Article:
Oliveira, F.S. and Jee, S.J. (Accepted: 2026) Tokenizing Creativity in Open Communities. European Journal of Operational Research. ISSN: 0377-2217 (In Press)

Tokenizing Creativity in Open Communities

Fernando S. Oliveira
Trade and Agriculture Directorate, OECD, France

Su Jung Jee *
Leeds University Business School, University of Leeds, LS2 9JT, UK
Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, OX1 3UQ, UK

Open communities, where individuals share and evaluate creative content, often face challenges in maintaining high-quality contributors over the long term. This study examines the effective governance of Decentralized Autonomous Organizations (DAOs), an emerging organizational paradigm that may address this issue. Built on blockchain technology, DAOs operate through decentralized and automated decision-making, using cryptographic tokens to reward contributors transparently. Assuming that transparent and collectively determined reward mechanisms can encourage participation, we develop an agent-based model with reinforcement learning to simulate the long-run implications of diverse DAO governance structures. Our models account for the cumulative process of reward acquisition and behavioral adaptation, shaping the evolution of such communities. We juxtapose three DAO scenarios, each characterized by its approach to distributing engagement rights that shape community evolution: Egalitarian (equitable rights distribution), Merit-based (rights based on reputation), and Property-based (rights contingent on investment). These models are compared with traditional open communities that lack reward mechanisms. Our findings underscore the potential of tokenizing creativity to sustain open communities. Notably, the Merit-based DAO exhibits high engagement rates and weak concentration of governance decisions compared to other DAO governance structures. Meanwhile, the Property-based DAO presents a viable alternative by fostering exploratory learning and inclusive reward distribution for creators, albeit with relatively low engagement.

Key words: Open Community; DAO; Rewards; Agent-based Simulation; Creativity

---

1. Introduction

Online open communities dedicated to creating and sharing digital content have thrived over the past few decades. For example, *Wikipedia* (wikipedia.org) relies on volunteers who collaborate to build an encyclopedia and *Thingiverse* (thingiverse.com) is sustained by individuals with a passion for creating and sharing 3D printing designs. Scholars have investigated a spectrum of motivations to comprehend the voluntary engagement of community participants (Von Krogh et al. 2012). Intriguingly, intrinsic

* Corresponding author: s.jee@leeds.ac.uk

motivations, such as deriving enjoyment and fulfillment, have been posited as a primary driving force propelling people's contributions to the community (Boudreau and Lakhani 2011).

Although open communities have gained popularity, many of them face a common challenge: 'maintaining a core group of high-quality contributors' (Khern-am nuai et al. 2018). Statistical evidence (e.g., Halfaker et al. 2013) indicates that even well-established communities, such as Wikipedia, experience a decline in the number of deeply engaged contributors over time, jeopardizing the community's long-term vibrancy. Given this difficulty, the sufficiency of relying on intrinsic motivation, which is considered the primary driving force behind community participation, to sustain contributions in the long run comes into question.

Some communities have experimented with reward mechanisms to boost engagement. Empirical evidence on their effectiveness has been mixed so far. Some scholars have argued for the existence of crowd-out effects, in which rewards decrease intrinsic motivation (Frey and Jegen 2001, Conti et al. 2023), while other studies have shown the effectiveness of properly recognizing and rewarding quality contributions (Khern-am nuai et al. 2018, Yu et al. 2022). The evidence so far indicates that the potential role and impact of rewards depend on how they are designed and managed within the community.

In this context, the blockchain-based Decentralized Autonomous Organization (DAO) has gained increasing attention as an emerging model with the potential to revolutionize the incentive system of open communities (e.g., Krishnan 2020, Zichichi et al. 2019). Leveraging the transparent and automated transaction capabilities blockchain offers via smart contracts, DAOs create purpose-driven tokens to reward participants for their contributions to the community. Through well-crafted token systems and a sound governance structure, DAOs can distribute rewards to community contributors in a transparent manner.

The transparent mechanism of DAOs allows all participants not only to see the reward rules recorded in the software code but also to collectively decide on reward governance within the community (e.g., how rewards are managed and distributed). Hence, incentive mechanisms are shaped by community participants, a distinctive feature compared with communities in which rewards are centrally managed by community owners. A pertinent real-world instance of a DAO is *steemit* (steemit.com), where participants receive rewards for their contributions to content creation, sharing, and rating, all tied to cryptographic tokens. Since DAOs for open communities are still in their infancy, numerous aspects regarding governance and token designs remain unclear (Lumineau et al. 2021).

Different reward methodologies influence participant behavior in distinct manners, collectively shaping the trajectory of the community's development. Given this context, the central research question addressed in this study is as follows: *How can we design a DAO governance structure capable* *of guiding a community toward achieving its shared objectives and ensuring the sustained engagement of contributors?*

Using agent-based simulation (e.g., Axelrod 1997, Rosa et al. 2019, Oliveira 2023, Ponta et al. 2023, Nguyen et al. 2024), which is well-suited for studying emerging phenomena, we investigate the evolution of communities where individuals engage in content production and rating under various DAO governance structures. In the proposed governance structures, participants receive rewards through cryptographic tokens for their contributions to content production and rating. Based on the existing evidence on the positive role of transparent rewards in motivating people (e.g. Gutierrez et al. 2025) and the potential effectiveness of rewards in certain open community settings (e.g., Yu et al. 2022), we assume that the transparent and collectively shaped reward mechanism enabled by DAO can encourage the participants' contribution.

Based on this assumption, the DAO governance structures we investigate are characterized by how they allocate the essential rights for community engagement, which in turn influences the quality and activity levels of contributions. We explore three distinct DAO governance structures: 1) Egalitarian, 2) Merit-based, and 3) Property-based. These structures distribute rights equally among all participants, in proportion to reputation based on previous contributions, and dollar investments in the community, respectively. We compare the community's evolution under DAO scenarios with that of the existing open community scenario (referred to as *Open Community* in our model; Section 4), which is used as a reference point to understand the role of DAO governance structures. Our model can be adapted to test DAO governance structures beyond the scenarios addressed in this study.

Our investigation unveiled the potential of tokenizing creativity for open communities. In particular, we show that the Merit-based DAO achieves the highest level of overall community engagement. In addition, the Property-based DAO maintains a high learning rate and inclusiveness in rewarding content producers, although the overall engagement rate was relatively low. The high learning rate is essential in the creativity-driven community because it helps the community maintain content (or design) diversity. In our simulations, the Egalitarian DAO did not outperform other governance structures regarding overall engagement levels and design development.

Our paper has three key contributions. First, we extend the existing literature on the diverse motivations of contributors in creativity-driven open communities (Von Krogh et al. 2012). Given the current open communities' struggles in retaining high-quality contributors, we showcase the potential promise of a cryptographic token-based incentive mechanism to address this issue. Our results initiate a thought-provoking future discourse regarding the incentive mechanisms within open community settings.

Second, we propose an agent-based model (ABM) to experiment with promising DAO governance. Given that blockchain is considered a fundamental technology shaping a new social and economic system (Iansiti et al. 2017), there exists a demand for innovative tools to aid in comprehending the future driven by this technology. Agent-based models allow the study of emergent behavior (e.g., Axelrod 1997, Sawyer 2001), being ideal in the study of decentralized task coordination in complex systems (Gudmundsson et al. 2023). Building upon the mechanisms within existing open communities, we introduce diverse governance structures that can be incorporated into DAO settings. The proposed model can be utilized by firms, entrepreneurs, and researchers intrigued by the potential of introducing a new DAO to foster creative works. The model can be extended to conduct experiments with various tokens and governance designs to find an adequate DAO structure aligned with their objectives.

Third, as a methodological contribution, the article uses reinforcement learning (RL) in an ABM to simulate emergent reputations in the DAO, considering two types of agents (authors and raters) with individual features. While reinforcement learning has been integrated into ABM in other domains, to our knowledge, this is one of the first applications of an ABM plus RL approach to DAO governance in creative communities. This novel integration builds on prior work by enabling agents to adapt based on past outcomes, capturing dynamic reputation effects that static ABM frameworks might miss. This technique allows comparison of different organizational structures by describing how individual participants learn and how the community evolves. The inclusion of RL is justifiable from a technical perspective, as it converges to the optimal policy given a sufficiently large number of learning opportunities (e.g., Sutton and Barto 1998). It has been used, among other topics, to study several important social science problems, including auction theory (Bandyopadhyay et al. 2008), investment in oligopolistic markets (Bunn and Oliveira 2008, Oliveira and Costa 2018), equilibria selection in game theory (Daskalova and Vriend 2021) and for policy optimizations in dynamic scheduling methods for autonomous factories (Didden et al. 2024).

To motivate the article from a real-world open community, we commence in Section 2 by offering descriptive empirical insights from *Thingiverse*. Section 3 delves into existing research concerning the motivations of open community participants, elucidating the potential for DAOs to contribute to the thriving and long-term sustainability of creativity-driven open communities. Sections 4 and 5 introduce an agent-based model crafted to simulate the ramifications of varied DAO governance mechanisms on community evolution. With the outcomes of simulation outlined in Section 6, we conclude in Section 7 by discussing ways in which DAOs can revolutionize traditional open community governance via integrating incentive mechanisms inspired by the token economy.

2. Description of an Existing Open Community: A 3D Printing Case

To ground our study in real-world open communities, we begin by examining Thingiverse. Thingiverse was established by MakerBot, a 3D printer company, as part of its strategic initiative to expand the consumer 3D printing market. The company aimed to catalyze this expansion by fostering an online open community where people could freely exchange 3D printable design files (West and Kuk 2016). Although MakerBot's commercial interests drove the community's inception, Thingiverse has evolved into a thriving open community where numerous individuals enthusiastically share their designs and evaluate the creations of other participants.

Creators on Thingiverse engage with other contributors' designs by clicking the Like button, leaving comments, or sharing pictures of physically printed outputs crafted from the provided design files. The platform encompasses eleven broad categories - 3D Printing, Art, Fashion, Gadgets, Hobby, Household, Learning, Models, Tools, Toys and Games, and Other - into which individuals can upload their design files and assess the work of their peers. These categories encompass diverse attributes: some, like 3D Printing or Tools, showcase strong engineering characteristics, while others, such as Fashion or Art, highlight aesthetic qualities. In addition, many categories exhibit a blend of both aesthetic and functional features. The category information, which indicates the domain in which a participant produces and appreciates 3D printing designs, can reveal distinctive facets of a participant's creative and evaluative activities.

Between November 2008 (establishment of the platform) and August 2022, about 5 million 3D printing design files were uploaded to the community. We randomly sample 5,000 distinct design-specific identification codes (about 0.1% of all existing identification codes), then web-crawl the activity history of the author of each design. The main information we collect is the category of designs each participant has produced and liked because it demonstrates the participant's unique style of production and rating, respectively. Several design files are removed from the platform when we access them through the design-specific identification code. In addition, collected author information can be partly redundant because each author can produce multiple items. As a result, we obtain the engagement history of 1,532 unique participants. In reality, many people in the open communities are relatively passive participants who use others' shared work without sharing their original work. Therefore, the 1,532 participants we observed can be considered active participants because they uploaded at least one design file.

Figure 1 depicts the community activities of participants in our sample from various perspectives. Figures 1(a) and (b) show the distribution of authors based on their popularity measured by the number of followers and the average number of Likes a produced design receives, respectively. Figures 1(c) and (d) demonstrate the diversity-wise contribution of each creator through the number of distinct design categories to production and rating. Participants who have engaged in multiple categories have expanded their range of participation over time, either through developing designs or interests. Figures 1(e) and (f) show the distribution of quantity-wise contribution in production

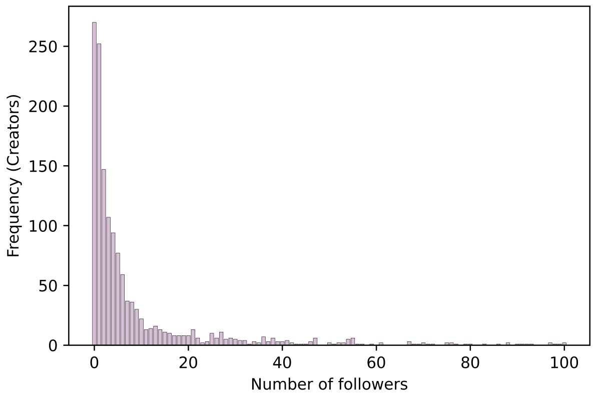

(a) Number of followers

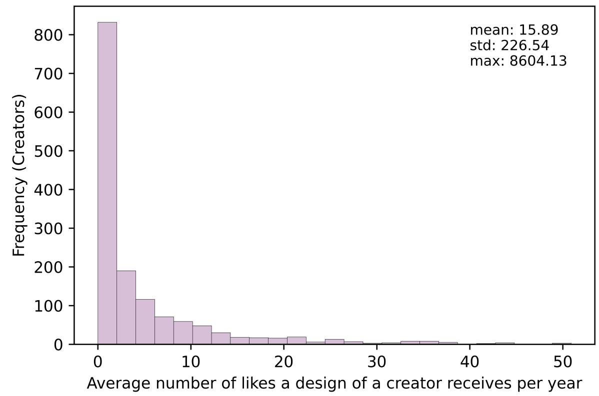

(b) Average likes a design receives (per year)

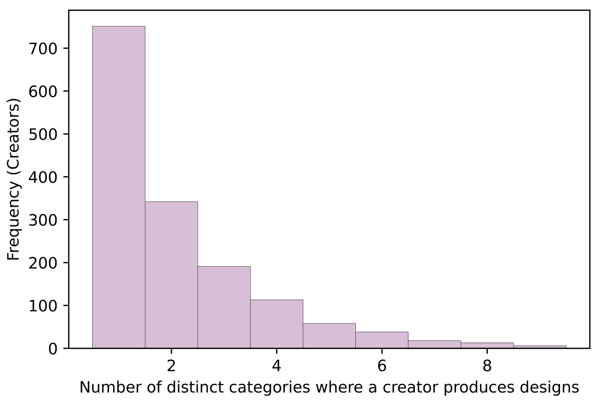

(c) Number of categories (produced)

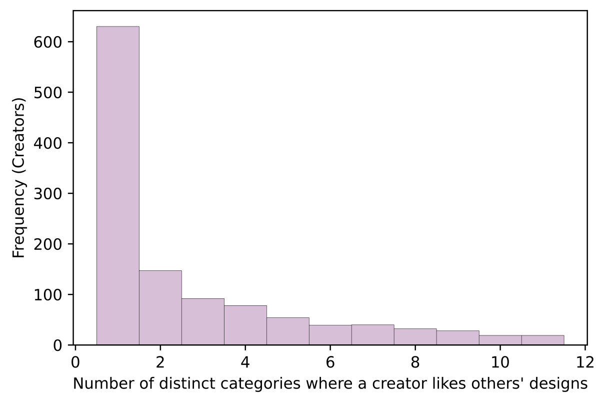

(d) Number of categories (liked)

