32–36.
- Collins. (2021). *Get your crypto at the ready: NFTs are big in 2021*. Collins Dictionary Language Blog.
- Cong, L. W., Landsman, W., Maydew, E., & Rabetti, D. (2022). *Tax-loss harvesting with cryptocurrencies* [NBER Working Paper, ID 30716].
- Cong, L. W., Li, Y., & Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3), 1105–1155.
- Corbet, S., Hou, Y. G., Hu, Y., Larkin, C., Lucey, B., & Oxley, L. (2022). Cryptocurrency liquidity and volatility interrelationships during the COVID-19 pandemic. *Finance Research Letters*, 45, 102137.
- Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10), 1407–1424.
- Demirer, R., & Kutan, A. (2006). Does herding behavior exist in Chinese stock markets? *Journal of International Financial Markets, Institutions and Money*, 16(2), 123–142.
- Diebold, F. X., & Rudebusch, G. D. (1999). *Business cycles: Durations, dynamics, and forecasting*. Princeton University Press.
- Diewert, W. E. (1976). Exact and superlative index numbers. *Journal of Econometrics*, 4(2), 115–145.
- Dowling, M. (2022a). Fertile LAND: Pricing non-fungible tokens. *Finance Research Letters*, 44, 102096.
- Dowling, M. (2022b). Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters*, 44, 102097.
- Drehmann, M., Oechssler, J., & Roeder, A. (2005). Herding and contrarian behavior in financial markets: An internet experiment. *The American Economic Review*, 95(5), 1403–1426.
- Economou, F., Kostakis, A., & Philippas, N. (2011). Cross-country effects in herding behaviour: Evidence from four South European markets. *Journal of International Financial Markets, Institutions and Money*, 21(3), 443–460.
- Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.
- Financial Times. (2021a). *How digital art dominated 2021*.
- Financial Times. (2021b). *How NFTs became a \$40bn market in 2021*.
- Fu, J., & Wu, L. (2021). Regime-switching herd behavior: Novel evidence from the Chinese A-share market. *Finance Research Letters*, 39, 101652.
- Hartwich, E., Ollig, P., Fridgen, G., & Rieger, A. (2023). Probably something: A multi-layer taxonomy of non-fungible tokens. *Internet Research*, 34(1), 216–238.
- Hategan, V. (2023). *Dead NFTs: The evolving landscape of the NFT market*. dappGambler.
- Hemenway Falk, B., Tsoukalas, G., & Zhang, N. (2022). Economics of NFTs: The value of creator royalties. Available at SSRN: <https://ssrn.com/abstract=4284776>
- Holden, R., & Malani, A. (2022). An examination of velocity and initial coin offerings. *Management Science*, 68(12), 8515–9218.
- Horky, F., Rachel, C., & Fidrmuc, J. (2022). Price determinants of non-fungible tokens in the digital art market. *Finance Research Letters*, 48, 103007.
- Huang, D., & Goetzmann, W. N. (2023). *Selection-neglect in the NFT bubble* [NBER Working Paper, ID 31498].
- Huo, X., Liu, X., & Zheng, W. (2023). The term structure of mutual fund herding. *European Financial Management*, 29(3), 901–929.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.

- Kapoor, A., Guhathakurta, D., Mathur, M., Yadav, R., Gupta, M., & Kumaraguru, P. (2022). TweetBoost: Influence of social media on NFT valuation. *Companion Proceedings of the Web Conference 2022*, 22, 621–629.
- Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3), 967–998.
- Kong, D.-R., & Lin, T.-C. (2021). Alternative investments in the fintech era: The risk and return of non-fungible token (NFT). Available at SSRN: <https://ssrn.com/abstract=3914085>
- Korteweg, A., Kräussl, R., & Verwijmeren, P. (2016). Does it pay to invest in art? A selection-corrected returns perspective. *The Review of Financial Studies*, 29(4), 1007–1038.
- Kuminoff, N., Parmeter, C., & Pope, J. (2010). Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities? *Journal of Environmental Economics and Management*, 60(3), 145–160.
- Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of Political Economy*, 74(2), 132–157.
- Li, S., & Chen, Y. (2023). How non-fungible tokens empower business model innovation. *Business Horizons*, 66(4), 543–544.
- Mandel, B. R. (2009). Art as an investment and conspicuous consumption good. *American Economic Review*, 99(4), 1653–1663.
- Mazur, M. (2021). Non-fungible tokens (NFT). The analysis of risk and return. Available at SSRN: <https://ssrn.com/abstract=3953535>
- McAndrew, C. (2022). *Collectors are spending more, not less, on NFT art, according to the art market report*. Art Basel and UBS.
- Mei, J., & Moses, M. (2002). Art as an investment and the underperformance of masterpieces. *The American Economic Review*, 92(5), 1656–1668.
- Mekacher, A., Bracci, A., Nadini, M., Martino, M., Alessandretti, L., Aiello, L. M., & Baronchelli, A. (2022). Heterogeneous rarity patterns drive price dynamics in NFT collections. *Scientific Reports*, 12(1), 13890.
- Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5), 1089–1122.
- Muehlenbachs, L., Spiller, E., & Timmins, C. (2015). The housing market impacts of shale gas development. *American Economic Review*, 105(12), 3633–3659.
- Nadini, M., Alessandretti, L., Di Giacinto, F., Martino, M., Aiello, L. M., & Baronchelli, A. (2021). Mapping the NFT revolution: Market trends, trade networks, and visual features. *Scientific Reports*, 11(1), 20902.
- O'Brien, R. M. (2007). A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity*, 41(5), 673–690.
- Ofek, E., & Richardson, M. (2003). DotCom Mania: The rise and fall of internet stock prices. *The Journal of Finance*, 58(3), 1113–1137.
- Park, A., & Sabourian, H. (2011). Herding and contrarian behavior in financial markets. *Econometrica*, 79(4), 973–1026.
- Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science*, 51(10), 1494–1504.
- Renneboog, L., & Spaenjers, C. (2013). Buying beauty: On prices and returns in the art market. *Management Science*, 59(1), 36–53.
- Rojas-Lamoren, Á.J., Del Barrio-García, S., & Alcántara-Pilar, J. M. (2022). A review of three decades of academic research on brand equity: A bibliometric approach using co-word analysis and bibliographic coupling. *Journal of Business Research*, 139, 1067–1083.
- Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34–55.
- Rosen, S. (1981). The economics of superstars. *The American Economic Review*, 71(5), 845–858.
- Schaar, L., & Kampakis, S. (2022). Non-fungible tokens as an alternative investment: Evidence from Crypto-Punks. *The Journal of The British Blockchain Association*, 5(1), 1–12.
- Shiller, R. (1995). Conversation, information, and herd behavior. *American Economic Review*, 85(2), 181–185.
- So, E. C., & Wang, S. (2014). News-driven return reversals: Liquidity provision ahead of earnings announcements. *Journal of Financial Economics*, 114(1), 20–35.
- Sockin, M., & Xiong, W. (2022). *Decentralization through tokenization* [NBER Working Paper, ID 29720].
- Sotheby's. (2021). *NFTs: Redefining digital ownership and scarcity*.
- Stiglitz, J. E. (1989). Markets, market failures, and development. *The American Economic Review*, 79(2), 197–203.

- Taleb, N. N. (2021). Bitcoin, currencies, and fragility. *Quantitative Finance*, 21(8), 1249–1255.
- Triplett, J. (2004). *Handbook on hedonic indexes and quality adjustments in price indexes: Special application to information technology products* [OECD Working Papers Series, 9].
- Urquhart, A. (2016). The inefficiency of Bitcoin. *Economics Letters*, 148, 80–82.
- Valeonti, F., Bikakis, A., Terras, M., Speed, C., Hudson-Smith, A., & Chalkias, K. (2021). Crypto collectibles, museum funding and OpenGLAM: Challenges, opportunities and the potential of non-fungible tokens (NFTs). *Applied Sciences*, 11(21), 9931.
- Vayanos, D., & Wang, J. (2012). Liquidity and asset returns under asymmetric information and imperfect competition. *The Review of Financial Studies*, 25(5), 1339–1365.
- Wang, Y., Horky, F., Baals, L. J., Lucey, B. M., & Vigne, S. A. (2022). Bubbles all the way down? Detecting and date-stamping bubble behaviours in NFT and DeFi markets. *Journal of Chinese Economic and Business Studies*, 20(4), 415–436.
- Wilson, K. B., Karg, A., & Ghaderi, H. (2021). Prospecting non-fungible tokens in the digital economy: Stakeholders and ecosystem, risk and opportunity. *Business Horizons*, 65(5), 657–670.
- Yousaf, I., & Yarovaya, L. (2022). Herding behavior in conventional cryptocurrency market, non-fungible tokens, and DeFi assets. *Finance Research Letters*, 50, 103299.

How to cite this article: Fridgen, G., Kräussl, R., Papageorgiou, O., & Tugnetti, A. (2024). Pricing dynamics and herding behaviour of NFTs. *European Financial Management*, 1–41. <https://doi.org/10.1111/eufm.12506>

APPENDIX A: DATA EXTRACTION, CLEANING AND PREPARATION

For the on-chain part of our analysis, we use Dune, Etherscan and the Application Programming Interfaces (APIs) of Alchemy and OpenSea, while for the off-chain portion, we employ the price oracle of Chainlink and the APIs of Discord and Twitter.

Using Dune, we extract every on-chain transaction involving ERC-721 and ERC-1155 tokens on six of Ethereum's largest NFT marketplaces, specifically OpenSea, Rarible, Nifty, LooksRare, Foundation and SuperRare. The extraction period spans from the establishment of the ERC-721 standard in January 2018–February 2022.² From the retrieved transactions, we identify 14,580 collections that have sold at least one NFT. Using Etherscan and OpenSea, we identify the deployers and the artists behind each collection and gather all the available metrics on a collection level, such as floor price, the total number of owners, collection name, number of transfers, the total volume traded in ETH and the artists' characteristics, like, gender and the total number of artists.

In a second step, we exclude the collections whose NFTs were involved in fewer than 30 sales to establish statistical significance within collections. Following the categorization criteria of <https://Nonfungible.com>, we divide the collections into six categories: art, collectibles, game, metaverse, utility and other. Since our research focuses on art NFTs, we omit collectibles as they are comprised of a single artwork with minor variations per NFT. We also exclude NFTs associated with games and metaverses since their price is also influenced by the in-game or metaverse mechanics. Further, we remove utility NFTs such as liquidity provider positions in Uniswap V3 and other NFTs that do not fit in any of the preceding categories, like, music NFTs or those with unique features.

