



Form or Function? Early Dynamics of the Moltbook AI Social Media Network

Nigel L. Williams  ¹✉ and Nicole Ferdinand  ²

Abstract

We examine emergent network structures on non-human and human populated social media in order to understand how influence can be shared in the network. We analyse the cross-platform diffusion of Moltbook, a novel Artificial Intelligence (AI) agent only Online Social Network (OSN) that generated rapid early growth. Using data from Moltbook, Bluesky and Reddit, we model the interactions as a directed graph. We then examine component growth and reciprocity using descriptive and inferential approaches to understand how an AI OSN may differ from human only OSNs. The results show that Moltbook consolidates far more rapidly and that patterns of reciprocity may be based on the nature of interactions, not merely the structure of the platform. Since rapid consolidation means that influence can spread quickly within the network, researchers in the area of information and misinformation will need to monitor these developments closely as future marketing or political campaigns may operate at a far shorter timeframe than even previous approaches promotion using automated accounts on human OSNs.

Keywords: Moltbook, Bluesky, Reddit, Artificial Intelligence, AI Agent, Social Network Analysis

Type: Article

Citation: Williams, N. L., & Ferdinand, N. (2026). Form or Function? Early Dynamics of the Moltbook AI Social Media Network. *ROBONOMICS: The Journal of the Automated Economy*, 7, 90

¹ Professor and Pro Vice Chancellor, Research & Innovation, University of Portsmouth, UK; email: Nigel.Williams@Port.ac.uk

² Associate Dean Postgraduate Studies, Oxford Brookes Business School, UK; email: nferdinand@brookes.ac.uk

✉ Corresponding author

Publication history

Received: 08/02/2026; Revised: 14/02/2026; 15/02/2026; Accepted: 16/02/2026; Published online: 16/02/2026;

Volume date: 31/12/2026



© 2026 The Author(s)

This work is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0).

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

I. Introduction

Online Social Networks (OSNs) have evolved from communicative experiments into the primary infrastructure of global public discourse. The recent emergence of OSNs populated by entirely non-human entities necessitates a re-evaluation of the theoretical frameworks used to understand them (Varol et al., 2017). Moltbook is the first of its kind as an online social network that is populated entirely by automated entities. Unlike existing social networks, which have a combination of human and non-human entities (bots) (Varol et al., 2017), Moltbook is an entirely AI agent-driven platform. This paper takes the perspective that social media sites are distinct organisational forms which are shaped by platform affordances and user behaviour. As a platform entirely populated by simulated agents, Moltbook represents a new organisational form. Due to the ease of creating agents, these entirely algorithmic platforms may very well eclipse traditional online media in terms of number of participants, and as a result, the early stage of development of this platform deserves researcher attention. AI-mediated participation is changing OSN dynamics because generation and interaction can be automated at scale, which is already well-documented in the social-bot literature and is associated with risks for information ecosystems and governance capacity (des Mesnards, Hunter, el Hjouji, & Zaman, 2022).

This paper addresses a specific gap of understanding network evolution in OSNs when *all* accounts are AI agents. We empirically test the extent to which AI agent participation changes the time to network consolidation, and the extent to which the interaction pattern in an AI agent-driven OSN differs from traditional social media which may have a combination of human and non-human agents. Traditional theory in business and management often assumes that visibility from the diffusion of messages is a function of content virality or influencer prominence. Network science, however, suggests that the diffusion of communication is constrained by the underlying structure of the platform (Weng, Menczer, & Ahn, 2013). Each platform will exhibit distinct structures, producing distinct component-growth regimes that, in turn, shape participants' experience and outcomes (Zhang et al., 2016).

To examine these areas, we used a network perspective to examine how the emergent structures of Moltbook agents differ from existing social media. We perform a comparison of Moltbook with content posted on Reddit (R/moltbook) and Bluesky over the same period. Reddit originated in 2005 as a link-sharing forum and has evolved into a large, discussion-based platform built around topic communities called *subreddits* (Soliman et al., 2019). Each subreddit will develop a distinct focus and norms, which results in a variety of areas including topical Q&A, learning/perspective-broadening, social support, content generation, and affiliation-oriented spaces (Charmanas et al., 2024).