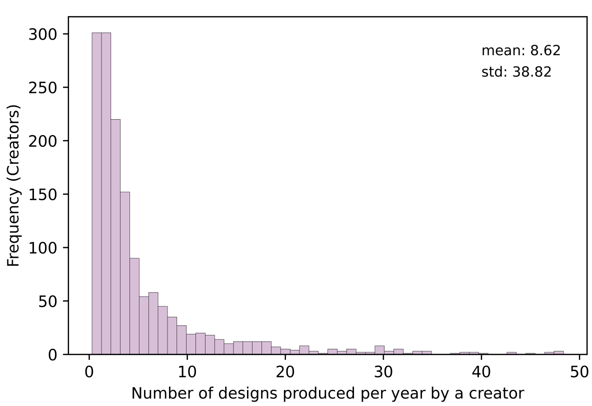

(e) Number of produced designs (per year)

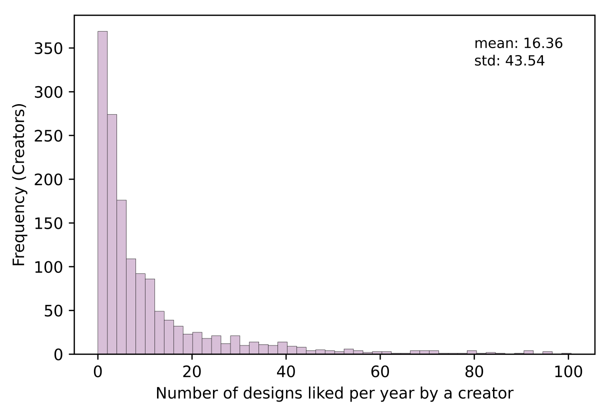

(f) Number of liked designs (per year)

Figure 1 Activity of Creators in Our Sample (n = 1,532)

Notes. (a) and (b) exhibit the distribution of authors based on the popularity of production proxied by the number of followers and the average number of Likes an author's design receives in one year, respectively. Values in the x-axis of (b) are calculated by considering the age of each design. Figures (c) and (d) present the distribution of participants based on the number of categories they contribute regarding production and rating, respectively. (e) and (f) show the distribution of participants based on the quantity of their production and rating activities, respectively. Values in the x-axis of (e) and (f) are calculated by considering the period of each participant's community engagement.

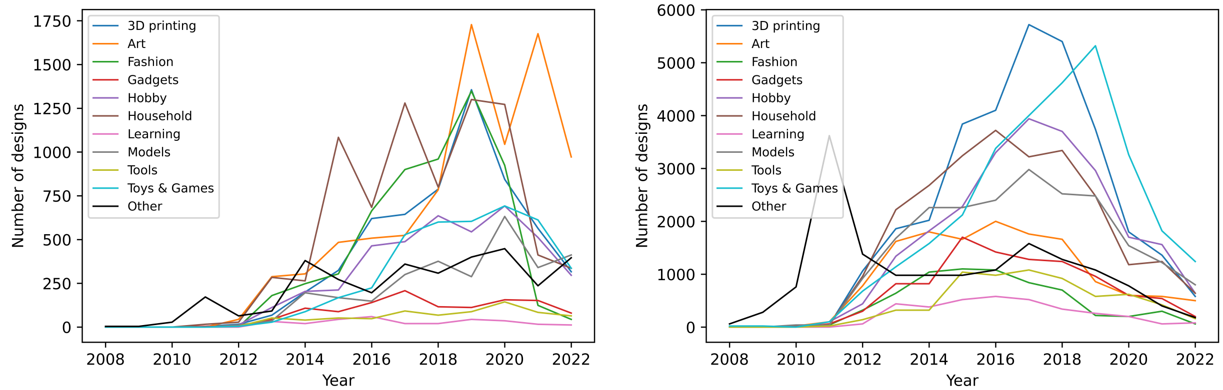

(a) Produced (b) Liked

Figure 2 Designs (a) Produced and (b) Liked by Creators in Our Sample (n = 1,532)

Notes. Year 2022 covers the data by mid-August. Engagement in both production (left) and evaluation (right) across all content categories shows a decreasing trend ten years after the community's establishment.

and rating, respectively. Overall, these figures show that participants' contributions follow a power-law distribution in which a small proportion of participants dominate the contribution while others contribute in a minor way. This aligns with the phenomenon observed in many other online open communities (e.g., Faraj and Johnson 2011, Johnson et al. 2014).

Figure 2 shows each category's aggregate number of produced and liked designs over time. The *Other* category was the most popular in the early phase of community development around 2010-11. This might be due to the ambiguity of the community's identity in this emerging period. Over time, other categories have emerged. Interestingly, the categories where people produce the most are aesthetic ones, such as Art and Fashion, while those where people like the most include more functional ones, such as 3D printing. This figure shows that the community experienced significant growth for the first ten years but recently showed a strikingly decreasing engagement of participants. This aligns well with the phenomenon observed in other open communities (Halfaker et al. 2013, Khern-am nuai et al. 2018).

There can be several reasons behind this recent decrease. One possible explanation is the lack of proper incentive mechanisms consistently motivating participants to contribute to the community. Building on our observations from the current open community for 3D printing, in Section 3, we review previous studies on the various motivations of participants who voluntarily contribute to online open communities and projects. Our review reveals the limitations of the current understanding in explaining and addressing the decline in participation in many open communities. We discuss why DAOs can be an alternative tool to address this issue.

3. Literature Review

3.1. Motivating Mechanisms in Open Communities

The online open communities dedicated to sharing knowledge and creative content has grown substantially in recent decades. These communities establish virtual realms wherein individuals with shared interests can engage in mutual interaction by exchanging and synthesizing their ideas and expressions (Blanchard and Markus 2004).

The relevant literature has discussed three distinct motivations of the community participants: intrinsic, extrinsic, and internalized extrinsic motivations (see the comprehensive review by Von Krogh et al. 2012). Intrinsic motivation springs from participants' inherent drives, such as enjoyment or self-efficacy, while extrinsic motivation emanates from rewards, encompassing financial compensation and opportunities for career advancement. Internalized extrinsic motivation, exemplified by peer recognition or reciprocity, derives from external origins but is perceived by contributors as an internal impetus. These motivational aspects can coexist and interact (Roberts et al. 2006).

In the early days of creative open communities, most communities relied on the voluntary participation of members who are largely driven by intrinsic motivation. This is indeed well in line with the findings from studies that showed that the primary driving force behind the contribution to the open community is intrinsic motivation (Boudreau et al. 2011). However, recent data reveals a decline in participation within open communities, indicating that the existing incentive structures might fall short of retaining long-term involvement (Halfaker et al. 2013). To enhance the engagement, some open communities have introduced features such as *likes* and *badges*, capable of fostering internalized extrinsic motivation by aiding in self-identity validation.

Recently, several open communities have experimented with rewards to understand whether extrinsic motivations can be relevant in the long run. Rewards can be distributed among participants based on their level of engagement, known as *completion-contingent* incentives used by, for example, Bestbuy.com. Alternatively, the rewards can be allocated proportionally to the quality of each participant's contribution. These are known as *performance-contingent* incentives. They are used, for example, by Epinions.com. Such incentives are grounded in the fundamental logic in economics that raising monetary incentives increases supply.

Meanwhile, other studies have argued for a potential crowding-out effect in which the reward may instead diminish intrinsic motivation (Frey and Jegen 2001). Empirical evidence in this area has so far been mixed and context-dependent. Some studies support the crowding-out effect (Conti et al. 2023), while others show that rewards, particularly those contingent on performance, can enhance both the quantity and quality of contributions by community participants (Yu et al. 2022). Recent research (e.g., Khern-am nuai et al. 2018) indicates that the *completion-contingent* approach leads to an uptick in the quantity of engagement. Yu et al. (2022) demonstrates that the *performance-contingent* reward system can result in heightened quantity and quality of the contributions. Therefore, so far, prior research suggests a nuanced role of reward mechanisms in open communities, shaped by the nature of the work and the specific design of incentives.

In this context, DAOs have garnered increasing attention as a means to transform existing open communities with a cryptographic token-based incentive system. In Section 3.2, we discuss how DAOs can innovate the existing open communities through a transparent and decentralized rewarding system.

3.2. DAO as an Innovative Tool

Blockchain was first designed by Satoshi Nakamoto in 2008 (see Nakamoto 2008) to innovate traditional financial models, which have relied on third-party intermediaries such as banks, by decentralizing the transactions. The blockchain chronologically records all the transactions, and everyone can look into the transaction histories. By doing so, it removes intermediaries such as banks, resolves the double spending problem (Chohan 2017), reduces transaction costs, and eliminates the risk that intermediaries' database is being forged, destroyed, or hacked by others.

Specifically, the Ethereum blockchain has allowed the emergence of a new organizational paradigm, the DAO, based on *smart contracts* (e.g., Antonopoulos and Wood 2018). Once designed, a smart contract is deployed to the Ethereum blockchain network, and different programs on the Ethereum network can interact with each other by calling specific functions, including sending and receiving cryptocurrencies. Transactions through smart contracts do not require intermediaries and executed automatically as programmed when specific criteria are met. The transactions are transparently recorded on the blockchain, allowing anyone to observe the details of each transaction.

These smart contracts are the fabric of DAO, a decentralized organization where like-minded people work together to achieve community-driven goals that benefit every participant. The DAO makes autonomous decisions based on the codified rules governing the organization. The essential concept for the development of the DAO is the empowerment of creative content producers in the digital environment (e.g., Krishnan 2020), such as the deployment of democratized social news, creative content production and evaluation, and crowdfunding (Zichichi et al. 2019).

This empowerment of the content producers is driven by rewarding people using a *token*, i.e., a smart contract that defines the rights assigned to someone who holds the token. For example, the rights can include access rights to a particular product or service, depending on the goal of the DAO. Or, the rights can reflect within-community influence in the community-level decisions. The tokens can be purposefully designed to incentivize an autonomous group of people to contribute to the community-level goal (Voshmgir 2020).

DAOs are already used for various purposes, including financial applications operating without intermediaries and creativity-driven communities. Specifically, *steemit* is an application of DAO for creative content production where participants share their content and evaluate others' work. *steemit* motivates participants to contribute to production and evaluation using different types of purpose-driven tokens, including a token for monetary rewards or within-community power rewards.

The creativity-driven DAOs share key characteristics with existing open communities. However, the main difference lies in their potential to sustain themselves by autonomously rewarding creators and evaluators through the innovative use of tokens and system transparency. Transparency here means not only that reward rules are open to all participants but also that the rules are collectively determined by community members. The fact that rewards are not simply imposed by a third party managing the community in a centralized manner, but are instead collectively decided by creators and evaluators, grants more authority to community participants. The transfer of authority in reward governance could, to some extent, mitigate concerns about the crowding-out effect.

Through adequately designed governance and purpose-driven tokens, DAOs can reward participants who make high-quality contributions and every participant can transparently see and collectively decide how the reward system is managed. The economic value of specially designed tokens increases as more participants acknowledge the value of creative activity within the community, as demonstrated by the increasing (or robust) level of community engagement. The explicit and transparent reward system can change people's behavior, shaping the direction of community evolution in the long run.

However, little is known about the effective governance structure of DAO that can be used to maintain people's creative contributions, including content production and evaluation. Simulation is an appropriate approach, enabling us to envision potential scenarios based on different DAO designs. The agent-based model has been previously employed to simulate the roles of smart contracts: for example, Diogo et al. (2018) and Rosa et al. (2019) analyze the impact of consensus mechanisms in smart contracts on individual behavior. We extend the previous approach by modeling interactions between content producers and raters engaging in open communities, thereby informing future designs of DAOs for creative content production. Section 4 elaborates on our agent-based model to experiment with various DAO governance structures.

4. Modeling an Open Community and the DAO

There are several reasons why the agent-based model has been chosen to model distributed decision problems related to DAOs and blockchain (e.g., Rosa et al. 2019): it can be used to model the detailed structure of the interactions between decisions of independent decision makers in the presence of non-linear and complex profit functions, in which people and organizations are not able to optimize their behaviors and instead learn by trial-and-error, for example using reinforcement learning (e.g., Oliveira 2023, Sutton and Barto 1998).

Table 1 Key Concepts Used in the Agent-Based Model

	Definition
Rating Variables	Variables used to define the rating behavior of a rater j, including the preference for products (w_j), the threshold for making rating decision (a_j), the probability of engaging in rating (P_j^r), and the number of products a rater j scans for the rating (Q_j^r).
Production Variables	Variables used to define the production behavior of an author i, including the radius to search for new designs (β_i), the probability of engaging in production improvement (P_i^p), and the number of rated works that i scans to improve its production (Q_i^p).
Rating	Rating awarded by rater j to author i, in the range 0 to 1 (r_{ji})
Design	Design produced by author i (s_i)
Rating Reputation	Reputation (ϕ_j) a rater j built based on its previous rating engagement; Computed by considering the extent to which j gives high ratings to the reputed authors
Production Reputation	Reputation (η_i) an author i built based on its previous production engagement; Computed by considering the ratings i received; *Product Visibility* (see Figure 3), which is the extent to which i's product is exposed to the raters, increases in proportion to η_i
Rating Reward (tokens)	Tokens given to raters as rewards for engagement in rating; Distributed in proportion to the Rating Reputation; This reward changes the Rating Variables of an agent in the next iteration.
Production Reward (tokens)	Tokens given to authors as rewards for engagement in production; Distributed in proportion to the Production Reputation; This reward changes the Production Variables of an agent in the next iteration.
Rating Rights	Rights to engage in the community through rating; Equivalent to *tickets to earn rewards* through community engagement as a rater; Distributed based on the governance rules of DAO
Voting Rights	Rights to influence decisions related to the DAO governance (e.g., choice among Egalitarian, Merit-based, and Property-based DAOs); Distributed based on the governance rules of DAO

Note. Rating Reputation, Rating Reward, Production Reward, and Voting Rights are not in the Open Community but are unique features of the DAO (see Figures 3 and 4).

Let us start by summarizing the concepts used in this section in Table 1. Let us consider a community comprising I members engaged in creating, advancing, and evaluating creative content. While general participants can function as both authors and raters, they also possess the flexibility to opt for specialization in either of these roles.

On the one hand, the Open Community is, by definition, online open space, lacking a formal framework for incentivizing participation. That is, it emulates the conventional form of community that largely relies on the intrinsic motivation of participants. On the other hand, the DAO builds upon the foundational concepts of the Open Community by incorporating methods for rewarding contributions to the community. Tokens are used to incentivize quality production and rating of creative content. A key assumption here is that the transparent and fair reward system of a DAO can change the behavior of participants to improve their contributions. Hence, DAO relies on not only intrinsic motivation but also extrinsic motivation of participants.

Many concepts and modeling techniques are shared between these two organizational structures: the Open Community and the DAO. Section 4.1 outlines the modeling of the Open Community and Section 4.2 leverages this foundation to expound upon modeling distinct types of DAOs in detail.