²The data extraction stopped in February 2022 due to the Russo-Ukrainian war, which introduced substantial levels of systematic volatility, making it impractical to examine the volatility of an individual asset class.

TABLE A1 Variable descriptions.

This table reports all the variables extracted from on and off blockchain sources alongside their short descriptions.

Variable	Description
<i>average price</i>	Average nonfungible token (NFT) price
<i>BGCI return</i>	Daily Bloomberg Galaxy Crypto Index last price in USD
<i>black</i>	Percentage of black color in the NFT image
<i>blue</i>	Percentage of blue color in the NFT image
<i>buyer</i>	ETH address of the NFT buyer
<i>buyer seller pair</i>	Unique buyer–seller combined ETH addresses
<i>collection name</i>	Name of the NFT collection
<i>CSAD</i>	Cross-sectional absolute deviation
<i>cum score</i>	Artist cumulative score
<i>cum score year rank</i>	Artist yearly cumulative score
<i>demand</i>	Demand of NFTs by each artist category, as ratio between <i>share</i> and its geometric mean
<i>deployer creator generalities</i>	Deployer type: Single artist (male, female), artists collaboration, company
<i>deployer creator name</i>	Name of the deployer
<i>discord account</i>	Presence of a Discord account representing the artist or the NFT collection
<i>discord server</i>	Name of the artist Discord account
<i>floor price</i>	NFT floor price
<i>followers count</i>	Number of followers per Twitter account
<i>following count</i>	Number of following per Twitter account
<i>gray</i>	Percentage of grey color in the NFT image
<i>green</i>	Percentage of green color in the NFT image
<i>last price</i>	NFT last selling price
<i>listed count</i>	Average number of public lists memberships per Twitter account
<i>market return</i>	Overall mean return across NFT collections
<i>marketplace collection</i>	NFT belonging to one between: Foundation (FND), Editorial, KnownOrigin, SuperRare
<i>max price</i>	Maximum NFT price
<i>newcomer proportion</i>	Proportion of newcomers in the NFT market
<i>nft</i>	Single NFT's ETH address
<i>nft type</i>	Belonging to the ERC721 or ERC1155 protocol per single NFT
<i>norm shannon entropy</i>	Normalized Shannon Entropy per NFT image
<i>num of colors</i>	Number of colors in the NFT image
<i>num of owners</i>	Number of NFT's owners overtime

(Continues)

TABLE A1 (Continued)

Variable	Description
<i>num of trades</i>	Number of trades per NFT
<i>orange</i>	Percentage of orange color in the NFT image
<i>platform of last sale</i>	Blockchain platform of last sale per NFT
<i>purple</i>	Percentage of purple color in the NFT image
<i>red</i>	Percentage of red color in the NFT image
<i>quote count month</i>	Average number of “Retweets” with comments per Twitter account per month
<i>score</i>	Artist score
<i>seller</i>	ETH address of the NFT seller
<i>share</i>	Artist market share, computed as ratio between artist specific <i>volume</i> and the overall market <i>volume</i>
<i>sold ratio</i>	Percentage of NFT sold per artist’s collection
<i>timediff</i>	Time difference (in days) between the collection creation and the selling dates per NFT
<i>twitter account</i>	Presence of a Twitter account representing the artist or the NFT collection
<i>twitter handle</i>	Name of the artist Twitter account
<i>usd amount</i>	NFT transaction specific USD amount
<i>volume</i>	Artist’s market capitalization
<i>verified</i>	Presence of “verified” badge per Twitter account
<i>white</i>	Percentage of white color in the NFT image
<i>year collection creation</i>	Year of creation per NFT collection
<i>yellow</i>	Percentage of yellow color in the NFT image

On the basis of this data set, we continue with the off-chain data extraction on Twitter and Discord. After manually identifying the Discord Servers and Twitter handles of the NFT collections (or in the case where these were missing their artists), we retrieve all publicly available metrics such as the number of followers, retweets, likes, replies, number of messages as well as discord messages and users that were involved in the general and announcement channels. At the end of the data extraction and cleaning on a collection level, our data set comprises 531 art NFT collections containing a total of 1,460,718 NFTs.

Using the Alchemy and OpenSea APIs, we retrieve all on-chain available information on these NFTs.³ Consequently, we further exclude the NFTs whose data are missing, those that are in video or animated Scalable Vector Graphics (SVG) formats and those with wrong data formats. In a final step, we download every file associated with each NFT and converted SVG

³We opted for the use of both APIs because we identified that each has missing data. By utilizing both, we were able to cover most of the gaps.

images to Portable Network Graphics (PNG) to allow for easier image analysis. Our final sample contains 2.15 terabytes (TB) of PNG, Joint Photographic Experts Group and Graphics Interchange Format (GIF) files of 875,389 NFTs.

From the extracted files, we determine their colour proportions and the total number of colours present by utilizing the Hue, Saturation, Value (HSV) colour model.⁴ We group every colour around nine commonly used colours: black, white, grey, red, green, blue, yellow, purple and orange. In addition, we assess the complexity of the files using Shannon's entropy and the block decomposition method, with the use of the coding theorem method, because it is regarded as one of the best approaches both from a statistical and an algorithmic point of view. In the case of GIF files, we average the values of each frame.

So-called wash trading relates to a form of market manipulation where a small number of investors repeatedly buy and sell the same asset, generating an inorganic market activity. It creates artificial trading volume and gives the appearance that the asset is more in demand than it actually is. Wash trading can greatly affect the price, the traded volume and the selling frequency. We exclude in our analysis all individual NFTs that have been involved in wash trading and all collections whose trading volume is greater than 90% due to wash trading, such as Terraforms by Mathcastles. Additionally, we eliminate NFTs whose first sale took place after 01-01-2022, since more than 75% of the total volume in January 2022 and 55% of the total volume in 2022 was a result of wash trading and identifying the genuine transactions would be complicated and risky.

Using Dune, we determine that 875,389 of the aforementioned art NFTs were involved in 385,884 sales. On the basis of these sales, we generate two data sets, one containing data pertaining to every sale of the extracted art NFTs (*Transaction data set* in Table 1), and the other comprising aggregated data per NFT, which was augmented with data from Discord and Twitter, and colour and image complexity data (*Aggregate data set* in Table 1). In both data sets, the sale price is measured in USD, for which we utilize the exchange rates provided by Chainlink at the time of the sale for the individual NFT sales and at the end of the day on 10 February 2022 for the collection floor price.

⁴HSV is a mathematical abstraction that describes how colours can be represented by electronic systems. In this model, colours are represented by a cylindrical geometry whose angle represents the hue, the x-axis determines saturation and the value depends on the y-axis.

3.5. Research Paper 5: *Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch*



Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch

JOAQUIN DELGADO FERNANDEZ, Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg, Luxembourg

TOM BARBEREAU, Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg, Luxembourg

ORESTIS PAPAGEORGIOU, Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg, Luxembourg

With advancements in distributed ledger technologies and smart contracts, tokenized voting rights gained prominence within decentralized finance (DeFi). Voting rights tokens (a.k.a. governance tokens) are fungible tokens that grant individual holders the right to vote upon the fate of a project. The motivation behind these tokens is to achieve decentral control within a decentralized autonomous organization (DAO). Because the initial allocations of these tokens is often undemocratic, the DeFi project and DAO of Yearn Finance experimented with a fair launch allocation where no tokens are pre-mined and all participants have an equal opportunity to receive them. Regardless, research on voting rights tokens highlights the formation of timocracies over time. The consideration is that the tokens' tradability is the cause of concentration. To examine this proposition, this article uses an agent-based model to simulate and analyze the concentration of voting rights tokens post three fair launch allocation scenarios under different trading modalities. The results show that regardless of the allocation, concentration persistently occurs. It confirms the consideration that the 'disease' is endogenous: the cause of concentration is the tokens' tradability. The findings inform theoretical understandings and practical implications for on-chain governance mediated by tokens.

CCS Concepts: • **Applied computing** → IT governance; • **Social and professional topics** → *Centralization/decentralization*; • **Computing methodologies** → **Agent/discrete models**;

Additional Key Words and Phrases: Blockchain, governance, decentralized finance, voting rights tokens, fair launch, agent-based modeling

ACM Reference Format:

Joaquin Delgado Fernandez, Tom Barbereau, and Orestis Papageorgiou. 2024. Agent-Based Model of Initial Token Allocations: Simulating Distributions post Fair Launch. *ACM Trans. Manag. Inform. Syst.* 15, 1, Article 5 (March 2024), 22 pages. <https://doi.org/10.1145/3649318>

This research was funded by the Luxembourg National Research Fund (FNR) and PayPal PEARL (grant reference 13342933), as well as by the FNR in the FiReSPArX (grant reference 14783405) and PABLO (grant reference 16326754) projects. Additionally, our research was supported by the European Union (EU) within its Horizon 2020 program, project MDOT (grant reference 814654). For the purpose of open access, the authors have applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any author-accepted manuscript version arising from this submission. PayPal's financial support was administered via the FNR, and by contractual agreement, PayPal has no involvement in the authors' research. Authors' address: J. Delgado Fernandez (Corresponding author), T. Barbereau, and O. Papageorgiou, Interdisciplinary Centre for Security, Reliability and Trust (SnT), 29 Av. John F. Kennedy, University of Luxembourg, Luxembourg, Luxembourg; e-mails: joaquin.delgadofernandez@uni.lu, tom.barbereau@uni.lu, orestis.papageorgiou@uni.lu.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2024 Copyright held by the owner/author(s).

ACM 2158-656X/2024/03-ART5

<https://doi.org/10.1145/3649318>

1 INTRODUCTION

The digital representation of value and ownership, in the form of fungible and non-fungible tokens respectively, provides the basis for the *token economy*. Unlike traditional economies, the token economy does not rely on trusted third parties to verify transactions [61, 64]; instead, **Distributed Ledger Technology (DLT)** and smart contracts ensure integrity in a pseudonymous peer-to-peer network of interactions [7, 73]. Advancements in the token economy with a focus on financial services and products materialized under the heading of **Decentralized Finance (DeFi)** [3, 6, 59, 71].

DLT enables people to coordinate themselves *on-chain*—that is, transactions and interactions are “mediated by a set of self-executing rules [i.e., smart contracts] deployed on a public blockchain” independently from central control [36, p. 1]. To achieve decentral control in DeFi, developers created and allocated so-called *voting rights tokens*—(fungible) tokens that stipulate voting entitlements to vote upon change proposals to a project [47]. An example is the Decentralized Exchange (DEX) Uniswap, a project which uses smart contracts to automate the exchange of fungible tokens of the Ethereum protocol. Its voting rights token, UNI, allows holders to cast votes and decide upon the use of resources stored in a treasury (defined in a smart contract). The participants and holders of voting rights tokens are typically referred to as members of a **Decentralized Autonomous Organization (DAO)**.