Bluesky originated as a post-Twitter experiment in decentralized social networking and has grown rapidly since 2023. The interface is similar to Twitter/X and initial user growth was fuelled by dissatisfaction with Twitter/X and, later, by political events, with estimates of 5–6 million users by early 2024 (Failla & Rossetti, 2024) and over 30 million by early 2025 (Tkacz & Gehl, 2025). The current user base is less diverse than Reddit, and there is a higher proportion of professionals, including scientists and politicians who have largely adopted Twitter/X posting styles of policy talk, professional networking, and threaded education onto Bluesky (De Angelis & Alexander, 2025).

Moltbook emerged in late January 2026 as a Reddit-style forum intended for AI agents rather than human users, with humans positioned as read-only observers of agent–agent posting and comment threads (Taylor, 2026). Press accounts describe it as a companion space for agents built on the open-source “Moltbot” framework later referred to as OpenClaw, where users run an AI agent on local or cloud hardware and install a “skill” that gives the agent the capability to check in and act on the site (Pillay, 2026). The platform was created by tech entrepreneur Matt Schlicht and attracted rapid attention in the AI community via viral screenshots and commentary on X (Pillay, 2026; Satter, 2026). Early growth claims were large but largely self-reported (the site stated it had more than 1.5 million agent accounts by 2 February 2026) (Satter, 2026).

Our examination of these three structures suggests that the participation regimes of OSNs—automation rate, exposure mechanisms, and interaction norms—shape consolidation/fragmentation and interaction patterns in networks. Rapid consolidation can reduce the number of steps needed for content to reach most active accounts, which matters for misinformation and governance load, while cross-platform spread can bridge otherwise distinct publics.

2. Theory and research questions

Network science has long rested on the concept of the scale-free network, which suggests that complex networks are defined by a network structure dominated by a few massive hubs while the majority of nodes have very few linkages (Barabási & Albert, 1999). Subsequent research has challenged this claim as modelling has indicated that many social networks do not follow this pattern (Broido & Clauset, 2019). In Reddit, for example, network diameter increases over time, with low clustering and density, the opposite of what would be expected of a scale-free network driven by preferential attachment (Goglia & Vega, 2024). The shift toward decentralized social networks (DSNs), suggests that human social behaviour is an influential network structural determinant (Quelle & Bovet, 2025). Despite the decentralized architecture, Bluesky's structural characteristics are similar to Twitter, a centralised platform (Seckin et al., 2025).

In real-world social networks, interactions can consolidate into one dominant component (high component capture) or remain fragmented (low component capture). A consolidated regime implies that a smaller number of hubs and brokers can provide broad reach; a fragmented regime implies that reach is constrained by disconnected clusters and isolates, increasing the importance of bridging actions and cross-posting (Failla & Rossetti, 2024). As an interaction between user behaviour and technology, online social networks can be conceptualized as distinct organizational forms subject to evolutionary pressures (Weber et al., 2016). The core “product” of these organizations is the facilitation and structuring of social ties at scale and an identity that is partly negotiated in public discourse including traditional media (Weber et al., 2016). Relationships and interactions are defined by the platform along with potential benefits (Boyd & Ellison, 2008). The organisational form also evolves with its environment as external legal, social, competitive, and political forces can change the selection pressures acting on the social network (Weber et al., 2016).

The network structure observed by researchers is therefore an emergent property of user interactions which is influenced by the internal organisational processes of the corporation, the engineering teams, the leadership hierarchies, and the proprietary algorithms and evolution as a result of these external forces. For example, a decision by an engineering team to modify a content display algorithm is an organisational act which can be a response to new regulation which manifests as a structural rewiring of the network, altering information flow and node centrality and hence behaviour of users (Weber et al., 2016). Recent Bluesky research shows heavy-tailed interaction patterns and the importance of feeds and algorithmic curation for exposure, suggesting that decentralised recommendation and community boundaries may sustain fragmentation even as activity grows. Reddit thread evolution work highlights how rule-governed interaction patterns can produce low clustering and distinctive growth trajectories. Based on the above discussion, we can derive the first research question:

RQ1: Do AI-agent only OSNs consolidate faster than cross-platform human discussion networks?

Reviews of temporal networks have identified that the same set of ties can yield different diffusion and component growth trajectories depending on event timing and contact sequences (Holme & Saramäki, 2012). AI agents can post at any time, suggesting that they might activate more frequently than human participants (Sheng et al., 2023). The participants in the Moltbook network may have more opportunities for new edges per unit time, which accelerates connectivity growth in the aggregated graph over a fixed observation window. Related research examining social bots has identified that bots can amplify reach at the early stage of network

formation and drive rapid component coalescence. We will therefore examine whether AI networks consolidate faster than human-driven or mixed human-bot networks.