To avoid confusion, we clarify our terminology for model components. In Table 2, we categorize the key elements of the model: parameters refer to fixed exogenous inputs (e.g., population size I, decay rate δ) or stochastic parameters (e.g., μ and ν guiding local search), state variables are agent attributes that can evolve (e.g., each author's search radius β_i), and other variables computed at each iteration denote quantities calculated from agents' states and decisions (e.g., an agent's Rating Rights l_j or contribution score ξ_j). We have indicated the possible value ranges for each item in Table A1 and noted which elements are fixed versus dynamically updated during the simulation. The terminology aligns with standard ABM practice (e.g. Grimm et al. 2020), where parameters represent exogenous inputs, state variables capture agent attributes, and calculated variables reflect endogenous outcomes computed at each iteration.

4.1. Modeling an Open Community

Within the Open Community, every author's work can be evaluated by other participants. A concise overview of the Open Community's organizational structure is illustrated in Figure 3. Figure 3 illustrates the interaction between an author i and a rater j. A defining attribute of the Open Community is the equitable distribution of Rating Rights, which denote an individual's allowed capacity to engage within the community by rating the works of various authors.

The rater is characterized by Rating Variables, which underlie their Rating Decisions. The author's characterization encompasses a set of Production Variables, Design Decision, and Production Reputation (see Table 1). The rater observes products showcased within the community space and assigns ratings to the products they find preferable from various authors. These ratings serve as the foundation for the author's cultivation of Production Reputation and development of Designs, enabling the refinement of product attributes to align with the community's preferences. The Production Reputation is pivotal in shaping the author's prominence within the community (Product Visibility). A higher Production Reputation increases the likelihood of receiving ratings for their products. Consequently, a feedback loop emerges encompassing Rating, Production Reputation, Design, and Product Visibility. When this loop operates positively, it guides the author in honing their designs through iterative revisions.

Table 2 Indexes, Parameters, and Variables

Indexes
i, j, k: stand for a community member, either an author or a rater
t: iteration index from 1 to T
Parameters
Fixed
T: number of iterations per episode of the simulation
I: number of community members
δ: idiosyncratic rating decay rate
σ: weight used to compute Voting Rights distribution considering contributions
Stochastic
μ and ν: guide exploration of authors; follow a uniform distribution, i.e., $\mu, \nu \sim U[0,1]$
State variables
β_i: radius used by the author i to search for new designs
P_i^p: probability that author i explores new product types
Q_i^p: number of products scanned by author i while exploration
$w_j \in [-1,1]$: idiosyncratic preference of rater j regarding design features (static)
a_j: rater j's idiosyncratic rating threshold above which j supports the author (static)
P_j^r: probability that rater j participates in rating
Q_j^r: number of products rated by rater j
Decision variables
s_i: design of author i, in the range $[-1,1]$
r_{ji}: rating j awards to the work of author i
l_{ji}: represents a *like* and it is equal to 1 if $r_{ji} > 0$ and 0, otherwise
Other calculated variables
l_j: Rating Rights possessed by j; exogenously fixed in the Egalitarian DAO, and endogenously updated in Merit-based and Property-based DAOs
ξ_j: how much j has supported the authors with good Production Reputation
η_i: Production Reputation of author i; Production Rewards are given to i in proportion to η_i
ϕ_j: Rating Reputation of rater j; Rating Rewards are given to j in proportion to ϕ_j

4.1.1. Modeling the Rater. Rating is an indispensable component of the community. Our analysis commences by delving into the process of rating and feedback.

Let the variable $s_i \in [-1,1]$ represent the design type the author i uses in producing an item. These designs undergo evolution as authors strive to garner greater appreciation from the raters. Each author can adjust their design set over time to optimize the received ratings. Let $w_j \in [-1,1]$ stand for the preferences of rater j regarding design features. For concreteness, we can interpret the design variable s_i as representing a spectrum of creative styles. For example, in a 3D-printing community $s_i = -1$ might correspond to a highly utilitarian design approach, whereas $s_i = 1$ denotes

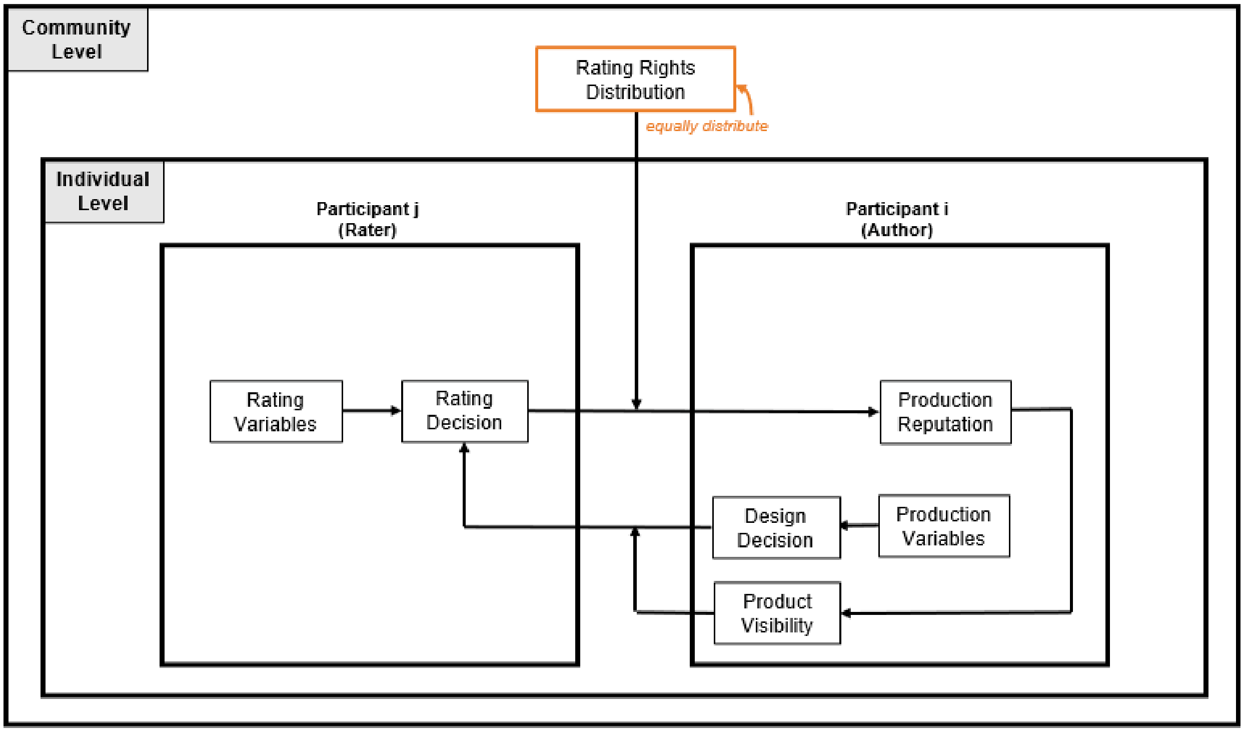

Figure 3 Open Community Governance

Note. Each participant can engage in the community as an author and a rater.

a highly experimental or aesthetic design. Each rater's preference w_j lies on the same spectrum, so rater j will only give a positive rating if an author's design is sufficiently close to w_j.

The r_{ji} denotes the rating that rater j assigns to the work of author i. The rating r_{ji} increases when s_i (design produced by the author i) and w_j (rater j's preferences) are congruent. Distinct ratings originate from the extent of proximity between the designs and the raters' preferences. The rating scheme adheres to a scale ranging from 0 to 1, with 1 being the highest and 0 the lowest.

We posit that the rating r_{ji} corresponds to the best work of i from j's standpoint. The rating is computed using equations (1): the community members only provide ratings for designs they possess a *strong* preference for. Each rater j possesses a unique threshold denoted as a_j. When the rating surpasses a_j, rater j extends support to the author and offers a *like* (l_{ji}); otherwise, no support is expressed.

$$r_{ji} = \begin{cases} \frac{2-\delta(s_i-w_j)^2}{2}, & \text{if } \frac{2-\delta(s_i-w_j)^2}{2} \geq a_j \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

The parameter $\delta > 0$ embodies the decay rate, signifying how swiftly r_{ji} diminishes as s_i deviates from w_j. When δ approximates zero, the community displays heightened tolerance to disparities, resulting in relatively elevated ratings even when s_i substantially differs from w_j.

Two additional Rating Variables define the rater's behavior, 1) P_j^r : the likelihood that a rater j engages in rating and 2) Q_j^r : the quantity of products scanned by rater j while arriving at a rating decision. A higher rating probability coupled with a larger sample size contributes to increased activity of rater j. In the Open Community model, these behavioral variables are considered exogenous.

Each rater j possesses an identical allocation of Rating Rights (l_j) within the Open Community: everyone can rate an author through a *like* expression. Since the Open Community restricts the rating decision to binary options (either 'like' or 'not'), let l_{ji} be equal to 1 if $r_{ji} > 0$, and 0 otherwise.

4.1.2. Modeling the Author. As depicted in Figure 3, an author's behavior is defined by the respective Production Reputation, a collection of Production Variables, and Design Decision. The raters continually assess each author's production. For each author i, the pursuit involves optimizing the ratings received from raters by adapting their designs s_i, which are dynamic and subject to learning over time.

The learning process involves a local search and selection phase, during which each author refines their design set to enhance the anticipated ratings received from others. The learning process for author i unfolds across three stages: 1) locally exploring potential directions and degrees of design adjustment; 2) discerning the optimal direction and extent of design adjustment by evaluating the projected impact on the rankings by each rater and the ensuing effect on author i's overall rating; 3) making a decision whether or not to adjust the design.

The Production Variables encompass β_i : the radius utilized by the author i when exploring new designs during the learning process; P_i^p : the probability that author i explores a new product type for improvement during a given iteration; and Q_i^p : the number of rated works scanned by i to acquire knowledge about production improvement (Table 1). All these variables are considered exogenous within the Open Community model.

The learning grounded in local search is delineated in equations (2a) and (2b), where μ and ν follow a uniform distribution, i.e., $\mu, \nu \sim U[0,1]$: these serve as stochastic exploration parameters guiding the author's local search within a specified radius $0 \leq \beta_i \leq 1$. Typically, this parameter is set to a small value (e.g., 0.1) to reflect the localized nature of the search and the cost associated with design changes. Equation (2a) embodies a search in the positive direction, signifying the author's assessment of adjusting designs closer to 1. Conversely, equation (2b) embodies a search for design enhancement in the negative direction, nearing -1.

When updating the design, an author i assembles a sample of raters and computes their ratings for s_i^- (the design adjusted negatively), s_i (the current design), and s_i^+ (the design adjusted positively). Subsequently, the author opts to pursue the design level yielding the highest rank. Therefore, a larger sample size translates to author i placing greater reliance on accurate information, thus facilitating a more informed revision of their designs.

$$s_i^+ = \min\left(s_i(1+\beta_i\mu), 1\right) \tag{2a}$$

$$s_i^- = \max\left(s_i(1-\beta_i\nu), -1\right) \tag{2b}$$

The final facet of the author's model is Product Visibility, which escalates in tandem with the author's Production Reputation. For instance, if an author possesses a Production Reputation of 5%, their products become visible to 5% of the members perusing through products. Essential to this dynamic is determining the probability of a rater evaluating a given product: the higher the visibility, the greater the likelihood that the author garners ratings.

4.1.3. Modeling the Feedback Process. All local design modifications are enacted after the authors' adjustments to their design sets. This updating process considers the concurrent design adjustments for a particular author i and the design alterations made by all other authors. Subsequently, the authors' revised ratings are computed.

The greater the number of likes received by author i, the higher their Production Reputation becomes. The authors possessing an elevated Production Reputation are more likely to receive ratings from raters. In the case of author i, the Production Reputation (η_i) is calculated using equation (3), which is based on the reinforcement learning algorithm (e.g., Sutton and Barto 1998). We have omitted the t index from the equation to simplify the notation. Instead, the symbol := denotes an iterative updating process, whereby the value of a variable at time $t+1$ is derived from the value of the variables at time t.

$$\eta_i := \frac{t}{t+1}\eta_i + \frac{1}{t+1}\frac{l_{.i}}{\sum_k l_{.k}} \tag{3}$$

Equation (3) delineates the update of Production Reputation (η_i) by author i. Here, $\frac{l_{.i}}{\sum_k l_{.k}}$ denotes the quantity of likes author i has received as a proportion of all likes accumulated by every author within the community during a specific iteration. The Production Reputation materializes as a cumulative outcome, persistently aligning with raters' expectations. This reputation evolves progressively over time, iteratively refining itself. While updating η_i, we allocate a weight of $\frac{t}{t+1}$ to the preceding η_i and a weight of $\frac{1}{t+1}$ to new rating information (i.e., $\frac{l_{.i}}{\sum_k l_{.k}}$). This reflects the reinforcement learning process, in which decisions made during the initial stages (e.g. when an author begins crafting 3D printing designs) hold greater relative significance than decisions (e.g. design set adjustments) made in later phases (Sutton and Barto 1998). During the early phases, the author's identity might be relatively obscure, and a subpar product could impede the author's chances of gaining visibility and traction among the public. Subsequently, as the author establishes a reputation for a specific product quality, a sub-optimal creation affects the reputation, albeit to a lesser degree.

4.2. Modeling a DAO

DAO can be perceived as an extension of the Open Community, imbued with additional structural elements. Substantial disparities emerge between the conventional Open Community and DAO due to the reliance of DAO on a token-based incentives framework affecting authors' and raters' behavior. Our proposition revolves around subjecting each member's production and rating involvement to evaluation through a smart contract, bestowing rewards tokens. The presumed impact of reward extends to influencing members' Production and Rating Variables (i.e., behaviors) by compensating them for their contributions (see Section 4.2.4 for insights into the modulation of behaviors resulting from tokens).

Figure 4 delineates the interactions between a rater j and an author i within a DAO. Noticeable shifts from the Open Community paradigm are highlighted in blue. At the same time, DAO governance mechanisms are illustrated through orange lines. Rating Rights are rights to engage in the community by rating others' work (Section 4.2.1) and Voting Rights are the rights to participate in the community-level governance setting decision (Section 4.2.5). In real-world DAOs, these rights can be referred to as another type of purpose-driven tokens. In our model, we call them *rights* instead of *tokens*, to avoid confusion with reward tokens that change the participants behavior (Sections 4.2.2 and 4.2.3).

4.2.1. Rating Rights Distribution. In the DAO, Rating Rights are equivalent to tickets to earn rewards through community participation as a rater, forming the foundation of interactions between raters and authors. While Rating Rights are distributed uniformly within the Open Community, we have characterized three distinct types of DAOs based on their allocation of Rating Rights (see orange lines within Community Level in Figure 4):

1) Egalitarian DAO: Each member is granted an equal number of Rating Rights, irrespective of their activity level.