The distribution of voting rights tokens, however, is considered as controversial given that the initial allocation often favors a minority of insiders (developers, investors, etc.). Tokens allocated to insiders are common within initial coin offerings [14], hence a differentiation between initial allocations that favor insiders (labeled, “private”) or those that do not (labeled, “public”) [28, p. 10]. History repeats itself as insider allocations of voting rights tokens are common in DeFi DAOs [6].

One outlier to insider allocations is Yearn Finance. Its core developer, Andre Cronje, denoted that at inception voting rights tokens are majoritarially allocated to “friends and family” [20]—impossibly leading to decentral control. As solution, he implemented a type of initial token allocation for the voting rights tokens of the Yearn Finance (YFI) DAO—the *fair launch*—whereby all community members have an equal opportunity to receive a portion of the initial supply [63]. Although in theory this allocation strategy achieves equity through principles of fairness (see [54]), reality looks dire in the long run: Barbereau et al. [5] demonstrate how, as with most other voting rights tokens, for YFI concentration of wealth and voting power persists. Because holders rarely use these to cast votes, Barbereau et al. [6] denote a common theme and propose a purposefully *descriptive* theory of voting rights tokens as justification for concentration: they are *tradable* assets on cryptocurrency markets. This description may not seem surprising against the consideration that wealth in the token economy is concentrated (c.f. concentration in Bitcoin and Ethereum [32]), and so are capital markets more broadly (see [52]). Indeed, the common feature of tradability appears to justify, on an intuitive level, the expectations that “wealth trickles up in free-market economies” [12].

The hypothesis whether the experiment of Andre Cronje’s fair launch was inherently doomed to fail, given that the underlying tokens are *tradable*, remains unaccounted for. Taking the principle of a ‘fair’ launch (n.b., all tokens are allocated ‘fairly’), one can consider whether alternative allocations are successful in achieving decentral control in the long run. Correspondingly, to challenge this consideration, this article addresses the following two research questions:

RQ1: Does trading behavior affect voting rights token distributions over time?

RQ2: Do alternative, ‘fair launch’ token allocations affect voting rights token distributions over time?

Provided the context in which these novel governance structures are deployed, our research topic—‘fair launch’ token allocations—is of importance to **Information System (IS)** and requires the adoption of multidisciplinary perspectives [66]. Hence, we lean on previous theory on

governance of public protocols/DAOs and token design, and use quantitative methods rooted in agent-based simulation—that is, **Agent-Based Modeling (ABM)**. Guided by an ambition “for discovery and explanation” [8, p. 516], we specifically adopt ABM to simulate the trade and eventual distribution of voting rights tokens post three distinct ‘fair launch’ allocations (n.b., scenarios denoted S_n). The developed model is going from “real world to simulation world” [8, p. 516], an approach that is particularly suitable for the exploration of understudied, novel phenomena—here, the fair launch allocation. Within IS, the utility of ABM for the study of phenomena with “nonlinear behavior” [35, p. 158] is well recognized. At large, the discipline is receptive of contributions emerging from simulation research [8, 21, 23, 72].

Davis et al. [21, p. 482] advise to ground the model within “simple theory”—theory that provides the “basic concepts and process that describe a phenomenon” [8, p. 506]. Here, we focus on governance of public-permissionless systems in general, and DeFi DAOs in particular (Section 2). To establish further “epistemic credibility in the simulation model” [8, p. 517], aside from theory (deductive approach) we use empirical data (inductive approach) from Yearn Finance (Section 3). The development of our model (Section 4) is informed by the artificial cryptocurrency markets designed in the work of Cocco et al. [19] and Roşu and Saleh [57]. Therein, agents represent traders that are endowed with an amount of fiat currency and seek to acquire the (artificially created) voting rights tokens (TKNs). The market rules are loosely based on understandings of clearing houses [41]. To investigate RQ2, we developed two alternative scenarios to Andre Cronje’s fair launch scenario (S_0) which respectively consider fairness in egalitarian terms (S_1) and ‘at random’ (S_2). The principles underlying these allocations build on political philosophy [43, 53]. We measure concentration in terms of the Gini coefficient [31] and Shannon entropy [62]. With our work, we make the following three contributions:

- We provide an agent-based model for the analysis of token distributions under various market conditions reflective of trading.
- Our simulation results show how over time, regardless of the ‘fair’ initial token allocation, concentration is imminent.
- We extend understandings on the governance of DAOs and tokenomics to formerly include initial token allocations as part of governance parameters.

2 RELATED WORK

2.1 Governance in Public-Permissionless Systems

Bitcoin [46] led to a burgeoning movement of developers who saw decentralization beyond technical terms, not least as an ambiguous mix of political, economic, and organizational ideals [60] that materialized in the emergence of alternative protocols. Commonly, the next generation of DLT includes *smart contracts* and the possibility to deploy tokens [7, 73]. Collectively, these technologies enable new, decentral business opportunities, modes of organization, and governance frameworks [61, 64].

While in public-permissionless systems ‘decentralization’ in technical terms is achieved, in political terms it is often contentious [11, 60]. Over the decision to increase the size of Bitcoin blocks, community debates led to an outright “civil war” [22, p. 8] that pitted parties against each other over socio-political motives. The core developers, who act as gatekeepers to protocol changes and are *de facto* in charge of governance, eventually took an autocratic decision against that increase—a behavior pejoratively described as “senatorial” [49]. Ethereum, too, faced its share of controversy. Following an exploited smart contract bug, the Ethereum Foundation’s leaders decided to irreversibly fork the ledger [26]. Against this backdrop, Penzo and Selvadurai [51, p. 19] denote how

Table 1. Excerpt of the Token Classification Proposed in the Work of Oliveira et al. [47]

Governance Parameters	Representation	Digital		Physical		Legal	
	Supply	Schedule based	Pre-mined, scheduled distribution	Pre-mined, one-off distribution		Discretionary	
	Incentive system	Enter Platform	Use Platform	Stay Long-Term		Leave Platform	

in public-permissionless systems, governing communities resort to informal adjudications, typically “immune . . . from state scrutiny.”

Consequently, scholars distinguish between *on-chain* and *off-chain* governance of DLT systems. The former refers to rules that enforce the ‘code-is-law’ dictum—in other words, using smart contracts (and tokens) to define governance mechanisms and structures (“now the code runs itself”) [55]. The greatest degree of on-chain governance is typically achieved within and as part of DAOs (see the definition proposed by Hassan and de Filippi [36]). The informal resolution mechanisms in Bitcoin and Ethereum, however, are examples that demonstrated the shortcomings of the dictum and shed light on ulterior powers and politics that are at play [22, 26]. Off-chain governance refers to the formalization of control via the intermediary of endogenous (through the foundation of institutions like consortia, cooperatives, etc.) or exogenous (national laws, regulations, standards, etc.) structures [55, 74] that are typically registered and held liable [4].

2.2 Voting Rights Token-Based Governance

Within DeFi, beyond improvements made to the financial value chain [3, 59], experiments were made at implementing governance structures fully on-chain—most notably by embedding voting rights into tokens. These tokens grant holders the ability to cast votes on proposals. While the features of these tokens are contextual to the individual project, the majority of these follow the fungible token standard ERC-20. Like most cryptocurrencies, they are tradable on regular and decentralized exchanges [5]. By nature, the study of these tokens is at the intersection of research of blockchain governance and *tokenomics*—subdomains of cross-disciplinary research on DLT—and contribute to research on DAOs (see [6, 9, 34, 36, 58]).

Oliveira et al. [47, p.8] define *Governance Parameters* of tokens as those parameters that “relate to what [it] effectively represents and how this connects to the way the platform is governed and managed.” The authors introduce three parameters (Table 1): (1) “Representation” (the type of asset represented by a token), (2) “Supply” (the way tokens are distributed), and (3) “Incentive system” (the way a token exerts influence over the network and/or its holder). Given the focus on on-chain governance, our scope is on token allocations and distributions—hence, the prime subject of study being on the “Supply” parameter. Oliveira et al. [47, p. 9] note how “Supply” strategies can either be on a one-time basis (“fixed”) or following increments (“schedule based”). Tokens can also be “pre-mined” (or “pre-sold” [28])—that is, a portion of the tokens is created and distributed before the official launch date. As such, they are means to overcome the chicken-and-egg problem [25].

For the supply of voting rights tokens, whose fair deployment is motivated by a normative ambition of political decentralization, the story is more ambiguous. Uniswap developers pre-mined a part of all voting rights tokens (UNI) and allocated some to a group of insiders. Among others, the DeFi projects SushiSwap (SUSHI) and MakerDAO (MKR) followed similar paths, opting for an allocation that favored insiders. Over time, in all of these cases, wealth concentration was eminent [6]. Voting rights tokens are not exclusive to DeFi projects—they are also used to distribute governance power in blockchain-based metaverse projects—where, similarly, concentration was observed [34].

Concentration of wealth and power is inherent to human societies and economic systems [52]. Pareto [48] exposed the land concentration in the Italian *novecento* (20th century), subsequently lending his name to the Pareto principle. Financial markets are no exception to the principle [42]. In cryptocurrency markets, the same phenomena may be observed. Both Bitcoin and Ethereum have centralized token distributions, and the trend appears to only increase [32]. This is true even in the case of proof-of-stake cryptocurrencies where the project's security is closely tied to the level of dispersion [57]. Likewise, Nadini et al. [45] and Klein et al. [38] observe concentration of wealth in markets for non-fungible tokens. When it comes to voting rights tokens, Barbereau et al. [5] identified that the level of concentration among voting rights tokens is even higher, highlighting cases where a handful of people hold more than 50% of all tokens. Interestingly, these tokens are barely used by their holders to vote [6].

Andre Cronje's project Yearn Finance (YFI) sought to eliminate favoritism and insider allocations [63]. By opting for the first, "fixed" supply strategy, YFI were *not* allocated to a minority of insiders. Instead, the implemented *fair launch* allocation followed the principle of 'fair equality of opportunity' (see [54])—effectively, the idea that each user has the same opportunity to obtain YFIs. Despite its failure to achieve an equitable distribution over time [5, 6], at least in theory, a fixed initial token allocation that is 'fair' would help achieve ambitions of decentral control. Evaluating this proposition is the subject of this study.