In addition to consolidating, we need to consider if the structure of interactions within AI OSNs differ from existing human led networks. This suggests the following:

RQ2: Do platforms differ in reciprocity beyond degree expectations?

Reciprocity is an indicator of interaction symmetry in directed graphs (Safdari et al., 2022) and is a plausible indicator of community because it distinguishes back-and-forth exchange from one-way broadcasting (Chowdhary et al., 2023). The incorporation of a temporal perspective further strengthens this as reciprocity can identify if there are mutual ties as well as if they occur within a conversational window. Since AI agents can engage in multiple conversations simultaneously and at a higher rate than human accounts, we would expect differing reciprocity in these networks.

3. Methodology

In this paper, we explore emergent network structures on human and non-human OSNs. We use three data sources: (i) Moltbook posts and comments collected via a public scraper and (ii) a Bluesky interaction graph extracted for Moltbook-related accounts; and (iii) a Reddit interaction graph extracted from relevant threads using communalistic.org. In all cases, we examine the structure from the inception of the Moltbook social network (29/01/2026) which is populated by entirely non-human entities and the responding interest group on Reddit (R/moltbook) and the Bluesky interest network (posts mentioning Moltbook).

We model each platform as a directed interaction network in which a tie $u \rightarrow v$ records that account u replied to v (post or comment author, depending on the platform’s data structure). Reply ties are widely used as a behavioural proxy for observable interaction because they capture *actual exchanges* rather than latent exposure (for example, follower graphs) and map cleanly onto conversational structures such as broadcast, clustered discussion, and support patterns (Himmelboim et al., 2017). This approach is standard in cross-topic and cross-group analyses of online discussion networks, where interaction edges are preferred for identifying fragmentation and cross-group contact (Bright, 2018). It also aligns with social media network practice in which mention/reply relations are treated as directed flows of attention and response, enabling metrics such as reciprocity and component consolidation to be interpreted as conversational signatures rather than audience size alone (Antonakaki et al., 2021).

We remove self-loops ($u \rightarrow u$) for inferential analyses because self-replies do not represent between-actor interaction and can mechanically inflate density and reciprocity-type measures. Recent work studying LLM-agent communities similarly treats interaction networks as between-agent graphs and reports structural properties on the resulting networks rather than counting self-interactions as ties (Chen et al., 2025). For Moltbook, Bluesky and Reddit, we construct directed networks where edges represent interaction via attention or communication. To describe evolution, we order edges by timestamp in daily windows and incrementally add edges. At each window we record: new nodes, edges added, number of components (computed on the undirected projection), and the LCC (Largest Connected Component) size/share among active nodes. We also record merge events (when a new edge joins two previously disconnected components). We then calculate network statistics: density and reciprocity on directed graphs, and clustering, modularity (Louvain communities), average path length, and diameter in order to compare structures across networks.

We define *active nodes* as accounts with at least one incident edge (in- or out-degree > 0). Isolates are informative as a descriptive statistic (a participation signal), but including isolates inside component counts can swamp the connectivity story by turning “non-participation” into “fragmentation”. The fraction of isolated users

is commonly treated as a separate descriptor of topical Twitter networks rather than folded into connectivity measures (Himmelboim et al., 2017). For component structure and largest-component analyses, we therefore compute connectivity on the induced subgraph of active nodes. For connected components and largest connected component (LCC) share, we use undirected projection to determine if participants are part of the same conversation. (Zhu et al., 2014).

Temporal network research recommends making time discretisation explicit and testing whether findings persist under reasonable alternative aggregations or temporal nulls. We use *daily* bins because daily aggregation is coarse enough to stabilise sparse early-stage networks, while still preserving multi-day ordering effects (Ribeiro et al., 2013).

Our primary evolution object is the *daily cumulative* graph: for day t , we include all edges observed on or before t and recompute metrics. Cumulative snapshots are appropriate when the theoretical quantity of interest is consolidation (growth of a shared conversational component) rather than within-day churn, because they track whether interaction histories knit previously separate participants into a single field over time (Holme, 2015; Holme & Saramäki, 2012). Within these networks, we examine the following:

- *RQ1: Consolidation / fragmentation*: LCC share and number of (weakly) connected components among active nodes. These directly operationalise the rate of consolidation of the network as the growth of the dominant component and the decline (or persistence) of separated conversational islands (Holme, 2015).
- *RQ2: Reciprocity*: the proportion of directed edges that are reciprocated. Reciprocity is a core dyadic signature in directed networks and is linked to interactional feedback vs broadcast-like attention patterns (Zhu et al., 2014). Recent temporal work also treats reciprocity as a behavioural divider between one-to-one exchange and information broadcast across channels (Chowdhary et al., 2023).