2) Merit-based DAO: The DAO allocates Rating Rights to participants in proportion to their current Rating Reputation.

3) Property-based DAO: The DAO allocates Rating Rights to participants in proportion to the capital distribution within the DAO (i.e. participants can acquire Rating Rights through capital investment).

The notation l_j denotes the Rating Rights possessed by rater j. In the Egalitarian DAO, $l_j = 1/I \times 100$, where I signifies the total number of DAO participants. Even if not actively engaged in rating, every participant enjoys an equal number of Rating Rights. Within the Merit-based DAO, l_j is proportional to j's Rating Reputation. Meanwhile, l_j is proportional to j's capital investment in the Property-based DAO.

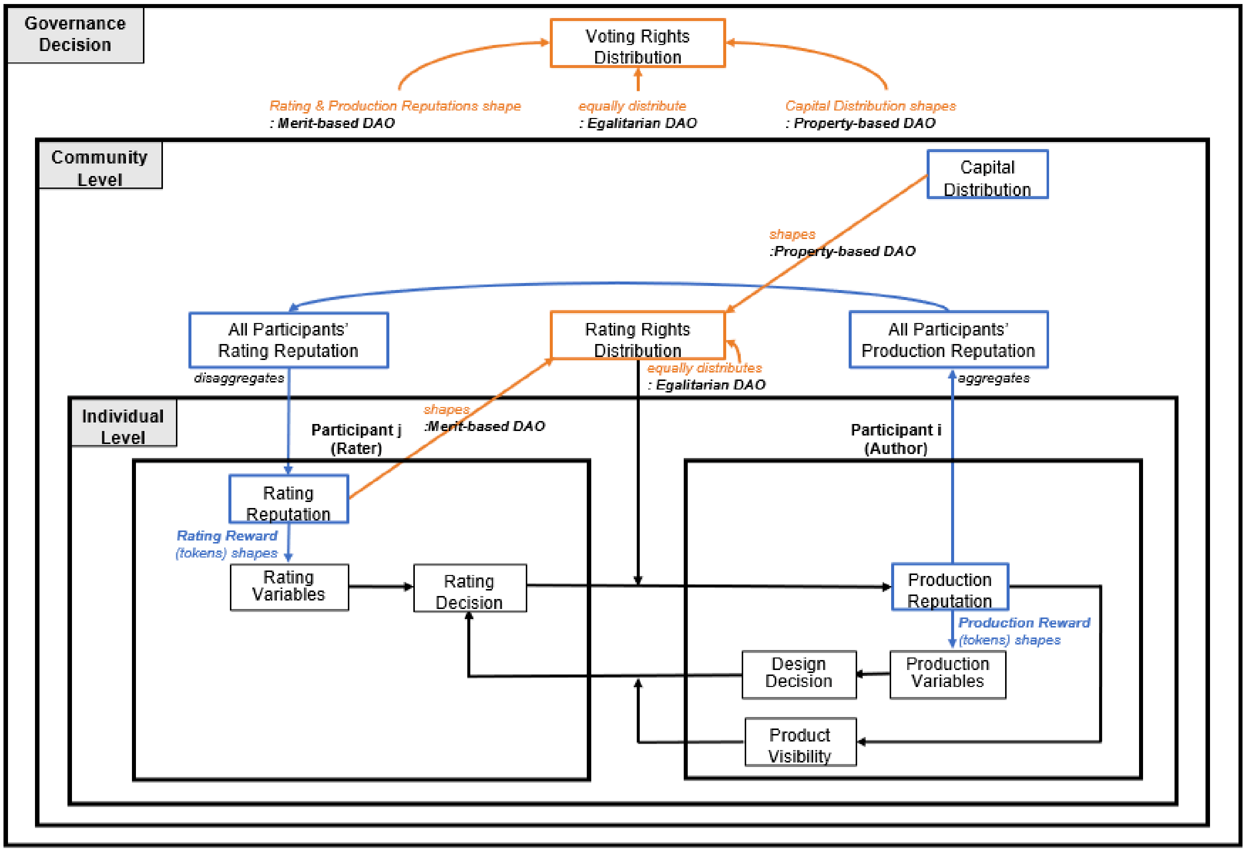

Figure 4 DAO Governance: Egalitarian, Merit-based, and Property-based

Note. Each participant can engage in DAO as an author and a rater.

Table 3 Comparison of DAO Governance Models

	Egalitarian	Merit-based	Property-based
Principle	Equal rights for all participants	Rights allocated based on prior contributions (reputation)	Rights allocated based on financial investment
Prioritized Value	Inclusiveness	Recognition of expertise and commitment	Investor influence
Authors' Expectation	Popular designs are likely to receive high ratings	Designs endorsed by reputed curators are likely to receive high ratings	Ratings may vary depending on the investment distribution (random)

Note: *Rights* in the table refer to both Rating Rights (for design evaluation) and Voting Rights (for governance decision).

4.2.2. Author i's Production Reputation. Production Reputation η_i holds a more extensive role beyond its representation within the DAO. Production Rewards (tokens) are allocated to each author i in proportion to η_i. These tokens possess the capacity to shape the Production Variables of every participant, thereby cascading their effects onto individual production behaviors (detailed in Section 4.2.4).

4.2.3. Rater j's Rating Reputation. For each rater j, Rating Reputation, ϕ_j, is determined based on j's ongoing and historical involvement in appraising others' work. The notation ξ_j signifies the extent of j's contribution to the DAO through their endorsements of authors possessing commendable Production Reputation. That is, as rater j assigns a higher rate to authors bearing higher Production Reputation, j's Rating Reputation increases. The measure of rating contribution (ξ_j) is derived from equation (4).

$$\xi_j = \frac{\sum_i \eta_i \frac{l_{ji}}{l_{.i}}}{\sum_k \sum_i \eta_i \frac{l_{ki}}{l_{.i}}} \tag{4}$$

Equation (4) elucidates that the extent of contribution hinges upon the allocation of Rating Rights to rater j (l_j) and how these Rating Rights are assigned to various authors possessing distinct Production Reputation (η_i). Here, l_{ji} denotes the Rating Rights apportioned from j to i (for an elaboration, refer to Section 4.2.1 for the specific delineation of l_j within each DAO scenario). The essence of contribution ξ_j is encapsulated in the weighted mean of the number of Rating Rights accorded by rater j to each author relative to the weighted mean of Rating Rights prevalent in the DAO. A heightened concentration of endorsement directed towards members with elevated Production Reputation invariably augments the value of contribution, denoted by ξ_j.

The accumulation of a rater's Rating Reputation, ϕ_j, occurs through a cumulative process of evaluations. Consequently, the assessments undertaken during a rater's initial foray into assessing authors hold greater significance than those conducted after establishing a robust reputation. This delineation is particularly relevant in emphasizing the ascendancy of the Rating Reputation, which undergoes a re-calibration process delineated by equation (5), representing the reinforcement learning algorithm (e.g., Sutton and Barto 1998) used to learn the reputation level of each author.

$$\phi_j := \frac{t}{t+1}\phi_j + \frac{1}{t+1}\xi_j \tag{5}$$

The accumulation of ϕ_j at time $t+1$ unfolds gradually. The most recent observation, ξ_j, is ascribed a weight of $\frac{1}{t+1}$ to signify the diminishing importance of supplementary information with the progression of iterations. On the other hand, the historical data, embodied by ϕ_j at time $t+1$, assumes a weight of $\frac{t}{t+1}$, underscoring its enduring relevance in the evolving context.

4.2.4. Endogenous Rating and Production Variables. As depicted in Figure 4, a pivotal attribute of DAO is incorporating token-based rewards to shape participants' behavior. This is achieved by rendering the Rating Variables and Production Variables that influence production and rating behaviors endogenous to the system (see Table 1).

The approach to render these state variables endogenous varies for rating and production behavior. In the context of raters, these variables are formulated as functions of the Rating Reputation, ϕ_j. Rating Rewards (tokens) are given in proportion to the Rating Reputation as a reward for the rating engagement. A heightened Rating Reward instigates greater motivation among raters to improve their rating behavior: prompts the rater to scan a more extensive sample of authors (Q_j^r) and to post ratings more frequently (P_j^r). Moreover, in the case of the Merit-based DAO, an augmented Rating Reputation translates into increased ownership of Rating Rights, l_j, thereby amplifying the rater j's influence within the DAO. Hence, in the Merit-based DAO, the elevation in participation is driven not only by the prospect of additional token rewards but also by the desire to accrue greater within-community authority through the proportional allotment of Rating Rights.

Consequently, the Rating Rewards (alongside the Rating Rights in the context of the Merit-based DAO) create a positive feedback system within DAO, incentivizing raters who consistently perform above average to become more active in the community. This increased activity enables them to gather more information. Simultaneously, the system reduces rewards for raters with subpar performance, diminishing their motivation to collect information and engage with the community. Over time, this may result in their exclusion from the DAO rating system. Thus, this mechanism is strategically designed to identify and empower high-performing raters, positioning them as DAO curators.

Similarly, in the context of production, as a member's Production Reputation, η_i, increases, they are allocated more Production Rewards (tokens). The Production Rewards positively affect the search radius for learning (β_i), the probability of production improvement (P_i^p), the number of rated works that i scans to improve its own production (Q_i^p). Hence, the DAO Production Reward allocation and Product Visibility mechanisms establish a feedback loop wherein popular authors are inclined to enhance their involvement in DAO by creating innovative products. Furthermore, these prominent authors are better aligned with raters' quality perception because they can gather information from a larger rater sample while contemplating product innovation. As a result, accomplished authors are more likely to accumulate influence over time, becoming integral to the essence of DAO's creative ecosystem.

4.2.5. DAO Governance: Voting Rights Distribution Finally, a fundamental aspect of DAO is the collective determination of the organization's governance structure through voting. The DAO allocates Voting Rights to each member, granting them the power to influence decisions about the governance setup (as depicted in Figure 4). In our context, for instance, participants can collaboratively decide whether a given DAO should adopt an Egalitarian, Merit-based, or Property-based governance structure. Moreover, within the Merit-Based DAO, participants can jointly decide on the criteria governing the significance of past activities when the DAO disburses Rating Reward (as illustrated by equation 5).

The viewpoints of participants possessing more Voting Rights will hold more significant sway in shaping these governance choices. The three distinct DAO governance structures and their associated Voting Rights allocations are as follows: 1) Egalitarian DAO: Every member is granted equal Voting Rights. 2) Merit-based DAO: Voting Rights are assigned proportionally to the combined weight of Rating Reputation and Production Reputation (refer to equation 6). The parameter σ dictates how Production and Rating Reputations influence Voting Rights calculation. Given that community development hinges more on authors' contributions than raters' input (as evidenced by real-world communities where most participants engage in rating rather than production; Section 2), σ should be close to 1 to incentivize increased production. However, alternative configurations are plausible, contingent upon the collective consensus of community members. 3) Property-based DAO: Voting Rights are distributed based on the level of capital invested in the community.

$$\text{Voting Rights} = (1-\sigma) \times \text{Rating Reputation} + \sigma \times \text{Production Reputation} \tag{6}$$

In summary, the system can be comprehended as comprising rewards for contributions and a means to amass influence in DAO governance via voting.

5. Agent-Based Simulation

This section provides a comprehensive description of the agent-based simulation. The Overview, Design concepts, and Detail (ODD) protocol is provided in Appendix A. The pseudo-code of the agent-based simulation algorithm is summarized in Table 4. The initialization stage defines the simulation parameters: T, I and δ. The designs s_i, preferences ω_i, and idiosyncratic rating threshold a_i, for each member i, are populated. The Rating Rights owned by each member i, l_i, are initialized using the power law, so that 20% of the members own 80% of the rights. This initial setup makes the Property-based DAO reflect reality, where the capital distribution follows a power law distribution (Gabaix 2009). The Production and Rating Reputations are initialized assuming a uniform distribution over the set of members, i.e., each member has a reputation $1/I$.

Then the algorithm proceeds by initializing the ratings for each community member. First, the rating probability (P_j^r) and the number of authors rated by j (Q_j^r) are dependent on the Rating Reputation of rater j. The higher the Rating Reputation of j, the more Rating Rewards (tokens) j receives, which in turn leads to a higher rating probability (P_j^r) and an increase in the number of authors rated by j (Q_j^r). Second, in the initialization stage, only when the potential rating of a given author i is higher than the minimum threshold for approval by rater j (a_j), an author i receives a positive rating (r_{ji}). Third, the initialization finishes with calculating the total ratings received and awarded by each community member.

Then, the agent-based game starts. For each iteration (1 to T), the simulation updates the Production Reputations (Section 4.1.3) and Rating Reputations (Section 4.2.3), distributes rights and rewards, and update the production (design) (Section 4.1.2) and rating (Section 4.1.1).

Reputation updates are calculated as a weighted average, reflecting the reinforcement learning process. As the game's history grows longer, the most recent observations become less significant in calculating reputation. We note that all agent updates in each simulation step are effectively applied in parallel to prevent any ordering bias. In each iteration, reputation updates for all agents are computed based on the prior state before any agents take new actions. This synchronous update scheme ensures that no agent gains a first-mover advantage due to update order.

Regarding the update of designs, for each author i, the probability of attempting to improve the designs (P_i^p) and the number of rated works that i scans to improve its production (Q_i^p) are endogenous and a direct function of the Production Reputation (η_i). Consequently, highly reputed authors are more likely to make greater efforts to improve their products.

The design development is a function of the endogenous radius (β_i): the higher the Production Reputation (η_i), the more the author ventures to search for new designs away from its current style. The search for new designs occurs in two directions (increase and decrease), based on the current design, s_i. The search step also includes a random component generated from a uniform distribution between zero and one, represented by ν and μ. The designs developed are between -1 and 1. The author i chooses to keep as the new design the one that receives the highest rating s_i^* from the Q_i^p rated works scanned.

In the ratings updating phase, for each rater j, there is a probability of j engaging in rating (P_j^r) and the number of products scanned by j to be rated (Q_j^r), which are a direct function of the Rating Reputation ϕ_j. And there is a minimum a_j required by j to give an author i a positive rating, r_{ji}.