2.3 Agent-Based Modeling

To evaluate the phenomenon of initial token allocations and concentration of tokens, we designed ABM, which is a computational method used to simulate the actions and/or interactions of autonomous agents to understand how systems behave and what determines outcomes [40]. As an analytical method, applications of ABM are found in a variety of disciplines from energy and pathology to risk management and finance. In IS research, the value of ABM was acknowledged given its methodological versatility to investigate systems whose "emergent properties unfold over time" [35, p. 158] and its supportive value in the development and/or verification of theory [10, 21, 23]. The literature studies of Beese et al. [8] and Dong [23] illustrate the breadth of ABM applications in IS research—citing the potential for scholars to embed theory in the exploration of complex phenomena.

For the study of cryptocurrencies and DLT-based systems, numerous works applied ABM. Bornholdt and Sneppen [13] proposed a model to study the emergence of cryptocurrencies vis-à-vis Bitcoin—considering factors such as trading, mining of new coins, and agent-to-agent interactions. Their findings show that Bitcoin may be interchangeable with cryptocurrencies of similar characteristics. Cocco et al. [19] built an artificial cryptocurrency marketplace based on an order book simulation of the Bitcoin market where agents trade autonomously. Their model is able to reproduce real price formations and market volatility—hence, our adaptation of it in this work. Roşu and Saleh [57] propose an environment to model the behavior of investors/agents in a proof-of-stake-based blockchain of cryptocurrency issuance. They denote, contrary to expectations, that agents seek to stabilize their portfolio instead of accumulating more wealth.

3 DATA PREPARATION

Given their open and auditable characteristics, DeFi projects, for the most part, are built on public-permissionless ledgers [3, 59]. These ledgers provide a rich source for the collection and analysis of quantitative data. Chen and Bellavitis [17] observe that 80% of DeFi platforms, are in fact built on the Ethereum ledger. Ethereum records a variety of details, not least on tokens, data about their creation, their initial distribution, and transaction histories.




Fig. 1. Graphs for the Crypto FGI and YFI price.

Table 2. Overview of Data Extraction

Extracted Addresses	96,227
Addresses Used in Analysis	86,752
Extraction Period	2020-07-17 to 2021-08-15

The fair launch was originally created as part of Yearn Finance, hence its practice informs our study inductively [8]. Specifically, data on Yearn Finance is used to (1) define the base scenario S_0 (“Cronje”), (2) ‘feed’ our model based on reality, (3) calibrate the model, and (4) validate our model. For (2), we also extracted the price of YFI (from [CoinGecko.com](https://www.coingecko.com)) and the Crypto **Fear & Greed Index (FGI)**. The graphs for these two additional data sources are presented in Figure 1.

Yearn Finance is built on Ethereum and uses the ERC-20 token standard for its voting rights token YFI. YFI was launched with a fixed supply whereby no tokens were allocated to insiders. Instead, the initial supply of 30,000 tokens in circulation was distributed via a liquidity providing scheme. Users could earn YFI by supplying liquidity into three distinct pools, allowing every user, regardless of their initial capital or other restrictions, to earn a portion of YFI’s supply proportionate to the contributed liquidity. This type of allocation strategy was coined as a *fair launch* [63].

To generate data for our model, we used [Dune](https://dune.com) to extract those addresses that have been holding YFI from Ethereum’s public ledger. Then, we organized the data such that we could determine how many tokens are owned daily by each address. Finally, we excluded a number of address ‘types’ from the dataset: smart contracts (since they never utilized their voting rights, despite holding YFI [6]), addresses holding YFI valued less than \$1 (since these rarely vote or trade their tokens owing to Ethereum’s gas fees being significantly higher than the token’s value), and addresses used to burn tokens (e.g., 0x000...0000) (since no one controls them and YFI is effectively taken out of circulation). Table 2 presents our final dataset.

Following the finalization of our dataset, we utilized Exploratory Data Analysis (EDA) to determine the model’s initial conditions and variables. We chose September 1, 2020 (i.e., 45 days after the project’s launch) as the starting date since at that point the Yearn Finance fair launch took place; in other words, all tokens were allocated to and claimed by liquidity providers. Using the Anderson-Darling test [2] and the Akaike information criterion [1], we identified that the probability distribution of the initial YFI allocation follows a Lomax distribution ($\lambda = 0.4, \alpha = 0.5$). Relying on the same methods, we found that the daily number of new addresses that have been holding YFI increases following an asymmetric Laplace distribution $ALap(0.71, 58, 76)$.

4 THE MODEL

The proposed model for initial allocations builds on an agent-based artificial cryptocurrency market (c.f. [19]). Subsequently, we describe the model in terms of the agents, the market rules, and the trading behavior. Then, we describe the initial token allocations of the three fair launch scenarios. Finally, we introduce the metrics used to evaluate concentration over time.

4.1 Agents

For our model, we take timesteps $t \in \mathbb{N}_+ = \{1, 2, 3, \dots\}$ which correspond to a single day and a new, individual trading round. The first timestep in our model is at $t = 45$ (all tokens were allocated, claimed, and are in circulation). For each timestep, we define agents $i \in I$ as the addresses that hold voting rights tokens (TKNs) at the beginning of each trading round. The number of agents at timestep t is given by $N_A(t) \in \mathbb{N}_+$. At the beginning of each trading round, a subset of I is selected to trade TKNs (the selection mechanism as well as the trading strategy of agents is described subsequently) and new agents (endowed *solely* with fiat currency) enter the market with the desire of placing buy orders to acquire TKNs. The new agents entering ($N_A(t+1) - N_A(t)$) follow $ALap(0.71, 58, 76)$ for every $t > 45$. In other words, at each trading round, the 95% confidence interval of the number of new agents is [114, 119]. We run our model for 347 days ($t = 392$), and the 95% confidence interval for the final number of agents ($N_A(392)$) is [47113, 50890].

Agents are endowed with fiat holdings $f_i(t)$ and TKN holdings $y_i(t)$. Based on Dragulescu and Yakovenko [24] and Brzezinski [15], the amount of fiat held by both individual agents at $t = 45$ and those agents entering the market at each trading round is drawn from a Pareto distribution with $\alpha = 2.1$ and $\min(f_i(t)) = \$400k$ for the richest 10% of our agents and from an $Exp(\frac{1}{40000})$ for the remaining, bottom 90%. The amount of TKNs held by agents at $t = 45$ depends on the chosen fair launch scenario S_i with $i \in \{0, 1, 2\}$.

Independently of fiat or TKN holdings, each agent i is assigned to one of two populations, **Diamond Hands (DHs)** and **Random Traders (RTs)**, representative of respective trading strategies. DHs are risk-averse traders, who pragmatically invest in the market and are more likely to not incur in trades. RTs, then, are agents who enter to market for a variety of reasons (portfolio diversification, gambling, etc.). Following Cocco et al. [19], the agent populations is split into 30% DHs and 70% RTs.

4.2 Market Rules

The TKN market is given by a mechanism *comparable* to a clearing house, whereby buy and sell orders are accumulated over time and cleared ('matched') periodically [41]. The purpose of the model and developed market is not exploring how price is formed; instead, it is to simulate how tokens circulate (and concentrate) based on clear conditions. The mechanism we utilize is *not* a formal clearing house, as it does not account for price formation, nor does it include adjustments of price after every transaction over time. Instead, at each timestep, the TKN price $T_p(t)$ is updated based on YFI's historical price data. Agents can autonomously decide whether they are willing to trade. Agents do not, however, have information about the orders other agents are placing. The scope of this work and the developed ABM is on token concentration, not the way price is formed. Clearing houses offer a simple and effective system for matching orders between agents, requiring limited computation overhead and facilitating a realistic flow of tokens, which makes them suitable for large ABMs [18].

At $t \geq 45$, the total number of tokens in circulation is given by the constant $T_s = 36666$. For the trade of TKNs, we model a two-sided market with a number of buyers, each willing to buy TKNs, and several sellers, each willing to sell TKNs. Additionally, at every timestep, the buy/sell orders

created by the agents are matched in a *first in first out* method, and at the end, the unmatched orders are canceled.

4.3 Trading Behavior

Depending on the population agents belong to, they exert a choice—to trade (T) or not to trade—at every timestep t . This decision is given by the probability $P_i(T)$. For RTs, who randomly wish to trade following a uniform distribution, $P_i(T) = 0.5$. For DHs, $P_i(T)$ is dependent on two independent variables. The first is the FGI ($FGI(t)$), which fluctuates between a value of 0 (“Extreme Fear”) and 100 (“Extreme Greed”) [70]. In our case, we consider the values of the index as “Extreme” (FGI_e) when $FGI(t) > th_h$ or $FGI(t) < th_l$, and “Normal” (FGI_n) when $th_h > FGI(t) > th_l$ with th_h and th_l thresholds for the extreme values of FGI. The second is the agent’s wealth (W), given by an agent’s individual holding denominated in fiat $f_i(t)$. An agent’s wealth at time t is considered “High” (W_h) when $f_i(t)$ is above the 90th percentile of the wealth distribution and “Low” otherwise. Therefore, the probability of a DH agent to trade is given by

$$P_i(T) = P(T|FGI_e, W_h)P(FGI_e)P(W_h) + P(T|FGI_e, W_l)P(FGI_e)P(W_l) \\ + P(T|FGI_n, W_h)P(FGI_n)P(W_h) + P(T|FGI_n, W_l)P(FGI_n)P(W_l). \quad (1)$$

If an agent is willing to trade, the subsequent decision to execute a buy or sell order depends on the population they belong to. For RT, the buy and sell orders follow a Bernoulli distribution with $p = 0.5$. Initially, the same holds for DH, but in their case the probability is calibrated at a later stage based on the data from Yearn Finance (c.f. Section 5). In the trading behavior, we do not consider protocols that allow to stake/sell voting rights token entitlements (e.g., Bribe Protocol), as these add yet another degree of complexity. At each timestep, the amount of fiat currency an agent spends on buying tokens follows a $\mathcal{N}(\mu = \frac{f_i(t)}{2}, \sigma = \frac{\mu}{3})$ and the number of TKNs an agent sells follows a $\mathcal{N}(\mu = \frac{y_i(t)}{2}, \sigma = \frac{\mu}{3})$. In our model, the average buy and sell values are considerably higher than those in the work of Cocco et al. [19]. This choice was intentional to increase the trading volume since agents trade at most once per day.