Table 4 Pseudo-code for the Agent-Based Model - Game

Initialization:

    Initialize the simulation parameters, T, I and δ.
    For $i = 1, ..., I$, initialize: s_i and $\eta_i = 1/I$
    For $j = 1, ..., I$, initialize: ω_j, a_j, $\phi_j = 1/I$, and $l_j = \frac{B_j}{\sum_j B_j} \times 100$, in which $B_j = \frac{100}{j^{1.215}}$
    Initialize ratings:
        With probability $P_j^r = a_j^r + b_j^r \phi_j$:
            Select a sample with a maximum of $Q_j^r = c_j^r + d_j^r \phi_j$ authors to vote for
            For each author i in the sample:
$$r_{ji} = \begin{cases} \frac{2-\delta(s_i - w_j)^2}{2}, & \text{if } \frac{2-\delta(s_i - w_j)^2}{2} \geq a_j \\ 0, & \text{otherwise} \end{cases}$$
    Calculate Rating Rights used by j (overall engagement):

For $t = 1$ to T:

    Update Production Reputation:
    For $i = 1, ..., I$:
        Production Reputation: $\eta_i := \frac{t}{t+1}\eta_i + \frac{1}{t+1}\frac{l_{.i}}{\sum_k l_{.k}}$
    Update Rating Reputation:
    For $j = 1, ..., I$:
        Contribution: $\xi_j = \frac{\sum_i \eta_i \frac{l_{ji}}{l_{.i}}}{\sum_k \sum_i \eta_i \frac{l_{ki}}{l_{.i}}}$
        Rating Reputation: $\phi_j := \frac{t}{t+1}\phi_j + \frac{1}{t+1}\xi_j$
    Update production (design):
    For $i = 1, ..., I$:
        Adjust production improvement probability:
            $P_i^p = a_i^p + b_i^p \times \eta_i$
        Adjust production improvement sample size:
            $Q_i^p = c_i^p + d_i^p \times \eta_i$
        Adjust design search radius: $\beta_i = a_i^\beta + b_i^\beta \times \eta_i$
        With probability P_i^p attempt to improve designs:
            Neighborhood search:
            Randomly generate ν and μ from $U[0, 1]$
            $s_i^+ = \min(s_i(1 + \beta_i \mu), 1)$
            $s_i^- = \max(s_i(1 - \beta_i \nu), -1)$
            $s_i^* = \arg\max\left[l_{.i}(s_i^-), l_{.i}(s_i), l_{.i}(s_i^+)\right]$
    Update ratings:
    For $j = 1, ..., I$:
        Adjust the rating probability: $P_j^r = a_j^r + b_j^r \phi_j$
        With probability P_j^r:
            Select a sample with a maximum of $Q_j^r = c_j^r + d_j^r \phi_j$ authors to vote for
            For each author i in the sample:
$$r_{ji} = \begin{cases} \frac{2-\delta(s_i - w_j)^2}{2}, & \text{if } \frac{2-\delta(s_i - w_j)^2}{2} \geq a_j \\ 0, & \text{otherwise} \end{cases}$$

6. A Co-Evolutionary Analysis of DAO Governance Structure

This section elucidates the computational experiments conducted using the agent-based model. It considers the influence of distinct governance alternatives on the progression of DAO. To capture a realistic representation of the art community, we initialize our model using the distributions of agents' designs s_i and preferences w_j depicted in Figure 5. Communities where people create art usually lack objective rating criteria but feature distinct design preferences by artists and audiences. Consequently, we assume that designs and preferences are diverse and distributed across the spectrum of potential outcomes from -1 to 1. In particular, in the real world, there are a few genres which the majority of the audience favors the most at a certain period, although the popularity evolves in the long term (Mauch et al. 2015). To reflect this, we assume that the distribution of preferences (w_j) depicted in Figure 5 encompasses mainstream art (concentrated around zero) and less prevalent other genres.

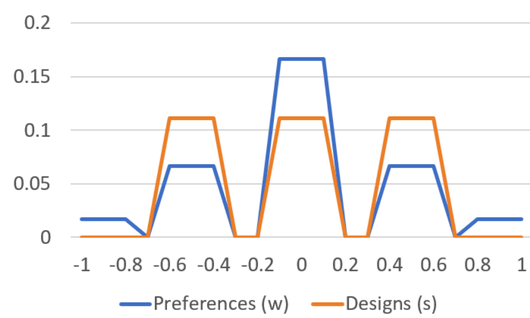

Figure 5 Distribution of Initial Designs (s_i) and Preferences (w_j)

Note. X-axis indicates the range of designs (s_i) and preferences (w_j), both are within -1 and 1. The Y-axis denotes the proportion of agents in each range of s_i and w_j.

The analysis begins by examining the convergence properties of the simulation, as well as the proportion of Rating Rights allocated to authors (indicative of overall engagement in rating) (Figure 6). A substantial engagement rate holds significance, linking to contributors' retention rates and the economic value of cryptographic tokens within the community. This, in turn, influences the long-term prosperity of the DAO. The subsequent analysis delves into how different governance scenarios impact the distribution of Production and Rating Reputations (illustrated in Figure 7). The reputation distribution showcases the community's inclusiveness and reveals whether it is primarily driven by a select few members while the contributions of the majority remain negligible. The third facet of simulation results outlines how distinct governance structures shape the distribution of designs acquired by authors (refer to Figure 8) and the likelihood of authors enhancing their designs (detailed in Table 6). Design evolution in the art community should foster greater diversity.

6.1. Parameters and Variables

The number of iterations is $T = 10,000$ and the population size is $I = 100$. The rating decay rate δ (see Equation (1)) was chosen through calibration so that the baseline Open Community scenario reproduces the heavy-tailed distribution of contributor activity observed in real communities (Figure 1) (see Appendix B for further details). We set the idiosyncratic rating decay rate to $\delta = 1000$, which not only reproduces real community distributions but is also large enough to enforce selective ratings, in which raters only endorse designs close to their own tastes. Small differences between an author's design and a rater's preference can lead to no rating.

All reported outcomes are computed as averages across the population of agents within a given simulation run. Each governance configuration is simulated over $T = 10,000$ iterations. While each configuration is typically run once due to computational intensity, we verified that repeated simulations with different random initializations converge to qualitatively similar steady-state outcomes for sufficiently large populations ($I = 100$). This reflects the stability of the reinforcement learning dynamics and reduces sensitivity to idiosyncratic initial conditions.

Rater j's preference w_j regarding design features is drawn from a specified distribution that depends on the scenario under examination, and j's idiosyncratic rating threshold a_j is uniformly distributed on $[0, 1]$. Rating Rights ownership by j is initialized via a power-law distribution (20% of members hold 80% of the rights), reflecting real-world wealth/contribution disparities: : $l_j = \frac{B_j}{\sum_j B_j} \times 100$, where $B_j = \frac{100}{j^{1.215}}$. Furthermore, the simulation integrates several state variables to endogenize probabilities and sample sizes in rating and production, as well as the radius of the production search process. Table 5 summarizes the variables used to endogenize the production and rating behavior as a function of the respective reputations.

Table 5 Production and Rating Variables

Function	Intercept	Slope
Production Improvement Probability (P_i^p)	0.02	$I/5$
Production Improvement Sample Size (Q_i^p)	2	$10 \times I$
Design Search Radius (β_i)	0.1	I
Rating Probability (P_j^r)	0.1	$I/5$
Rating Sample Size (Q_j^r)	2	$5 \times I$

There are some important justifications for the specific variable values listed in Table 5. These intercepts and slopes governing production and rating behaviors were calibrated to ensure plausible levels of agent activity. For instance, an initial rating probability of 0.1 means even new or low-reputation members occasionally rate content, and the chosen slope values increase this probability gradually with higher reputation so that it remains below 1. Similarly, the sample size variables were set to grow moderately with reputation, keeping search breadth reasonable. These values were determined via simulations to achieve dynamic yet stable behavior in the model, and to reflect engagement levels that are realistic for an active online community. Notably, we found that moderate variations in these variables do not qualitatively alter our results. We further discuss this robustness in the results (Section 6.3) and Appendices B and C.

6.2. Results

Figure 6 summarizes the overall engagement in rating through the evolution of the proportion of Rating Rights allocated to the authors (equivalent to the *likes* expressed in the case of the Open Community). The Figure illustrates the rating engagement dynamics, revealing that the Merit-based DAO governance structure surpasses all others, ultimately converging to an allocation of nearly 100% Rating Rights to the authors. Surprisingly, the remaining DAO structures exhibit closely similar Rating Rights allocations, ranging from 20% to 30%, and are eclipsed by the Open Community's performance, which maintains a rating engagement level of approximately 40%. This result shows that transparent recognition and rewards for high-quality raters through the distribution of Rating Rights could serve as incentives for participants to engage in the community in the long run.

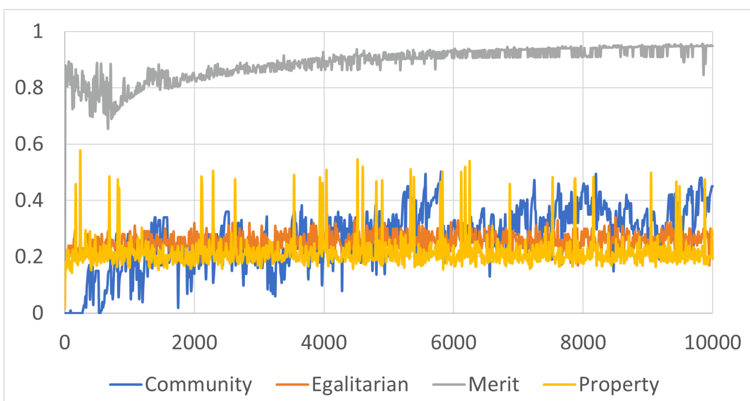

Figure 6 Overall Engagement in Rating (proportion of Rating Rights used)

Note. X-axis indicates the number of iterations.

Figure 7 presents the distributions of Production and Rating Reputations for both the Open Community and DAO scenarios. These simulations present an inequality in the distribution of Production Reputation within the Open Community (blue dots in Figure 7(a)), mirroring the patterns observed in real-world communities (Figure 1(a) and (b)). This Figure illustrates how our baseline model for the Open Community scenario emulates the real-world open community with no explicit reward. Notably, just two authors amass over 60% of the total Production Reputation, with only 10% of authors holding a positive Production Reputation. This stark concentration indicates that public attention is heavily focused on a few authors. In the case of Rating Reputation within the Open Community, the value is zero by definition (Figure 3), as evident from the outcomes displayed (blue dots in Figure 7(b)).

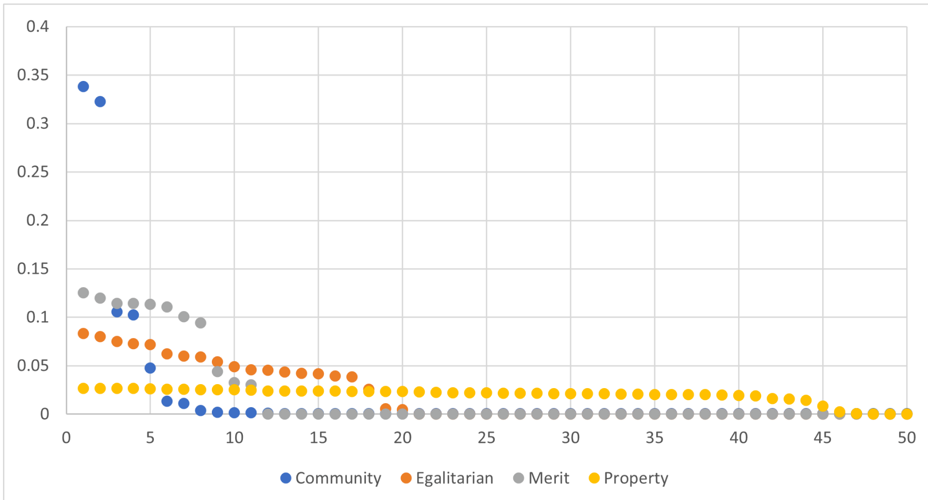

(a) Production Reputation

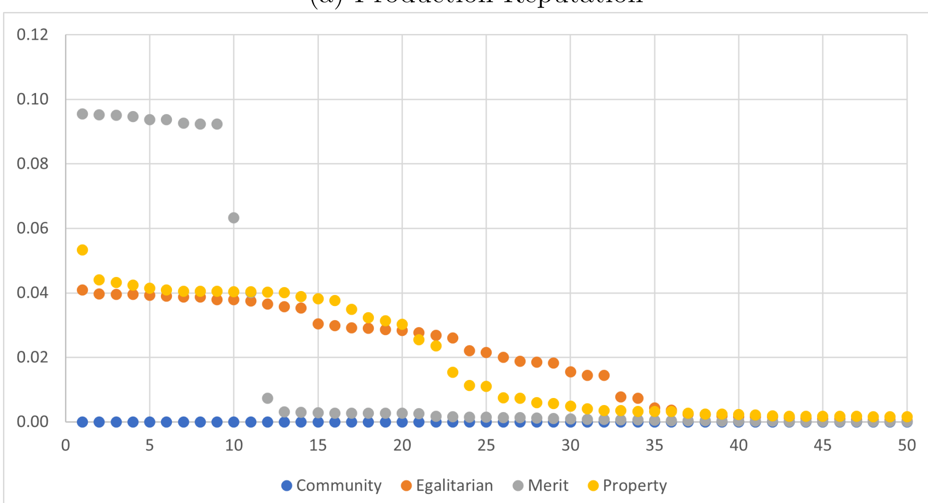

(b) Rating Reputation

Figure 7 Production and Rating Reputations (Y-axis) (Converged State)

Note. Each dot indicates an agent. The X-axis indicates an agent-specific ID in our simulation ($I = 100$, The Figures present 50 agents).

Egalitarian and Property-based DAOs exhibit notably larger numbers of authors with positive Production Reputations than the Open Community. Specifically, the Egalitarian DAO showcases 20% of authors possessing positive Production Reputation. Interestingly, the Property-based DAO demonstrates a more egalitarian progression than the Egalitarian DAO. In this case, around 45% of the authors attain positive Production Reputation, accompanied by a maximum reputation of merely 2.5%, which is relatively smaller in comparison to other scenarios (Figure 7(a)). These outcomes imply that an uneven and random distribution of Rating Rights may foster a less centralized DAO, at least in terms of rewarding authors.

The Merit-based DAO exhibits remarkable emergent behavior in its own right. It evolves as a system driven by relatively a small proportion of influential contributors, although not to the extent of the Open Community. Approximately 10% of the members attain positive Production Reputation. Similarly, the Rating Reputation is dominated by around 10% of DAO members, who have developed the role of curators responsible for overseeing the art presented within the DAO. Such results imply that the small population actively engaged in high-quality rating and production, as the Merit-based DAO resulted in nearly full engagement rates (Figure 6).

We are now well-positioned to elucidate the influence of distinct governance structures on authors' design improvement. Regarding the initial design distribution, we posit the existence of three prominent design clusters emblematic of the design diversity within the Art-oriented community (Figure 5). Commencing from these clusters, each scenario demonstrates varying degrees of learning engagement under diverse governance frameworks. As evident in Figure 8(a), the Open Community scenario exhibits minimal learning, with only a handful of elements transitioning from their initial design levels.

The learning process becomes more pronounced in the Egalitarian DAO, depicted in Figure 8 (b). Notably, there is a more substantial migration toward the center of the design spectrum (the trendy zone). At the same time, specific agents also demonstrate an aptitude for transitioning to less frequented positions. Furthermore, noticeable shifts in design levels are observed, with particular authors transitioning from a positive design set to a negative one, and vice versa. This trend is similarly discernible in the Merit-based DAO, as depicted in Figure 8(c).