Admittedly, the agents of our model have limited intelligence since the only market factor that influences their decisions is the FGI. Although agents with more sophisticated decision making (e.g., taking into consideration the price volatility of YFI) could potentially enhance the accuracy of the model, they are challenging to implement in our case due to the large number of agents considered. Despite that, existing literature indicates that ABMs that are composed of even less sophisticated agents (also known as *zero-intelligence* agents) have been able to capture some of the core characteristics of financial markets (see [27, 33]).

4.4 Fair Launch Scenarios

Our simulation is set up around three distinct scenarios representative of initial token allocations understood as ‘fair.’ Their design is informed on the basis of the epistemic dichotomy described in the work of Beese et al. [8] as well as in that of Dong [23].

The ‘base’ scenario, S_0 , is created following an inductive approach (its design is informed by data extracted from Yearn Finance), and the distribution of T_s is modeled to follow a Lomax distribution with $\lambda = 0.4$ and $\alpha = 0.5$. The artificially created fair launch scenarios S_1 and S_2 are designed following a deductive approach on the basis of theory—here, philosophical interpretations of what might be seen as ‘fair.’

The first alternative scenario, S_1 (“Bentham”), considers ‘fairness’ in egalitarian terms: equity is achieved in terms of uniformity such that the total supply of tokens is divided equally among the participants. Formerly, it considers Jeremy Bentham’s dictum that “everybody to count for

Table 3. ‘Fair’ Initial Token Allocation Scenarios

	Scenario	Allocation	Perspective
S_0	“Cronje” (Yearn Finance)	Everyone gets the same opportunity	Social liberalism
S_1	“Bentham”	Everyone gets the same	Egalitarianism
S_2	“Rawls”	Everyone gets a random amount	Darwinism

one, nobody for more than one” (see [43]), without consideration of individual interests or material situation. For S_1 , T_s is uniformly distributed such that each agent i at $t = 45$ holds $y_i(45) = \frac{T_s}{N_A(45)}$.

The second alternative scenario, S_2 (“Rawls”), considers randomness, and more specifically, the principle of a lottery as ‘fair’: equity is achieved in terms of a token allocation—at random—and in our case following a Normal distribution. It rehashes the idea that the outcome of each individual’s position, like the outcomes of ordinary lotteries, is a matter of good or bad “luck” [53, pp. 74–75]. Randomness and chance are central to the theory of Darwinian evolution [68]. For S_2 , T_s is distributed among agents following a truncated Normal distribution ($\mu = 0.103$, $\sigma = 0.192$) defined on $[0, \infty]$.

In sum, we investigate two additional scenarios aside from the base scenario (Table 3). While keeping the market conditions and parameters fixed, changing the initial allocation of tokens provides further insight into the concentration of wealth.

4.5 Metrics

Given the aim of analyzing the distribution of voting rights tokens post fair launch allocation, select metrics are computed at every timestep. These metrics are the Gini coefficient [31] and Shannon entropy [62]. This choice was made based on an evaluation of related works seeking to quantify and measure the distribution of tokens in a system—notably as discussed in the work of Gervais et al. [30], Gochhayat et al. [32], Barbereau et al. [6], and Klein et al. [38].

The Gini coefficient is typically used to assess the distribution of wealth in a given country. It was, however, also applied to study wealth distribution in Bitcoin and Ethereum [32], in non-fungible tokens markets [38], and the distribution of voting rights tokens in DeFi projects [5, 6]. For our model, the Gini G indicates the concentration of TKNs amid agents. The Gini is given by

$$G = \frac{\sum_{i=1}^{N_A} \sum_{j=1}^{N_A} |p_i - p_j|}{2N_A \cdot \sum_{j=1}^{N_A} p_j}, \quad (2)$$

where p_i corresponds to the share of TKNs held by agent i and N_A the total number of agents. It is maximized through the Dirac distribution δ_{i_0} (i.e., $p_{i_0} = 1$ for some $i_0 \in \{1, \dots, N_A\}$ and $p_i = 0$ for all $i \neq i_0$) and minimized through the uniform distribution (i.e., $p_i = \frac{1}{N_A}$ for all i).

Shannon entropy was initially developed to assess the information loss in telecommunication networks [62]. It has found application in multiple areas predominantly as a measure of uncertainty or randomness. In the context of blockchain technology, it was used to measure the level of decentralization in the consensus mechanism [32], network structure [29], and governance [67] of both Bitcoin and Ethereum.

Normalized Shannon Entropy (NSE) then takes values between 0 and 1, and determines the unpredictability of a distribution. We assume that a system where the voting tokens are distributed can exhibit high unpredictability (1), given that more agents influence the outcomes. The NSE is

Table 4. Optimal Parameter Values

Parameter	Description	Optimal Value
$P_{DH}(Buy)$	Buy probability of DH	0.7
$DH/N_A(t)$	Population share of DH	0.3
th_h	FGI threshold high	80
th_l	FGI threshold low	20
$P(T FGI_e)$	Trading probability under extreme market conditions	0.7
$P(T W_h)$	Trading probability under high wealth	0.7
$P(T FGI_n)$	Trading probability under normal market conditions	0.8
$P(T W_l)$	Trading probability under low wealth	0.9

given by

$$NSE = - \sum_{i=1}^{N_A} \frac{p_i \log(p_i)}{\log N_A}, \quad (3)$$

where $0 \log(0) \equiv 0$ by convention since $\lim_{p \rightarrow 0} p \log(p) = 0$. It is 0 for δ_{i_0} and 1 for the uniform distribution (i.e., the extremes are interchanged compared to the Gini coefficient). To ease graphical observation, we opted to consider 1-NSE instead of NSE such that, as in Gini, higher values correspond to higher degrees of centrality.

5 IMPLEMENTATION AND CALIBRATION

The model was implemented in Python using the MESA framework [37]. The simulation and calibration of the model was performed in a **High Performance Computing (HPC)** facility. The hardware provided, depending on the allocation of the HPC, was a Dual Intel Xeon Broadwell or Skylake with 128 GB of RAM.

For the calibration, we followed the recommendations of Richiardi et al. [56, p. 4], whereby a “full exploration” of the parameters is required. To do so, we implemented a **Grid Search (GS)** to find a set of optimal values of parameters. GS performs an exhaustive search over all the possible combinations of parameters until finding the optimal one. The goodness of the fit and the stopping condition of GS are computed using **Root Mean Squared Error (RMSE)** and **Mean Absolute Percentage Error (MAPE)** the actual (extracted from the dataset) and calibrated model. The respective equations are given by

$$RMSE = \sqrt{\left(\frac{1}{n}\right) \sum_{i=1}^n (x_i - \hat{x}_i)^2} \quad \text{and} \quad MAPE = \frac{100}{n} \sum_{i=1}^n \left| \frac{x_i - \hat{x}_i}{x_i} \right|, \quad (4)$$

where x_i is the actual observation and \hat{x}_i is the simulated value.

The optimization was executed over all eligible parameters. The FGI threshold takes values from 0 to 100. For all other parameters, we considered values between 0 and 1. The optimal parameter values are displayed in Table 4.

The optimal values of the DH/RT population ratio were close to the ratio used by Cocco et al. [19]. Therefore, we fixed it at 30% DHs and 70% RTs. Similarly, the buy probability was optimized to be 70% for DH. For the DHs, we found that high trading probabilities indeed lead to lower error rates (Table 5). This, as demonstrated in Roşu and Saleh [57], represents an expected behavior as more trading is linked with higher wealth concentration. Diametrically opposed to the high trading probability parameter set is an artificially created parameter set, with relatively low trading probabilities. The error rates for this parameter set are relatively worse than the optimal set of

Table 5. DH Trading Probabilities Parameter Sets with Their Correspondent Error Rates

Parameter	High	Medium	Low
$P(T FGI_e)$	0.7	0.3	0.1
$P(T W_h)$	0.7	0.4	0.1
$P(T FGI_n)$	0.8	0.3	0.2
$P(T W_l)$	0.9	0.5	0.2
$P_i(T)$	0.77	0.38	0.15
MAPE	0.1859	0.224	0.255
RMSE	0.007	0.009	0.012

high trading probabilities. We also artificially generated and investigated a compromise between the two sets (medium probability) without extreme trading probabilities.

In sum, we investigate three scenarios (see Table 3) under three trading probabilities (see Table 5). The results might vary due to the stochastic nature of ABM. In anticipation of this variance and to ensure the robustness of our results, we applied a Monte Carlo method [39] by repeating the experiment of the three simulation sets within HPC, resulting in more than 1,000 simulations (or approximately 300 per set of trading probabilities). For all simulations, the agents can place buy or sell orders depending on the probability defined in $P_{DH} = 0.7$ for DH (optimized value) and $P_{RT} = 0.5$ for RT (constant), respectively.

6 SIMULATION RESULTS

6.1 Effects of Trading Probabilities on the Three Scenarios

The first simulation considers the model's behavior under high trading probability (Figure 2). High trading probability refers to a relatively high likelihood for DH agents to place an order. The second simulation is the artificially created edge case with low trading probabilities (Figure 3). It is diametrically opposed to the former and explores the behavior of DH agents when the market dictates a relatively low likelihood to place a trade. (The graphs for S_0 and S_2 are visually coinciding.) The third simulation set was created artificially as middle ground between the high and low trading probabilities (Figure 4).

Across all trading probabilities, for S_0 and S_2 , the Gini values of the simulated data overlap considerably with the extracted YFI data for $t \geq 100$. Until the end of the simulation time frame, the Gini values of the extracted data diverge by at most by 1.1% on S_0 and 1.7% on S_2 across all trading probability scenarios. The close fit between the extracted data from Yearn Finance and the three simulations, and in particular for S_0 , is expected given the performed optimizations. The exception to the convergence is S_1 , whose graph is below the extracted data in all three simulations. It appears that an egalitarian initial token allocation would then lead to relatively less concentration in the time frame of the simulation.

We expect variations in the initial values of Gini and NSE, because even though our model starts with the same number of agents at each simulation, the initial token allocation is not fixed, and as discussed earlier, the token distribution in S_0 follows a Lomax distribution. Regardless of the scenario and trading probability, after the 1-NSE stabilizes, all distributions move in a lateral and parallel direction with regard to the extracted 1-NSE values from YFI. Although our simulation results seem to coincide with the Yearn Finance data for Gini, we observe high variations of the NSE. This divergence may be interpreted in two ways.




Fig. 2. Simulations for the parameter set representative of high trading probabilities.




Fig. 3. Simulations for the parameter set representative of low trading probabilities.

At inception of Yearn Finance some contracts held a large amount of YFI and distributed them shortly after. Given the scope of our model, we did not consider such behavior. We argue that the smart contracts that emitted YFI rapidly result in sharper rises in the metrics' values. This is consistent with our findings, which indicate that when there is a large amount of YFI accessible for trading in a short period of time, the metrics rise.