The Property-based DAO exhibits remarkable and unexpected shifts in designs. Our simulation reveals that the Property-based DAO structure ignites authors' enthusiasm for exploring and learning new designs. We also observe substantial leaps where many authors transition from positive to negative design sets and vice versa. This outcome suggests that the distribution of Rating Rights determined by capital investment could potentially act as the seeds of attraction for authors seeking design exploration. This could be driven by the fact that influential raters' preferences are relatively random rather than concentrated on certain designs, which contrasts with the Egalitarian or Merit-based DAO (see Table 3). When influential raters' preferences are somewhat random rather than aligned with mainstream taste or reputed curators' taste, it is sensible that authors are likely to explore more new designs across domains rather than sticking to popular ones.

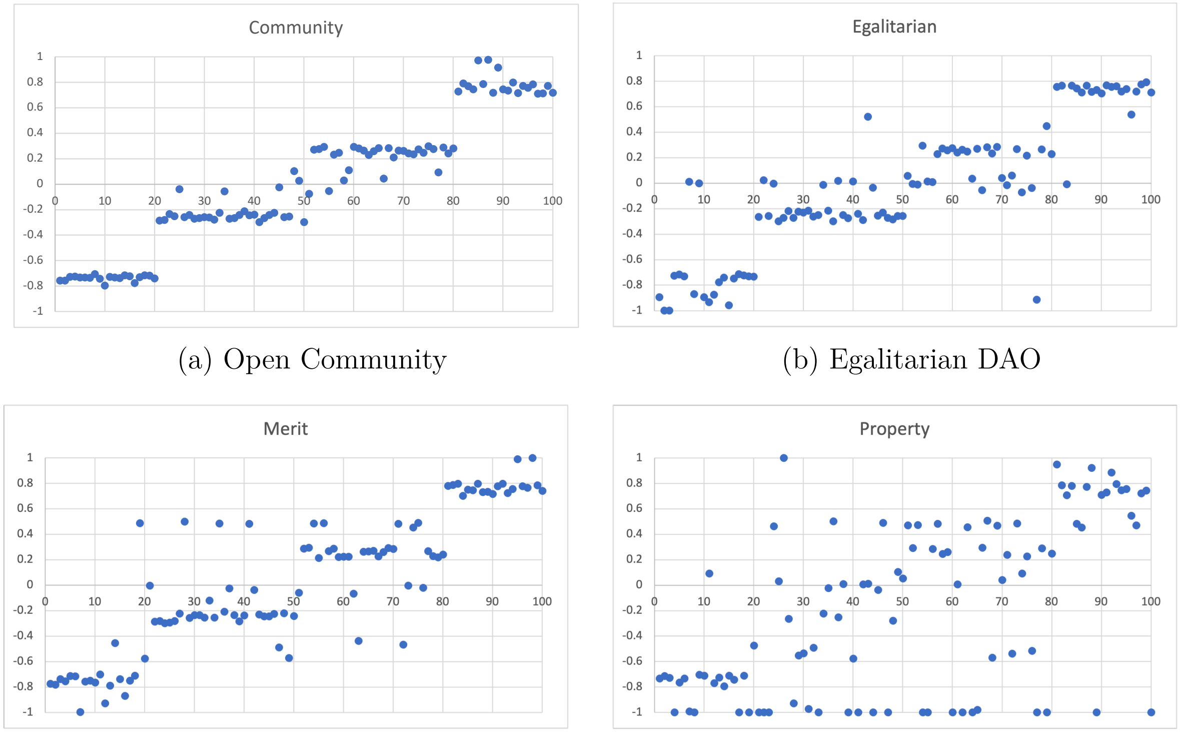

(a) Open Community (b) Egalitarian DAO

(c) Merit-based DAO (d) Property-based DAO

Figure 8 **Designs (Y-axis) Under Different Governance Structures (Converged State)**

Note. Each dot indicates an agent.The X-axis indicates an agent-specific ID in our simulation ($I = 100$).

The design learning outcomes are summarized in Table 6. This table provides insights into the percentage of authors who made sufficient production improvements to transition to design areas where their work gains higher value. In the Open Community scenario, only 12% of authors demonstrated learning. The Property-based DAO exhibited the highest level of learning, as previously discussed, with 58% of authors enhancing their designs, contributing to increased design diversity within the community. The Egalitarian and Merit-based DAOs showed proportions of 28% and 23% of authors, respectively, who engaged in learning to refine their designs.

Table 6 **Production Improvement Probability (P_i^p) Across Different Governance Structures (Converged State)**

	Community	Egalitarian	Merit	Property
Probability	0.12	0.28	0.23	0.58

Figure 9 depicts the distribution of Voting Rights within the Property-based and Merit-based DAOs, considering varying values of σ (namely, $\sigma = 0.1, 0.5,$ and 0.9) (see equation 6). It's worth noting that in the Egalitarian DAO, all members enjoy equal Voting Rights. In the Property-based

DAO, the distribution becomes highly skewed, with the top 5% members holding control over more than 50% of the Voting Rights. Within the Merit-based DAO, the Voting Rights are relatively less concentrated than the Property-based DAO, although the rights are still distributed among 10% of the participants. The influence of σ on the Voting Rights distribution is minor.

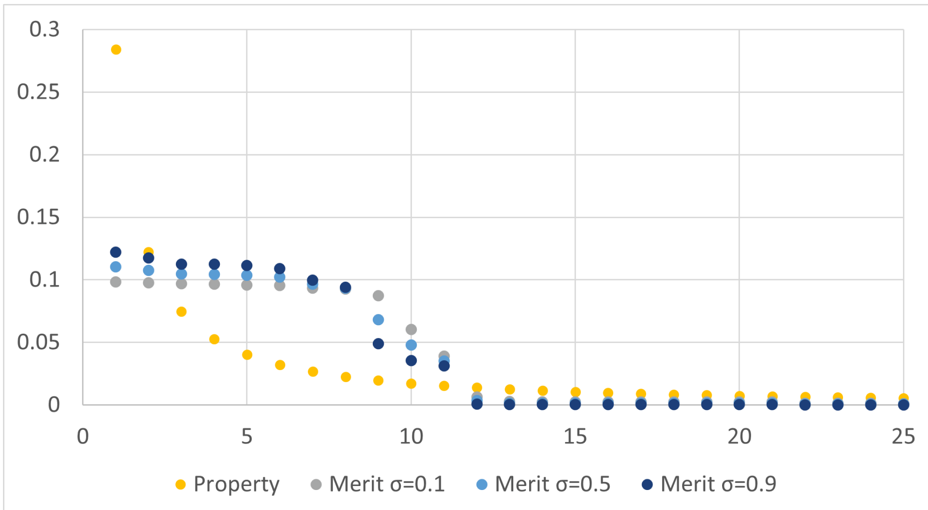

Figure 9 Distribution of Voting Rights (Converged State)

Note. Each dot indicates an agent. The X-axis indicates an agent-specific ID in our simulation ($I = 100$, The Figure presents 25 agents).

6.3. Summary

In summary, simulation results provide insights into the performance of different governance structures in community evolution. First, the Merit-based DAO allocates the most Rating Rights to the participants (Figure 6). To sustain the token-based reward mechanism, the economic value of the cryptographic tokens should increase or at least remain stable over time, and the value is correlated with the participants' community engagement. In this regard, Merit-based DAO may be an effective governance structure in maintaining the community's long-term prosperity, given its high engagement level over time. Moreover, the Merit-based DAO creates a significant class of highly reputable raters acting as elite curators (Figure 7(b)) and leads to the development of a group of authors fully rewarded by DAO (Figure 7(a)). A potential concern is the concentrated distribution of rewards, which implies relatively weak incentives for a large proportion of participants compared to other DAO models.

Second, one interesting result is the Property-based DAO's performance with a skewed capital distribution. We had expected this DAO to perform poorly in every dimension, but this was not the case. The Property-based DAO was appropriate for incentivizing learning (Table 6) leading to community content diversity. However, our results also indicate that the Property-based DAO leads to a relatively low overall engagement compared to the Merit-based DAO (Figure 6).

We examined the sensitivity of our findings to key parameters, variables and assumptions. We conducted additional simulation runs varying parameters such as σ (the weight balancing reputation and investment in the voting rule) (see Figure 9) and the initial distribution of designs (s) and preferences (w) (see Appendix C). Encouragingly, the main qualitative outcomes remained consistent across these variations. For example, in every tested scenario the Merit-based DAO still achieved the highest long-term engagement levels, and the Property-based DAO continued to foster a diverse range of productions (albeit with low overall engagement). This gives us confidence that our conclusions are robust to reasonable changes in parameter values and model structure.

7. Discussion and Conclusions

This article considers the introduction of DAO as an alternative to the conventional open community model where people share their creative works and evaluate others' contributions. Our motivation stems from observing how existing open communities have evolved, relying heavily on participants' intrinsic motivation, which often has not been sustained in the long term. This study breaks new ground by designing and simulating DAO governance structures tailored to creativity-driven open communities. For simplicity, we focused on three archetypal governance models (Egalitarian, Merit-based, and Property-based). We acknowledge that real-world communities might adopt hybrid or more nuanced governance approaches – such as quadratic voting, rotating leadership roles, or delegated rating rights – which lie beyond our present scope. These alternative designs are an interesting avenue for future research building on our framework.

We demonstrate how different designs of DAO governance can lead to the emergence of distinct behavior and community forms in the long run. Our paper contributes to the literature on incentive mechanisms in open communities by providing one of the first explorations of the governance mechanisms of DAO. Our simulation design provides a baseline tool for further exploration of a suitable DAO governance design, depending on the purpose of organizing a community. The proposed model can be applied when there is a requirement for the crowds' creativity to design new products or services.

As a methodological contribution, the article uses reinforcement learning in an agent-based model to simulate the emergent reputations in the DAO, considering two types of agents (authors and raters), with individual features. This technique allows the comparison of the different organizational structures, describing how the individual participants learn and the community evolves. In particular, our model proposes a novel *feedback mechanism* to study the relationship between the rating decisions and the associated rewards for *both* raters and authors. The proposed feedback mechanism allows us to study how the rating activity of an agent not only shapes the agent's Rating Reputation but also affects other agents' production activity.

This article bears some interesting implications for firms operating their platform or online open community to leverage people's creativity. While firms have been increasingly relying on the wisdom of crowds to innovate (Bogers et al. 2017), the effective management of the external sources of creativity remains largely unknown (Wei et al. 2015). This study is an early exploration to understand how firms can better utilize open communities by employing the cryptographic token-based reward mechanism. Properly designed DAO will incur less governance management efforts than conventional approaches, enabling firms to gain creativity with much lower transaction costs. The DAO structure explored in this study bears instrumental implications in that it can reshape existing patterns and costs of broadening the firm boundary (Holcomb and Hitt 2007, Leiblein and Miller 2003).

Moreover, crowd-sourcing has been increasingly implemented by not only firms but also public sector organizations. However, how to *manage* crowds to drive innovation remains a less understood question in the crowd-sourcing and innovation contest literature (Lakhani 2016). Given the significant increase in demand for social and environmental innovation through the utilization of crowds' creativity, this study holds important practical implications for the future of crowd-sourcing. Our simulation can be adapted to experiment with various crowd-sourcing scenarios. For example, when crowd-sourcing aims to solve technological or engineering problems, it is more likely that the community can apply relatively objective criteria for rating compared to communities focused on non-technical creativity presented in this study. In such communities, we can expect less diversity in preferences in terms of rating, as the evaluation criteria are more obvious. Variables can be adjusted to reflect scenarios in which participants anticipate a few outstanding or competing engineering solutions. We have conducted additional sensitivity analyses for such communities (Appendix C).

Despite the potential benefits, the negative aspects of open communities for creative works should be carefully considered when designing DAO governance. For example, in a community where reputation is linked to rewards, there is a risk that some participants may attempt to game the reputation system. This concern is increasingly relevant given the prevalent use of generative AI in creative work, which substantially reduces the marginal cost of content production and may lead to an exponential increase in submissions. In the present model, each author is represented by a stylized design type rather than by a stream of individual productions. Therefore, production volume does not directly influence reputation dynamics in our model. This modeling abstraction isolates governance effects from quantity-based competition.

However, in practice, the proliferation of AI-generated content may increase submission volume and create evaluation bottlenecks. In response, DAOs may need to impose governance safeguards, such as limiting the number of submissions within a given period, introducing reputation decay mechanisms, or evaluating contributions based on representative design quality rather than sheer quantity. In addition, while our model assumes that rating mechanisms serve as quality assessors, supplementary mechanisms for quality control, abuse detection, and prevention of coordinated manipulation could be useful in practical implementations. The advances in AI may support anomaly detection or reputation validation, thereby strengthening system resilience. Future extensions of the model could explicitly incorporate strategic overproduction or AI-induced productivity asymmetries to formally evaluate governance resilience under adversarial conditions.

In addition, there is a risk that a small number of highly reputed participants dominate the community's trends and decisions, deteriorating diversity in the long run. If a community aims to maintain creativity through diversity, the reputation-based reward system needs to be complemented by alternative mechanisms. For example, the DAO can add additional reward criteria, such as providing exceptional rewards to the authors who produce *atypical* but high-quality content and the raters who recognize and properly evaluate such authors. Such complementing mechanisms for identifying hidden stars could be used to overcome the dominance issue and motivate learning in diverse directions. This approach is also in line with our results from the Property-based DAO, which show that learning occurs in more diverse directions in the long run when authors expect to gain rewards even for contributions to designs that are distant from the most popular ones.

There are other dimensions of DAO that we have not considered in this study, as our work is one of the earliest attempts to experiment with DAO governance for creativity.

First, our simulation assumes that cryptographic tokens have sufficient economic value to influence agents' behavior. An important open question is how a new DAO can attract a critical mass of participants in its early stages before the token has established value. Moreover, token value can fluctuate over time, potentially leading to speculative booms or crashes that alter participant behavior. We did not model these early-stage bootstrapping challenges or volatile token dynamics. They remain important topics for future investigation into real-world DAO evolution and its implications for community sustainability.

Second, as with other simulation models, there is a trade-off between complexity and realism that we had to consider in model design. Our simulation is a simplified representation of decision-making in real-world DAOs and open communities. For example, while the rating decision in our model is computed based on the gap between design and preference, this gap may also be influenced by factors such as peers, network dynamics, and group norms. Building on the baseline model we provide, future studies could deepen understanding by exploring scenarios in which these additional factors play a critical role in agents' decisions.

Third, the power to decide DAO governance remains concentrated within a relatively small group of participants, although the Merit-based DAO shows a relatively mitigated level of concentration of Voting Rights. The concentration of power is a practical limitation of current DAOs, as discussed by other scholars and practitioners (e.g., Hsieh et al. 2018). Future studies need to understand the conditions in which the DAOs can reach a more decentralized state to achieve their original intention.

References

Antonopoulos, A. M. and Wood, G. (2018), *Mastering Ethereum. Building Smart Contracts and DApps.*, O'Reilly Media, Sebastopol, CA. ISBN: 1491971940.