The second interpretation pertains to the token supply (T_s). Formerly, the supply of YFI was “schedule based” [47]: it began with a supply of 30,000 YFI allocated following the fair launch, and subsequently an additional 6,666 YFI tokens were distributed. For simplicity, we start with 36,666 supply dispersed to the starting holders in our simulations. Again, the distribution of 6,666 YFI in a short period of time theoretically results in higher concentration than our model, which, in contrast, distributes YFI more slowly over time. The schedule-based supply of YFI can be observed




Fig. 4. Simulations for the parameter set representative of medium trading probabilities.

Table 6. Share of Agents That Control 90% of TKNs in Circulation at $t = 392$

Scenario	High Probability		Medium Probability		Low Probability		Extracted	
	Percentage	Actual	Percentage	Actual	Percentage	Actual	Percentage	Actual
S0 Cronje	2.59%	1137/43830	2.63%	1188/45092	3.63%	1499/41248	2.02%	849/41926
S1 Bentham	10.80%	4777/44214	10.73%	4847/44950	11.38%	4999/43895		
S2 Rawls	0.76%	376/49397	1.29%	572/44300	2.83%	1250/44056		

in the ‘bumps’ around $t = 100$. While the change is more subtle in Gini, within 1-NSE the change is more clearly observable. This is due to the comparatively higher sensitivity of the latter metric with regard to minor fluctuations (see [6]), which can also be observed in the preceding figures, where the standard deviation of NSE is substantially higher than that of Gini.

6.2 Actual Concentration of Wealth amid *Whales*

Following the three simulation sets focusing on the trading probabilities, we performed a more granular analysis of the actual concentration of TKNs amid the population of agents ($N_A(392)$). In particular, we sought to investigate the share of agents that hold 90% of all tokens in circulation. These agents are so-called *whales*, “‘wealthy,’ above-average token-holders” [6]. In consideration of the amount of available data following the Monte Carlo simulations, in Table 6 we present a more feasible analysis on the basis of the results from the first simulation round. For the sake of comparison, the extracted column refers to reality (i.e., the Yearn Finance data).

Unsurprisingly, in consideration of the values metrics took in the previous analysis, we observe a concentration of TKNs in the hands of the few. These few individuals are *de facto* in control, as they may exert significant political pressure. In relative terms, as reflected in the metrics, the egalitarian allocation S_1 shows that the actual number of whales is higher. Regardless, our results align and support the timocratic description of DeFi governance by Barbereau et al. [6] and the observations on whales (in metaverse DAOs) by Goldberg and Schär [34].

6.3 Extending the Simulation of S_1 “Bentham”

After running the first set of simulations, we opted to run a separate simulation to explore whether S_1 indeed demonstrates more or less concentration over time. To do so, we extended the simulation




Fig. 5. Simulation for the Bentham scenario under the three trading probabilities.

rounds from $t = 392$ (August 15, 2021, the last data point extracted from Yearn Finance) to $t = 545$ (March 1, 2022, the last point of the simulations). This represents an extension of 44.39%. The results of the simulation are presented in Figure 5.

From our previous simulations on the effects of different trading probabilities, we observe how S_1 Bentham's initial token allocation positively affects both metrics. It is to be expected that an equal distribution of tokens at origination will reduce concentration, at least early on. In this simulation, we observe similar phenomena to what was demonstrated in the work of Roşu and Saleh [57]: even though the delayed effect of an egalitarian initial token allocation (like the one simulated) might generate, concentration is, judging from our simulation results, inevitable in the long run. To corroborate this observation, we fitted linear regressions on Gini from the three trading scenarios. In the worst-case scenario, the slope of the linear regression is $4 * 10^{-4}$.

7 VALIDATION OF THE MODEL

Validation is an essential part of ABM [8, 21, 23]. There are numerous techniques for validation, all of which are used to establish credibility in the simulations [8]. To validate our model, we opt to use three different techniques: *event validity*, *parameter variability (sensitivity analysis)*, and *extreme condition tests*.

7.1 Event Validity

For the event validity, simulated events are compared with those occurring in real-world systems [8]. Specifically, we take the share of agents that control 90% of TKNs in circulation between $t = 1$ and $t = 392$ for S_0 Cronje. (The values at $t = 392$ are identical to those displayed in Table 6.) The real-world (extracted) data is taken from Yearn Finance (c.f. Section 3). The comparison between these datasets is displayed in Figure 6.

The model during calibration was not given any information regarding the token concentration. From Figure 6, we visualize how both the simulated and real-world values converge after approximately 100 steps (100 natural days). At the end of the simulations, the difference is 0.4 percentage




Fig. 6. Validation through event validity for the share of agents that control 90% of the circulation between Yearn Finance and S_0 .




Fig. 7. Impact of different population allocations on Gini and NSE.

points. These results present a solid base for the validity of the model, as it closely replicates reality [21].

7.2 Parameter Variability

For the parameter variability, input parameters are modified and the resulting changes analyzed [8]. We evaluated the impact of alternative DH/RT population ratios—with 10%, 30%, 50%, 70%, and 90% DH—on the scenario S_0 relative to reality. The change was observed in terms of the metrics. Here, too, we applied the Monte Carlo method using HPC. Figure 7 gives the simulated metrics for the five ratios along the actual distribution of the Yearn Finance token.

The procedure yields variability between the different population ratio and reality. We define the Δ as the difference between the simulated scenario and the extracted data. The simulations with 90% ($\Delta_{Gini} = 1.95\%$; $\Delta_{NSE} = 34.89\%$), 70% ($\Delta_{Gini} = 1.30\%$; $\Delta_{NSE} = 29.42\%$), and 50% DH




Fig. 8. Impact of extreme population allocations on Gini and NSE.

($\Delta_{Gini} = 0.71\%$, $\Delta_{NSE} = 25\%$) perform relatively worse than those with 30% ($\Delta_{Gini} = 0.015\%$; $\Delta_{NSE} = 19.65\%$) and 10% ($\Delta_{Gini} = 0.4\%$; $\Delta_{NSE} = 16.52\%$). In consideration of the Δ values and Cocco et al. [19] (who take 70% irrationality), the 30% DH is most appropriate and therefore justifies the models' validity [21].

7.3 Extreme Conditions Test

For the extreme conditions tests, we tested whether our model behaves reasonably when extreme values are selected for specific parameters [8]. To do so, we selected parameters to be tinkered with, all while keeping the other parameters at their optimal value (refer to Table 4). Here, the considered scenario is S_0 relative to reality.

In the first analysis, we considered extreme values for the ratio of DH. Figure 8 displays our model if the population share of DH ($DH/N_A(t)$) is at 0.01 and 0.99. The results displayed here support our argument that the more the tokens are traded, the more concentration is to be observed. Conversely, fewer trades yield less concentration.

In the second analysis, we considered extreme values for the FGI threshold (th). Figure 9 displays our model under different th_h and th_l . As we note minor differences ($0.33 < \Delta_{Gini} < 0.54$; $2.71 < \Delta_{NSE} < 4.56$) between the extreme FGI scenarios and the optimal scenario, we conclude that the impact of extreme values for the FGI threshold is moderate on the model. This meets expectations as the model ought to behave “reasonably” when extreme values are selected [8, p. 512].

8 DISCUSSION

Using ABM, we evaluated how trading probabilities affect distribution over time within three distinct scenarios representative of ‘fair’ initial token allocations. Our findings are consistent with Barbereau et al. [6] timocratic description, as the ability to trade voting rights tokens appears to be one of the causes of concentration (RQ1). Amid all three simulation sets with high, medium, and low trading probabilities, the three scenarios tend toward concentration (RQ2). Subsequently, we discuss our contributions, implications, and limitations.




Fig. 9. Impact of extreme FGI thresholds on Gini and NSE.

Table 7. Italicized Refinements to the Token Classification of Oliveira et al. [47]

Governance Parameters	Representation		Digital		Physical		Legal	
	Supply	Distribution	Schedule based	Pre-mined, scheduled distribution		Pre-mined, one-off distribution		Discretionary
		Allocation	Fair launch			'Unfair' launch, pre-sale		
	Incentive system		Enter Platform		Use Platform		Stay Long-Term	

8.1 Contributions and Implications

The concentration of wealth in the long term, as observed in our constructed ABM, aligns with findings on the concentration of wealth in public-permissionless [6, 32, 34], and general understandings on the concentration of wealth [52]. The implications of our findings are of theoretical and practical nature.

Our findings allow to “sharpen” theory [21, p. 440] on tokenomics and DAO governance. Specifically, we contribute to the token classification of Oliveira et al. [47], as we refine the “Governance Parameters” in favor a distinction of the “Supply” parameter in terms of “Distribution” and “Allocation.” The “Distribution” parameter is accounted for already, as it is equivocally used for the “Supply” of tokens. For “Allocations,” we distinguish between “Fair Launch” allocations (like the ones described) and all other token allocations that may favor a minority of insiders (e.g., like Uniswap did). This contribution parallels research on initial coin offerings which account for these ‘unfair’ allocations as so-called “private pre-sale[s]” [28, p. 10]—a terminology we adopt here.

Table 7 showcases our refinement vis-à-vis the original classification of Oliveira et al. [47]. While our findings do not allow to distinguish causation or correlation between allocation and concentration of tokens over time, the inclusion of “Allocations” in the token classification provides an indication for the normative ambitions of on-chain governance frameworks. Certainly, these are of value to research on DAOs, their governance, and concentration of power (see [6, 9, 34, 36, 58]).

The practical implications of our findings are for the design of future governance frameworks that leverage voting rights tokens. Our work provided additional evidence that trading largely determines the extent to which governance power is concentrated. Hence, beyond the alternative to move governance of DAOs ‘off-chain’ (which carries a high risk of concentration of power (see [22, 26, 49, 55])), the possibility to transfer tokens must be addressed.

In practice, this can be achieved through a new class of tokens described as *soulbound*. The introduced definition refers to “accounts, or wallets, that hold publicly visible, non-transferable (but possibly revocable-by-the-issuer) tokens” [69, p. 2]. In other words, the (albeit pseudonymous) identity of a holder is encrypted into a **Soulbound Token (SBT)** that is linked to the respective wallet. The opportunities for on-chain governance are promising:

- They mitigate Sybil attacks.
- They (could) grant more voting power to reputable holders.
- They enable for “proofs-of-personhood.”
- They allow to correlate between SBTs which support particular causes and prevent a “tyranny of the majority” [44].

These opportunities provide avenues for research as they require contextual analysis. To date, we note the intended application of SBTs for Know-Your-Customer processes and user credentials as the cryptocurrency exchange Binance stipulated the intent to explore SBT on its native blockchain. Binance’s SBT would grant access to specific functions of the BNB Chain [50]. Another, experimental application of SBTs is in electronic health records (see [65]). It remains to be seen how these are implemented in practice and to what extent they achieve said promises.

8.2 Limitations

The research is subject to a number of limitations. In turn, some also offer potential research directions in the field of modeling cryptocurrency/token markets and organizational works in IS research.

The first limitation pertains to the defined market rules. Although these followed the principles stipulated in the work of Mendelson [41], the clearing mechanism lacks a formal price clearing method. As we aim to replicate the behavior of YFI and evaluate the concentration of tokens, we acknowledge that the defined market mechanics may result in an oversimplification of reality. Fitting a model that accurately replicates the price of the mentioned asset poses challenges due to the high volatility and stochasticity observed in cryptocurrency markets. (This limitation was acknowledged by Cocco et al. [19].)

The second limitation pertains to the awareness of agents. In the designed model, the decision making of individual agent’s does not depend on past decisions or those of other agents. To address these shortcomings, future work may build upon and extend our model to include a public order book where agents are aware about other orders. Further, a distinction may be made between trading mechanisms and clearing methods on centralized and decentralized exchanges.

The third limitation pertains to the trading behavior of agents. For our simulations, we heavily rely on the FGI as a proxy for market conditions. Subsequent work could opt for the use of more granular indicators, such as the price or the volatility of different DeFi assets or social media data.

Last, because this study focused on initial token allocations, we did not consider the different techniques to cast the actual votes. While the allocation might lead to concentration, specific voting techniques may prevent whales from skewing the outcomes of elections. *Quadratic voting*—whereby preferences in terms of strength instead of a simple ballot are cast [16]—is now effectively used in some DAOs (e.g., Synthetix) and discussed in others (e.g., OlympusDAO, CurveDAO). That method, however, is yet to prove fruitful in the long run: one of the considered cases in the work

of Barbereau et al. [6], Synthetix, uses quadratic voting, and although the findings reveal relatively less concentration, all metrics are still high. The evaluation of different voting techniques represents a notable research direction for interdisciplinary IS research.

9 CONCLUSION

Within the DeFi space, recent scholarship observed the implementation of on-chain governance frameworks for DAOs that leverage tokens embedded with voting rights. The initial allocation of these voting rights tokens ought to follow principles of fairness to achieve normative goals of political decentralization. The fair launch allocation of Andre Cronje gained prominence, as it did not allocate any tokens to a minority of insiders. However, in practice it fell short, as over time YFI tokens became highly concentrated.

The contributions of this study are threefold. First, on the basis of Cocco et al. [19] and Roşu and Saleh [57], we proposed ABM to simulate fair launch initial token allocation. Using the model, we simulated alternative initial token allocation scenarios understood as ‘fair’ [43, 53]. Second, as our simulation results showed, over time, independently of market conditions and agents’ willingness to trade, concentration is imminent. At last, the implications of our results allowed to extend understandings on DAOs and tokenomics to formerly include allocations as formal part of developed governance understandings.

ACKNOWLEDGMENTS

The authors thank Reilly Smethurst and Joachim Geske for their friendly reviews. The authors also thank Gilbert Fridgen for his valuable feedback as part of the first submission.

REFERENCES

- [1] H. Akaike. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19, 6 (Dec. 1974), 716–723. <https://doi.org/10.1109/TAC.1974.1100705>
- [2] T. W. Anderson and D. A. Darling. 1952. Asymptotic theory of certain “goodness of fit” criteria based on stochastic processes. *Annals of Mathematical Statistics* 23, 2 (June 1952), 193–212. <https://doi.org/10.1214/aoms/1177729437>
- [3] Raphael Auer, Bernhard Haslhofer, Stefan Kitzler, Pietro Saggese, and Friedhelm Victor. 2023. *The Technology of Decentralized Finance (DeFi)*. Technical Report. Bank for International Settlements. <https://www.bis.org/publ/work1066.htm>
- [4] Tom Barbereau and Balázs Bodó. 2023. Beyond financial regulation of crypto-asset wallet software: In search of secondary liability. *Computer Law & Security Review* 49 (2023), 105829. <https://doi.org/10.1016/j.clsr.2023.105829>
- [5] Tom Barbereau, Reilly Smethurst, Orestis Papageorgiou, Alexander Rieger, and Gilbert Fridgen. 2022. DeFi, not so decentralized: The measured distribution of voting rights. In *Proceedings of the 55th Annual Hawaii International Conference on System Sciences*. <https://doi.org/10.24251/HICSS.2022.734>
- [6] Tom Barbereau, Reilly Smethurst, Orestis Papageorgiou, Johannes Sedlmeir, and Gilbert Fridgen. 2023. Decentralised Finance’s timocratic governance: The distribution and exercise of tokenised voting rights. *Technology in Society* 73 (May 2023), 102251. <https://doi.org/10.1016/j.techsoc.2023.102251>
- [7] Roman Beck, Christoph Müller-Bloch, and John Leslie King. 2018. Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems* 19, 10 (2018), 1. <https://aisel.aisnet.org/jais/vol19/iss10/1>
- [8] Jannis Beese, M. Kazem Haki, Stephan Aier, and Robert Winter. 2019. Simulation-based research in information systems. *Business & Information Systems Engineering* 61, 4 (2019), 503–521. <https://doi.org/10.1007/s12599-018-0529-1>
- [9] Cristiano Bellavitis, Christian Fisch, and Paul P. Momtaz. 2022. The rise of decentralized autonomous organizations (DAOs): A first empirical glimpse. *Venture Capital* 25, 2 (2022), 187–203. <https://doi.org/10.1080/13691066.2022.2116797>
- [10] Nicholas Berente, Stefan Seidel, and Hani Safadi. 2018. Research commentary—Data-driven computationally intensive theory development. *Information Systems Research*. Published Online, December 20, 2018. <https://pubsonline.informs.org/doi/epdf/10.1287/isre.2018.0774>
- [11] Balázs Bodó, Jaya Klara Brekke, and Jaap-Henk Hoepman. 2021. Decentralisation: A multidisciplinary perspective. *Internet Policy Review* 10, 2 (2021), 1–21. <https://doi.org/10.14763/2021.2.1563>
- [12] Bruce M. Boghosian. 2019. Is inequality inevitable? *Scientific American* 321, 5 (2019), 70–77.

- [13] Stefan Bornholdt and Kim Sneppen. 2014. Do Bitcoins make the world go round? On the dynamics of competing crypto-currencies. *arXiv preprint arXiv:1403.6378* (2014).
- [14] Thomas Bourveau, Emmanuel T. De George, Atif Ellahie, and Daniele Macciocchi. 2022. The role of disclosure and information intermediaries in an unregulated capital market: Evidence from initial coin offerings. *Journal of Accounting Research* 60, 1 (2022), 129–167. <https://doi.org/10.1111/1475-679X.12404>
- [15] Michal Brzezinski. 2014. Do wealth distributions follow power laws? Evidence from ‘rich lists.’ *Physica A: Statistical Mechanics and Its Applications* 406 (July 2014), 155–162. <https://doi.org/10.1016/j.physa.2014.03.052>
- [16] Vitalik Buterin. 2019. Quadratic Payments: A Primer. Retrieved February 29, 2024 from <https://vitalik.eth.limo/general/2019/12/07/quadratic.html>
- [17] Yan Chen and Cristiano Bellavitis. 2020. Blockchain disruption and decentralized finance: The rise of decentralized business models. *Journal of Business Venturing Insights* 13 (June 2020), e00151. <https://doi.org/10.1016/j.jbvi.2019.e00151>
- [18] Silvano Cincotti, Marco Raberto, and Andrea Tegli. 2010. Credit money and macroeconomic instability in the agent-based model and simulator Eurace. *Economics* 4, 1 (Dec. 2010), 1–32. <https://doi.org/10.5018/economics-ejournal.ja.2010-26>
- [19] Luisanna Cocco, Giulio Concas, and Michele Marchesi. 2017. Using an artificial financial market for studying a cryptocurrency market. *Journal of Economic Interaction and Coordination* 12, 2 (2017), 345–365. <https://doi.org/10.1007/s11403-015-0168-2>
- [20] Andre Cronje. 2021. Fair launches, decentralized collaboration, and Fixed Forex. *Medium*. (July 2021). Retrieved February 29, 2024 from <https://andrecronje.medium.com/>
- [21] Jason P. Davis, Kathleen M. Eisenhardt, and Christopher B. Bingham. 2007. Developing theory through simulation methods. *Academy of Management Review* 32, 2 (2007), 480–499. <https://doi.org/10.5465/amr.2007.24351453>
- [22] Primavera De Filippi and Benjamin Loveluck. 2016. The invisible politics of Bitcoin: Governance crisis of a decentralised infrastructure. *Internet Policy Review* 5, 3 (Sept. 2016). <https://doi.org/10.14763/2016.3.427>
- [23] John Qi Dong. 2022. Using simulation in information systems research. *Journal of the Association for Information Systems* 23, 2 (2022), 408–417. <https://doi.org/10.17705/1jais.00743>
- [24] Adrian Dragulescu and Victor M. Yakovenko. 2001. Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. *Physica A: Statistical Mechanics and Its Applications* 299, 1-2 (Oct. 2001), 213–221. [https://doi.org/10.1016/S0378-4371\(01\)00298-9](https://doi.org/10.1016/S0378-4371(01)00298-9)
- [25] Benedict J. Drasch, Gilbert Fridgen, Tobias Manner-Romberg, Fenja M. Nolting, and Sven Radszuwll. 2020. The token’s secret: The two-faced financial incentive of the token economy. *Electronic Markets* 30, 3 (2020), 557–567. <https://doi.org/10.1007/s12525-020-00412-9>
- [26] Quinn DuPont. 2017. Experiments in algorithmic governance—A history and ethnography of “The DAO,” a failed decentralized autonomous organization. In *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance*. Routledge, Taylor & Francis Group, London, 1–21.
- [27] J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. 2005. The predictive power of zero intelligence in financial markets. *Economic Sciences* 102, 6 (2005), 2254–2259.
- [28] Gilbert Fridgen, Ferdinand Regner, André Schweizer, and Nils Urbach. 2018. Don’t slip on the initial coin offering (ICO): A taxonomy for a blockchain-enabled form of crowdfunding. In *Proceedings of the 26th European Conference on Information Systems (ECIS ’18)*.
- [29] Adem Efe Gencer, Soumya Basu, Ittay Eyal, Robbert van Renesse, and Emin Gün Sirer. 2018. Decentralization in Bitcoin and Ethereum networks. *arXiv:1801.03998 [cs.CR]* (2018).
- [30] Arthur Gervais, Ghassan O. Karame, Vedran Capkun, and Srdjan Capkun. 2014. Is Bitcoin a decentralized currency? *IEEE Security & Privacy* 12, 3 (May 2014), 54–60. <https://doi.org/10.1109/MSP.2014.49>
- [31] C. Gini. 1912. *Variabilità e Mutabilità*. Libreria Eredi Virgilio Veschi, Rome. <https://ui.adsabs.harvard.edu/abs/1912vamu.book....G>
- [32] Sarada Prasad Gochhayat, Sachin Shetty, Ravi Mukkamala, Peter Foytik, Georges A. Kamhoua, and Laurent Njilla. 2020. Measuring decentrality in blockchain based systems. *IEEE Access* 8 (2020), 178372–178390. <https://doi.org/10.1109/ACCESS.2020.3026577>
- [33] Dhananjay K. Gode and Shyam Sunder. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101, 1 (Feb. 1993), 119–137. <https://doi.org/10.1086/261868>
- [34] Mitchell Goldberg and Fabian Schär. 2023. Metaverse governance: An empirical analysis of voting within decentralized autonomous organizations. *Journal of Business Research* 160 (2023), 113764. <https://doi.org/10.1016/j.jbusres.2023.113764>
- [35] Kazem Haki, Jannis Beese, Stephan Aier, and Robert Winter. 2020. The evolution of information systems architecture: An agent-based simulation model. *MIS Quarterly* 44, 1 (2020), 155–184.