Axelrod, R. (1997), *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration.*, Princeton University Press, Princeton, New Jersey. doi: 10.1515/9781400822300.

Bandyopadhyay, S., Rees, J. and Barron, J. (2008), 'Reverse auctions with multiple reinforcement learning agents.', *Decision Sciences* **39**(1), 33–63. doi: 10.1111/j.1540-5915.2008.00181.x.

Blanchard, A. L. and Markus, M. L. (2004), 'The experienced "sense" of a virtual community: Characteristics and processes', *ACM Sigmis Database: the Database for Advances in Information Systems* **35**(1), 64–79. doi: 10.1145/968464.968470.

Bogers, M., Zobel, A.-K., Afuah, A., Almirall, E., Brunswicker, S., Dahlander, L., Frederiksen, L., Gawer, A., Gruber, M., Haefliger, S. et al. (2017), 'The open innovation research landscape: Established perspectives and emerging themes across different levels of analysis', *Industry and Innovation* **24**(1), 8–40. doi: 10.1080/13662716.2016.1240068.

Boudreau, K. J., Lacetera, N. and Lakhani, K. R. (2011), 'Incentives and problem uncertainty in innovation contests: An empirical analysis', *Management Science* **57**(5), 843–863. doi: 10.1287/mnsc.1110.1322.

Boudreau, K. J. and Lakhani, K. R. (2011), 'How to manage outside innovation', *MIT Sloan Management Review Sloanselect Collection* pp. 50–57. url: https://sloanreview.mit.edu/article/how-to-manage-outside-innovation.

Bunn, D. and Oliveira, F. (2008), 'Modeling the impact of market interventions on the strategic evolution of electricity markets.', *Operations Research* **56**(5), 1116–1130. doi: 10.1287/opre.1080.0565.

Chohan, U. W. (2017), The double-spending problem and cryptocurrencies, *in* 'Public Governance on the Blockchain', Routledge, pp. 71–84. doi: 10.4324/9781003610601.

Conti, A., Gupta, V., Guzman, J. and Roche, M. P. (2023), Incentivizing innovation in open source: Evidence from the github sponsors program, Technical report, National Bureau of Economic Research. doi: 10.3386/w31668.

Daskalova, V. and Vriend, N. (2021), 'Learning frames.', *Journal of Economic Behavior & Organization* **191**, 78–96. doi: 10.1016/j.jebo.2021.08.020.

Didden, J. B., Dang, Q.-V. and Adan, I. J. (2024), 'Enhancing stability and robustness in online machine shop scheduling: A multi-agent system and negotiation-based approach for handling machine downtime in industry 4.0.', *European Journal of Operational Research* **316**(2), 569–583. doi: 10.1016/j.ejor.2024.02.006.

Diogo, A., Fernandes, B., Silva, A., Faria, J., Neves, J. and Analide, C. (2018), 'A multi-agent system blockchain for a smart city', *CYBER 2018 : The Third International Conference on Cyber-Technologies and Cyber-Systems* pp. 68–73.

Faraj, S. and Johnson, S. L. (2011), ‘Network exchange patterns in online communities’, *Organization Science* **22**(6), 1464–1480. doi: 10.1287/orsc.1100.0600.

Frey, B. S. and Jegen, R. (2001), ‘Motivation crowding theory’, *Journal of Economic Surveys* **15**(5), 589–611. doi: 10.1111/1467-6419.00150.

Gabaix, X. (2009), ‘Power laws in economics and finance’, *Annual Review of Economics* **1**(1), 255–294. doi: 10.1146/annurev.economics.050708.142940.

Grimm, V., a. B. U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J. and et al. (2020), ‘The odd protocol for describing agent-based and other simulation models: A second update.’, *Journal of Artificial Societies and Social Simulation (JASSS)* **23(2)**(7). doi: 10.18564/jasss.4259.

Gudmundsson, J., Hougaard, J. L. and Platz, T. T. (2023), ‘Decentralized task coordination.’, *European Journal of Operational Research* **304**(2), 851–864. doi: 10.1016/j.ejor.2022.04.026.

Gutierrez, C., Obloj, T. and Zenger, T. (2025), ‘Pay transparency and productivity’, *Strategic Management Journal* **46**(8), 1831–1860. doi: 10.1002/smj.3707.

Halfaker, A., Geiger, R. S., Morgan, J. T. and Riedl, J. (2013), ‘The rise and decline of an open collaboration system: How wikipedia’s reaction to popularity is causing its decline’, *American Behavioral Scientist* **57**(5), 664–688. doi: 10.1177/0002764212469365.

Holcomb, T. R. and Hitt, M. A. (2007), ‘Toward a model of strategic outsourcing’, *Journal of Operations Management* **25**(2), 464–481. doi: 10.1016/j.jom.2006.05.003.

Hsieh, Y.-Y., Vergne, J.-P., Anderson, P., Lakhani, K. and Reitzig, M. (2018), ‘Bitcoin and the rise of decentralized autonomous organizations’, *Journal of Organization Design* **7**(1), 1–16. doi: 10.1186/s41469-018-0038-1.

Iansiti, M., Lakhani, K. R. et al. (2017), ‘The truth about blockchain’, *Harvard Business Review* **95**(1), 118–127.

Johnson, S. L., Faraj, S. and Kudaravalli, S. (2014), ‘Emergence of power laws in online communities: The role of social mechanisms and preferential attachment’, *MIS Quarterly* **38**(3), 795–A13. doi: 10.25300/MISQ/2014/38.3.08.

Khern-am nuai, W., Kannan, K. and Ghasemkhani, H. (2018), ‘Extrinsic versus intrinsic rewards for contributing reviews in an online platform’, *Information Systems Research* **29**(4), 871–892. doi: 10.1287/isre.2017.0750.

Krishnan, A. (2020), ‘Blockchain empowers social resistance and terrorism through decentralized autonomous organizations.’, *Journal of Strategic Security* **13**(1), 41–58. url: https://www.jstor.org/stable/26907412.

Lakhani, K. (2016), *Managing communities and contests to innovate with crowds. Revolutionizing Innovation: Users, Communities, and Open Innovation*, MIT Press, Cambridge, MA. doi:10.7551/mitpress/9439.003.0010.

Leiblein, M. J. and Miller, D. J. (2003), 'An empirical examination of transaction-and firm-level influences on the vertical boundaries of the firm', *Strategic Management Journal* **24**(9), 839–859. doi: 10.1002/smj.340.

Lumineau, F., Wang, W. and Schilke, O. (2021), 'Blockchain governance—a new way of organizing collaborations?', *Organization Science* **32**(2), 500–521. doi: 10.1287/orsc.2020.1379.

Mauch, M., MacCallum, R. M., Levy, M. and Leroi, A. M. (2015), 'The evolution of popular music: Usa 1960–2010', *Royal Society Open Science* **2**(5), 150081. doi: 10.1098/rsos.150081.

Nakamoto, S. (2008), 'Bitcoin: A peer-to-peer electronic cash system', *Decentralized Business Review* p. 21260. url: https://assets.pubpub.org/d8wct41f/31611263538139.pdf.

Nguyen, L. K. N., Howick, S. and Megiddo, I. (2024), 'A framework for conceptualising hybrid system dynamics and agent-based simulation models.', *European Journal of Operational Research* **315**(3), 1153–1166. doi: 10.1016/j.ejor.2024.01.027.

Oliveira, F. (2023), 'The emergence of social inequality: A co-evolutionary analysis.', *Journal of Economic Behavior & Organization* **215**, 192–206. doi: 10.1016/j.jebo.2023.06.003.

Oliveira, F. and Costa, M. (2018), 'Capacity expansion under uncertainty in an oligopoly using indirect reinforcement-learning', *European Journal of Operational Research* **267**(3), 1039–1050. doi: 10.1016/j.ejor.2017.11.013.

Ponta, L., Puliga, G., Lazzarotti, V., Manzini, R. and Cincotti, S. (2023), 'To copatent or not to copatent: An agent-based model for firms facing this dilemma.', *European Journal of Operational Research* **306**(3), 1349–1363. doi: 10.1016/j.ejor.2022.07.003.

Roberts, J. A., Hann, I.-H. and Slaughter, S. A. (2006), 'Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects', *Management Science* **52**(7), 984–999. doi: 10.1287/mnsc.1060.0554.

Rosa, E., D'Angelo, G. and Ferretti, S. (2019), 'Agent-based simulation of blockchains', *Proceedings of the 19-th Asia Simulation Conference* . url: arxiv.org/pdf/1908.11811.

Sawyer, R. (2001), 'Emergence in sociology: Contemporary philosophy of mind and some implications for sociological theory.', *American Journal of Sociology* **107**(3), 551–585. doi: 10.1086/338780.

Sutton, R. and Barto, A. (1998), *Reinforcement-learning: An introduction.*, MIT Press. ISBN: 0-262-19398-1.

Von Krogh, G., Haefliger, S., Spaeth, S. and Wallin, M. W. (2012), 'Carrots and rainbows: Motivation and social practice in open source software development', *MIS Quarterly* pp. 649–676. doi: 10.2307/41703471.

Voshmgir, S. (2020), *Token economy: How the Web3 reinvents the internet*, Vol. 2, Token Kitchen, Berlin. ISBN: 9783982103815.

Wei, X., Chen, W. and Zhu, K. (2015), Motivating user contributions in online knowledge communities: virtual rewards and reputation, *in* '2015 48th Hawaii International Conference on System Sciences', IEEE, pp. 3760–3769. doi: 10.1109/HICSS.2015.452.

West, J. and Kuk, G. (2016), 'The complementarity of openness: How makerbot leveraged thingiverse in 3d printing', *Technological Forecasting and Social Change* **102**, 169–181. doi: 10.1016/j.techfore.2015.07.025.

Yu, Y., Khern-am nuai, W. and Pinsonneault, A. (2022), 'When paying for reviews pays off: The case of performance-contingent monetary rewards.', *MIS Quarterly* **46**(1). doi: 10.25300/MISQ/2022/15488.

Zichichi, M., Contu, M., Ferretti, S. and D'Angelo, G. (2019), 'Likestarter: a smart-contract based social dao for crowdfunding', *Proceedings of the 2st Workshop on Cryptocurrencies and Blockchains for Distributed Systems (CryBlock'19)* . doi: 10.48550/arXiv.1905.05560.

Appendix A. ODD Protocol: ABM for the DAO governance

1. Overview

Purpose

Blockchain-based decentralized autonomous organizations (DAOs) have the potential to innovate existing open communities, where people share and evaluate creative works, through transparent and collectively determined reward mechanisms. Our model simulates how different governance structures of DAOs shape participant engagement level, reward allocation, reputation distribution, and creative content production in open communities. These dimensions are collectively considered to evaluate the long-run evolution of the community from multiple angles. We use the traditional open community model as a baseline to assess the viability of DAO governance structures.

Entities, State Variables, and Scales

Entities: Each agent can act as both an author and a rater, reflecting real-world communities where participants engage in both content creation and evaluation. Authors are content creators who produce designs and adapt them over time in response to feedback. Raters are evaluators who assess authors' outputs based on individual preferences and thresholds. The environment represents the community's governance setting, defining the rules by which rewards for productions and ratings are updated.

State Variables: Table A1 provides an overview of the state variables for authors and raters, including their types, ranges, and whether they are dynamic or static. Detailed descriptions of the state variables and their interactions are provided in Section 4.

Table A1 State Variables of Authors (i) and Raters (j)

Variable	Description	Type	Range	Dynamic/Static
β_i	Search radius	Continuous (float)	[0, 1]	Dynamic in DAO; Static in Open Community
P_i^p	Probability of exploring new product types	Probability (float)	[0, 1]	Dynamic in DAO; Static in Open Community
Q_i^p	Number of products scanned while exploration	Integer	≥ 0	Dynamic in DAO; Static in Open Community
s_i	Design decision	Continuous (float)	$[-1, 1]$	Dynamic
w_j	Preference	Continuous (float)	$[-1, 1]$	Static
a_j	Rating threshold	Continuous (float)	[0, 1]	Static
P_j^r	Probability of rating participation	Probability (float)	[0, 1]	Dynamic in DAO; Static in Open Community
Q_j^r	Number of products rated	Integer	≥ 0	Dynamic in DAO; Static in Open Community
r_{ji}	Rating	Continuous (float)	[0, 1]	Dynamic
l_{ji}	Rating decision	Binary	$\{0, 1\}$	Dynamic

Process Overview

Comprehensive details of the process are described in Table 4 and Section 5. Briefly, each iteration includes:

1. Update Production Reputations (η_i) and Rating Reputations (ϕ_j) simultaneously via reinforcement learning
2. Distribute rewards, Rating Rights, and Voting Rights based on governance rules
3. Update production behavior (local search characterized by β_i, P_i^p, and Q_i^p; and design decision s_i) and rating behavior (scanning characterized by P_j^r and Q_j^r; and rating decision r_{ji} and l_{ji})

2. Design Concepts

Basic Principles

Reinforcement learning–based evolution of a system is a well-developed concept in evolutionary computation and has been applied in various contexts (see Introduction for examples). This study is a novel application of this concept to DAO governance for a creativity-driven community. DAOs are an emerging but nascent form of organization whose governance effects have been rarely studied. We simulate the consequences of different DAO governance structures in terms of long-run community engagement, emergent reputation, and exploratory learning. In the simulation model, agents adaptively interact through rating and production. DAO governance structures mediate these interactions through reward mechanisms and the distribution of rating and voting rights.

Emergence

The key outcomes of the model are emerging patterns of engagement rate, reputation distributions for production and rating, and creative diversity at the community level driven by individual adaptive strategies.

Adaptation

Agents alter their behaviors based on feedback from reward tokens. As authors, they update production behavior, including the probability of making any improvement (P_i^p), sample size for exploring new designs (Q_i^p), the radius for local search (β_i), and their design (s_i). As raters, they update rating behavior, including the probability of providing a rating (P_j^r), rating sample size (Q_j^r), and their rating decisions (r_{ji} and l_{ji}).

Objectives

Agents incrementally adapt their production design (s_i) towards a direction that they expect to get higher ratings in the future. Rating behavior does not have a particular objective because preference (w_j) is fixed, while token rewards motivate more active engagement in rating (P_j^r and Q_j^r).

Learning

We implement a reinforcement learning mechanism to update agents' reputations and subsequent behavior in the DAO. This represents a creative application of a mechanism that has been used in other contexts (e.g., Bunn and Oliveira 2008, Didden et al. 2024). Each agent updates its production and rating reputations by taking previously undertaken actions as reference points. New ratings are incorporated through updating, while the previously established reputation state is given greater weight than the new rating, reflecting path-dependent evolution.

Prediction

When authors decide whether and how to update their design, they make predictions by estimating the expected ratings of a sampled set of new designs that deviate positively or negatively from their current design. Based on these predictions, they choose whether to update their design and the direction of the update that is expected to yield the highest rating. Details of the process are presented in Section 4.1.2.