- [36] Samer Hassan and Primavera De Filippi. 2021. Decentralized autonomous organization. *Internet Policy Review* 10, 2 (2021), 1–10. <https://doi.org/10.14763/2021.2.1556>
- [37] Jackie Kazil, David Masad, and Andrew Crooks. 2020. Utilizing Python for agent-based modeling: The mesa framework. In *Social, Cultural, and Behavioral Modeling*, Robert Thomson, Halil Bisgin, Christopher Dancy, Ayaz Hyder, and Muhammad Hussain (Eds.). Springer International Publishing, Cham, 308–317.
- [38] Niklas Konstantin Klein, Fritz Lattermann, and Dirk Schiereck. 2023. Investment in non-fungible tokens (NFTs): The return of Ethereum secondary market NFT sales. *Journal of Asset Management* 24 (May 2023), 241–254. <https://doi.org/10.1057/s41260-023-00316-1>
- [39] Thomas Lux. 2018. Estimation of agent-based models using sequential Monte Carlo methods. *Journal of Economic Dynamics and Control* 91 (2018), 391–408. <https://doi.org/10.1016/j.jedc.2018.01.021>
- [40] Charles M. Macal. 2016. Everything you need to know about agent-based modelling and simulation. *Journal of Simulation* 10, 2 (2016), 144–156. <https://doi.org/10.1080/14649350802481470>
- [41] Haim Mendelson. 1982. Market behavior in a clearing house. *Econometrica: Journal of the Econometric Society* 50, 6 (1982), 1505–1524. <https://doi.org/10.2307/1913393>
- [42] Loretta J. Mester. 2007. Some thoughts on the evolution of the banking system and the process of financial intermediation. *Economic Review-Federal Reserve Bank of Atlanta* 92, 1/2 (2007), 67.
- [43] John Stuart Mill. 1864. Chapter V—Of the connexion between justice and utility. In *Utilitarianism*. Cambridge University Press, 62–96. <https://doi.org/10.1017/CBO9781139923927.005>
- [44] John Stuart Mill. 1998. *On Liberty and Other Essays*. Oxford University Press.
- [45] Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. 2021. Mapping the NFT revolution: Market trends, trade networks, and visual features. *Scientific Reports* 11, 1 (2021), 20902. <https://doi.org/10.3389/fbloc.2019.00012>
- [46] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. Retrieved February 29, 2024 from <https://bitcoin.org/en/bitcoin-paper>
- [47] Luis Oliveira, Liudmila Zavolokina, Ingrid Bauer, and Gerhard Schwabe. 2018. To token or not to token: Tools for understanding blockchain tokens. In *Proceedings of the 39th International Conference of Information Systems*.
- [48] Vilfredo Pareto. 1964 [1919]. *Cours d'économie Politique*. Vol. 1. Librairie Droz. <https://doi.org/10.3917/droz.paret.1964.01>
- [49] Jack Parkin. 2019. The senatorial governance of Bitcoin: Making (de)centralized money. *Economy and Society* 48, 4 (Oct. 2019), 463–487. <https://doi.org/10.1080/03085147.2019.1678262>
- [50] Helen Partz. 2022. First Binance Soulbound Token BAB Targets KYC User Credentials. Retrieved February 29, 2024 from <https://cointelegraph.com/news/first-binance-soulbound-token-bab-targets-kyc-user-credentials>
- [51] Stephen Penzo and Niloufer Selvadurai. 2021. A hard fork in the road: Developing an effective regulatory framework for public blockchains. *Information & Communications Technology Law* 31, 2 (2021), 240–266.
- [52] Thomas Piketty. 2014. *Capital in the Twenty-First Century*. Belknap Press.
- [53] John Rawls. 1971. *A Theory of Justice*. Harvard University Press Boston. <https://doi.org/10.4159/9780674042605>
- [54] John Rawls. 1991. Justice as fairness: Political not metaphysical. In *Equality and Liberty*. Springer, 145–173. https://doi.org/10.1007/978-1-349-21763-2_10
- [55] Wessel Reijers and Mark Coeckelbergh. 2018. The blockchain as a narrative technology: Investigating the social ontology and normative configurations of cryptocurrencies. *Philosophy & Technology* 31, 1 (March 2018), 103–130. <https://doi.org/10.1007/s13347-016-0239-x>
- [56] Matteo G. Richiardi, Roberto Leombruni, Nicole J. Saam, and Michele Sonnessa. 2006. A common protocol for agent-based social simulation. *Journal of Artificial Societies and Social Simulation* 9 (2006), 1–20.
- [57] Ioanid Roşu and Fahad Saleh. 2021. Evolution of shares in a proof-of-stake cryptocurrency. *Management Science* 67, 2 (2021), 661–672. <https://doi.org/10.1287/mnsc.2019.3515>
- [58] Carlos Santana and Laura Albareda. 2022. Blockchain and the emergence of decentralized autonomous organizations (DAOs): An integrative model and research agenda. *Technological Forecasting and Social Change* 182 (2022), 121806. <https://doi.org/10.1016/j.techfore.2022.121806>
- [59] Fabian Schär. 2021. Decentralized finance: On blockchain-and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review* 103, 2 (2021).
- [60] Nathan Schneider. 2019. Decentralization: An incomplete ambition. *Journal of Cultural Economy* 12, 4 (July 2019), 265–285. <https://doi.org/10.1080/17530350.2019.1589553>
- [61] Jan Schwiderowski, Asger Balle Pedersen, and Roman Beck. 2023. Crypto tokens and token systems. *Information Systems Frontiers* 26 (March 2023), 319–332. <https://doi.org/10.1007/s10796-023-10382-w>
- [62] C. E. Shannon. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27, 3 (July 1948), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>

- [63] Laura Shin and Andre Cronje. 2020. Andre Cronje of Yearn Finance on YFI and the fair launch: ‘I’m Lazy.’ *Unchained Podcast* (July 2020). Retrieved February 29, 2024 from <https://unchainedpodcast.com/andre-cronje-of-yearn-finance-on-yfi-and-the-fair-launch-im-lazy/>
- [64] Ali Sunyaev, Niclas Kannengießer, Roman Beck, Horst Treiblmaier, Mary Lacity, Johann Kranz, Gilbert Fridgen, Ulli Spankowski, and André Luckow. 2021. Token economy. *Business & Information Systems Engineering* 63 (2021), 457–478.
- [65] Namrta Tanwar and Jawahar Thakur. 2023. Patient-centric soulbound NFT framework for electronic health record (EHR). *Journal of Engineering and Applied Science* 70, 1 (Dec. 2023), 1–11. <https://doi.org/10.1186/s44147-023-00205-9>
- [66] Horst Treiblmaier, Melanie Swan, Primavera De Filippi, Mary Lacity, Thomas Hardjono, and Henry Kim. 2021. What’s next in blockchain research? An identification of key topics using a multidisciplinary perspective. *ACM SIGMIS Database: The DATABASE for Advances in Information Systems* 52, 1 (2021), 27–52. <https://doi.org/10.1145/3447934.3447938>
- [67] Rowan van Pelt, Slinger Jansen, Djuri Baars, and Sietse Overbeek. 2021. Defining blockchain governance: A framework for analysis and comparison. *Information Systems Management* 38, 1 (2021), 21–41. <https://doi.org/10.1080/10580530.2020.1720046>
- [68] Andreas Wagner. 2012. The role of randomness in Darwinian evolution. *Philosophy of Science* 79, 1 (2012), 95–119.
- [69] E. Glen Weyl, Puja Ohlhaver, and Vitalik Buterin. 2022. Decentralized Society: Finding Web3’s Soul. Retrieved February 29, 2024 from <https://www.microsoft.com/en-us/research/publication/decentralized-society-finding-web3s-soul/>
- [70] Jackson Wood. 2022. The Crypto Fear and Greed Index, Explained. Retrieved February 29, 2024 from <https://www.coindesk.com/learn/the-crypto-fear-and-greed-index-explained/>
- [71] Dirk A. Zetzsche, Douglas W. Arner, and Ross P. Buckley. 2020. Decentralized finance. *Journal of Financial Regulation* 6, 2 (Sept. 2020), 172–203. <https://doi.org/10.1093/jfr/fjaa010>
- [72] Meng Zhang and Guy Gable. 2014. Rethinking the value of simulation methods in the information systems research field: A call for reconstructing contribution for a broader audience. In *Proceedings of the 35th International Conference on Information Systems*.
- [73] Rui Zhang, Rui Xue, and Ling Liu. 2019. Security and privacy on blockchain. *ACM Computing Surveys* 52, 3 (2019), 1–34.
- [74] Rafael Ziolkowski, Gianluca Miscione, and Gerhard Schwabe. 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 (2020), 316–348. <https://doi.org/10.1080/07421222.2020.1759974>

Received 5 August 2022; revised 17 October 2023; accepted 17 February 2024