Sensing

In DAOs, in principle, all agents can observe products, ratings, reputations, and governance rules. However, given each agent's limited cognitive capacity and level of engagement, the information to which agents pay attention when changing their behavior is limited. Authors pay attention to different product designs within the scope of their local search radius (β_i) and sample size (Q_i^p), and raters pay attention to products within the scope of their sample size (Q_j^r). These variables are updated in every iteration (see Table 4 for details).

Interaction

The key interactions between raters and authors are as follows: raters evaluate authors' designs, and authors update their production behavior in response to these evaluations. These interactions are local, as authors learn and produce within a radius of their current design. The interactions are mediated by the token-based reward mechanism.

Stochasticity

Rating threshold (a_j) is initialized randomly to process the threshold without causing variability.

Collectives

A DAO, which is a community of agents, exhibits evolving emergent properties that arise from individual agents' behavior. These community-level properties feed back into agents' behavior. As the community-level properties are emergent, they are not explicitly represented by variables in the model. Community-level outcomes (e.g., engagement rate, reputation concentration, and design diversity) are computed as aggregate statistics at each iteration and indirectly influence subsequent agent behavior through the reward mechanism.

Observation

Model outputs include an overall level of engagement, production and rating reputation distributions, design diversity and change, and Voting Rights distributions (governance decision rights).

3. Details

Initialization

- 100 agents are initialized in terms of design (s_i), preferences (w_j), and rating thresholds (a_j).
- Initial ratings are computed based on the initialized state variables.
- Reputations (η_i, ϕ_j) start uniform at $1/I$

Input Data

The model does not incorporate time-varying exogenous input data. All dynamics arise endogenously from agent interaction, reinforcement learning, and governance-mediated reward allocation.

Submodels

(1) Rating: Each rater j possesses a unique threshold denoted as a_j. When the rating surpasses a_j, rater j extends support to the author and offers a support (l_{ji}); otherwise, no support is expressed.

$$r_{ji} = \begin{cases} \frac{2-\delta(s_i - w_j)^2}{2}, & \text{if rating } \geq a_j \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

$\delta > 0$ embodies the decay rate, signifying how swiftly r_{ji} diminishes as s_i deviates from w_j. The calibration of $\delta > 0$ is explained in Appendix B.

(2) Production Reputation: The greater the number of supports received by author i ($l_{\cdot i}$), the higher the i's Production Reputation (η_i). The reputation accumulated at an early stage carries greater weight than support received after a robust reputation has been established. This dynamic is reflected in the update of Production Reputation (η_i) through a reinforcement learning algorithm (e.g., Sutton and Barto 1998).

$$\eta_i^{t+1} = \frac{t}{t+1}\eta_i^t + \frac{1}{t+1}\left(\frac{l_{\cdot i}}{\sum_k l_{\cdot k}}\right) \tag{2}$$

(3) Production Design Learning: The author i's learning process involves a local search and decision to adjust their current design (s_i). $\mu, \nu \sim U[0,1]$ serve as stochastic exploration parameters guiding the author's local search within a specified radius $0 \leq \beta_i \leq 1$. The local search identifies the design adjusted negatively (s_i^-) and the design adjusted positively (s_i^+).

$$s_i^+ = \min(s_i(1+\beta_i\mu), 1) \tag{3-1}$$

$$s_i^- = \max(s_i(1-\beta_i\nu), -1) \tag{3-2}$$

Author i samples raters and computes their expected ratings to pursue the design change yielding the highest expected rating.

(4) Rating Contribution: Rater j's contribution to DAO is its endorsements of authors possessing commendable Production Reputation (η_i). As rater j assigns a higher rate to authors bearing higher Production Reputation, j's Rating Contribution (ξ_j) increases. Reflecting this logic, ξ_j is computed as the weighted mean of the number of Rating Rights accorded by rater j to each author, relative to the weighted mean of Rating Rights prevalent in the DAO.

$$\xi_j = \frac{\sum_i \eta_i \cdot l_{ji}/l_{\cdot i}}{\sum_k \sum_i \eta_i \cdot l_{ki}/l_{\cdot i}} \tag{4}$$

(5) Rating Reputation: Based on the Rating Contribution (ξ_j), a rater j's Rating Reputation (ϕ_j) is updated using a reinforcement learning algorithm (e.g., Sutton and Barto 1998). Similar to Production Reputation, equation (5) reflects that ratings undertaken at an early stage carry greater significance than those conducted after a robust reputation has been established.

$$\phi_j^{t+1} = \frac{t}{t+1}\phi_j^t + \frac{1}{t+1}\xi_j \tag{5}$$

(6) Voting Rights: In the Merit-based DAO, Voting Rights are assigned proportionally to the combined weight of Rating Reputation and Production Reputation. The parameter σ indicates how the relative importance of production and rating reputations influences voting rights.

$$\text{Voting Rights} = (1-\sigma) \times \text{Rating Reputation} + \sigma \times \text{Production Reputation} \tag{6}$$

Appendix B.

To reproduce stylized empirical features of the real-world open communities in our baseline model, we calibrated the rating decay parameter (δ) using two criteria. First, the interaction between δ, author designs (s), and rater preferences (w) should generate a heavy-tailed distribution of Production Reputation consistent with empirical patterns observed in Figure 1, given the initial distribution of preferences described in Section 6. Second, δ should be sufficiently large to enforce selective evaluation, such that raters assign positive ratings primarily to designs that are close to their own preferences (Equation 1).

We explored values of δ across a range of magnitudes. As illustrated in Figure B1, values of δ below approximately 100 produce diffuse rating patterns that fail to generate concentrated reputation distributions. For δ above this threshold, rating behavior becomes increasingly selective and the resulting reputation distribution exhibits heavy-tailed characteristics. We therefore select $\delta = 1000$, which lies well within this selective regime and produces stable dynamics consistent with the empirical stylized facts described in Section 6.1.

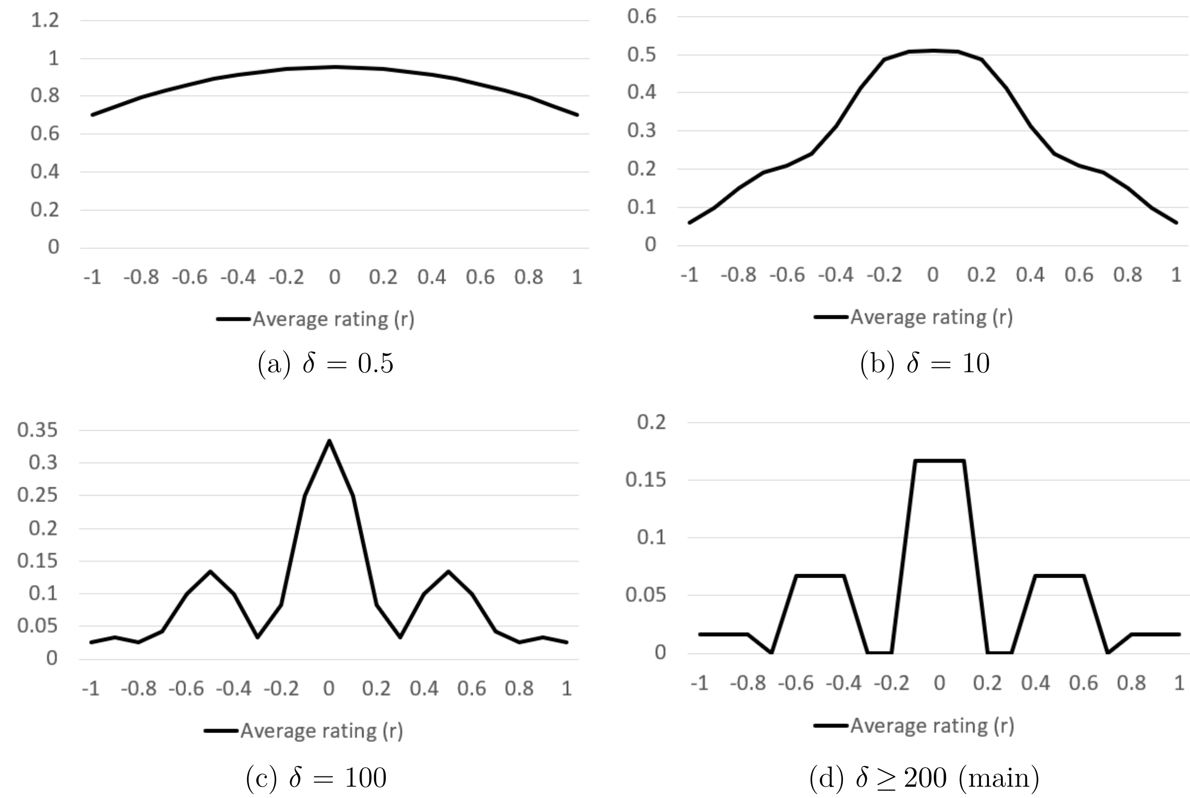

(a) $\delta = 0.5$ (b) $\delta = 10$

(c) $\delta = 100$ (d) $\delta \geq 200$ (main)

Figure B1 Average rating (r) distributions at different rating decay rate (δ)

Note. (a) $\delta = 0.5$: close to a uniform distribution; (b) $\delta = 10$: a unimodal distribution with higher ratings for the central design; (c) $\delta = 100$: a multimodal distribution with peaked ratings decreasing quickly when moving away from the local maxima; and (d) $\delta \geq 200$: a multimodal distribution with flat tops, showing tolerance to deviations from the local optima.

Appendix C.

As part of our sensitivity analyses, we consider the context of an engineering-oriented community characterized by initial designs converging around the best-performing technical solution and less diverse preferences (Figure C1), instead of diversified designs and preferences in an art-oriented community. In this scenario, the merit-based DAO still achieves the highest long-term engagement levels (Figure C2), and the property-based DAO continues to present a viable alternative fostering learning (Figure C3), showing that the main qualitative results remain robust. Across both art-oriented and engineering-oriented initial conditions, the relative ranking of governance structures remains unchanged, indicating robustness of comparative results to preference dispersion.

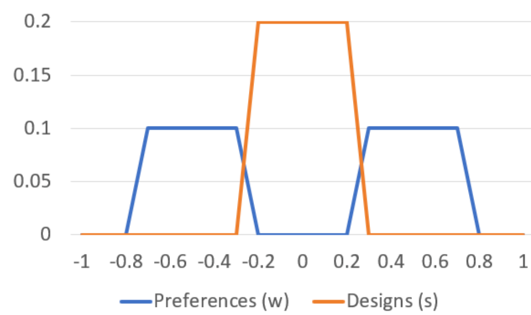

Figure C1 Distribution of Initial Designs (s_i) and Preferences (w_j)

Note. X-axis indicates the range of designs (s_i) and preferences (w_j), both are within -1 and 1. The Y-axis denotes the proportion of agents in each range of s_i and w_j.

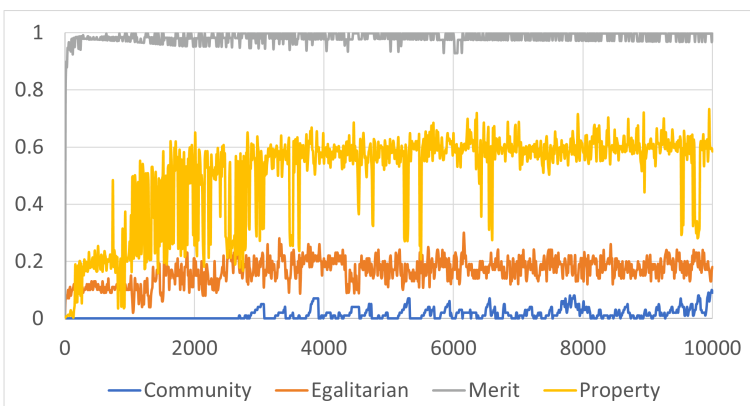

Figure C2 Overall Engagement in Rating (proportion of Rating Rights used)

Note. X-axis indicates the number of iterations.

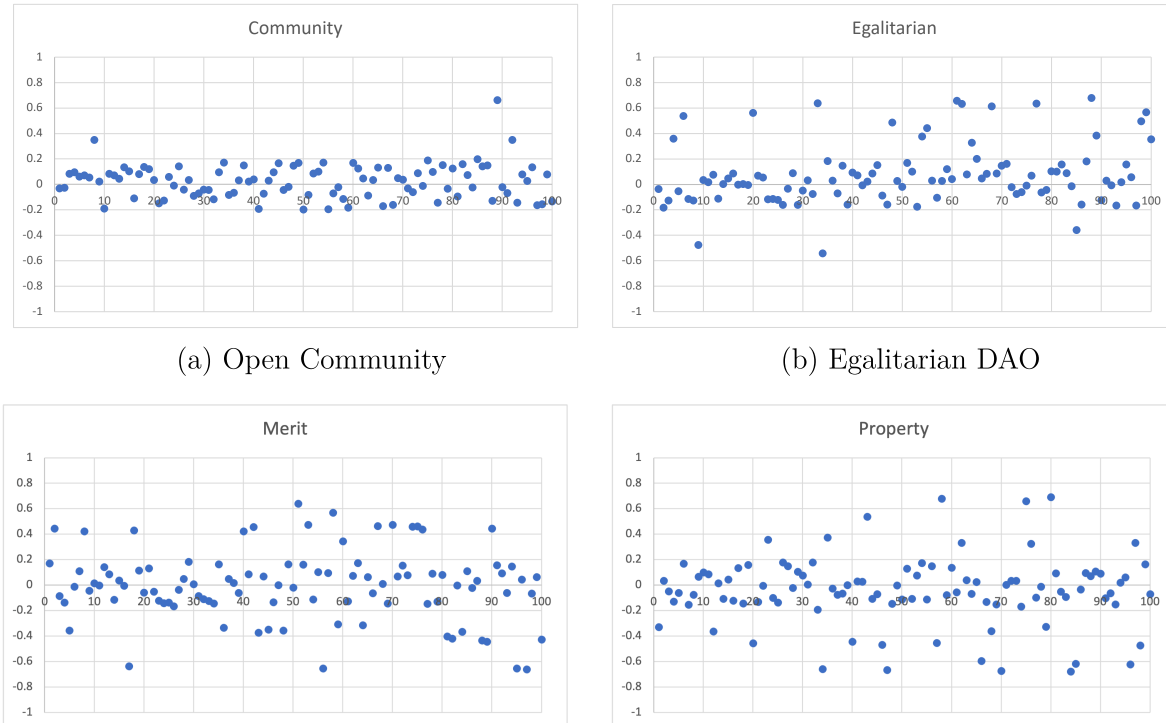

(a) Open Community

(b) Egalitarian DAO

(c) Merit-based DAO

(d) Property-based DAO

Figure C3 Designs (Y-axis) Under Different Governance Structures (Converged State)

Note. Each dot indicates an agent.The X-axis indicates an agent-specific ID in our simulation ($I = 100$).