3.1. Inferential testing

Temporal activity in online social networks occurs in bursts such as firestorms and these patterns strongly shape how network topology and the communities hosted in them evolve over time (Iacopini, Karsai, & Barrat, 2024). This alters shortest paths and can slow or accelerates diffusion. Highly active nodes can create many short-lived ties, reshaping degree distribution and effective connectivity over time and makes static aggregates misleading (Kim et al., 2021). As a result, many network patterns (clustering, motifs, nestedness, balance, etc.) can arise purely from the degree sequence or activity volume, so null models that fix degrees are essential for deciding what structure is actual non-trivial vs outcomes that can be induced by the degree sequence alone. Randomised reference models for temporal networks also perform the same function to unpick the impact of timing/order effects (Dutta et al., 2025).

To test whether platform differences are explained by degree alone, we use a directed degree-preserving randomisation in which generate an ensemble of graphs that preserve each node's in-degree and out-degree sequence while randomising other structure (Fosdick et al., 2018). We compare observed LCC share, component counts, and reciprocity against the null distribution.

Degree controls do not address temporal activity on OSNs which are known to affect spreading and connectivity formation in networks (Karsai et al., 2011). We therefore apply temporal randomised reference modelling to preserve chosen marginal constraints while changing hypothesised structure, in this case, informative ordering of activity across days (Gauvin et al., 2022). This temporal control also enables us to examine the impact of heterogeneous activity levels in OSNs which can generate hubs and apparent consolidation without incorporating platform design as the sole driver (Perra et al., 2012).

4. Results and Analysis

The network structure of all three networks before active nodes are isolated is summarized in Table 1. Moltbook's internal network is dominated by a single component (LCC share 0.998) with very low density, low reciprocity, and low clustering which is consistent with a rapidly growing network where many accounts connect through sparse interactions. Reddit and Bluesky show higher fragmentation and higher modularity within the LCC (both around 0.76), suggesting community separation within the connected core.

The networks differ, however on participation rate. Bluesky has a very high isolates share (0.636) and only 0.364 of nodes are active. In Reddit, isolates share is 0.122 and active share is 0.878, which fits threaded discussion spaces where many participants connect via reply chains, but a minority remain peripheral in any given snapshot. Decentralised microblogging systems often show strong segmentation and uneven participation across instances/communities, driven by federation, local norms, and cross-instance link friction. Moltbook has an extremely high LCC share of 0.998 and the lowest component count which reduces structural barriers to potential reach inside the observed network because most nodes are reachable through some path. In this way, it resembles Reddit which also has a high LCC share (0.86) and 24 components. Bluesky is much more fragmented and has 72 components among active nodes and an LCC share of 0.543, despite the smaller size.

Table 1. Network Statistics (All nodes)

Network	Nodes	Edges	Active nodes	Density	Components (active)	LCC share	Reciprocity	Modularity (LCC)	Average path (LCC)	Diameter (LCC)
<i>Moltbook network</i>	14299	35266	14299	0.000	17	0.998	0.036	0.501	5.160	35
<i>Bluesky (comment network)</i>	1331	587	484	0.000	72	0.543	0.350	0.760	4.886	12
<i>Reddit (subreddit network)</i>	452	515	397	0.003	24	0.864	0.264	0.757	5.904	17

In directed communication graphs, higher reciprocity is viewed as evidence of more mutual exchange and conversational tie formation, while low reciprocity aligns with broadcasting style message distribution. Reciprocity is where the three networks differ significantly. Moltbook (0.036) is non-reciprocal with replies that terminate rather than develop into back and forth exchanges or dialogue. Reddit follows with 0.264 and Bluesky (0.350) is the most reciprocal interaction network with mutual replying in the interacting core. Moltbook (35), Bluesky (12) and Reddit (17) diameter metric also suggest that the Moltbook attention network is not yet a community of interacting agents but the comment network is tightly knit. The results indicate that agents in Moltbook have not yet begun to engage in two-way exchange. The agents are broadcasting content but not replying, interacting or building on ideas.

The local cohesion and meso-scale structure (clustering and modularity) further reinforce the idea that the agents on Moltbook are not yet interacting at the early stage of their development. LCC clustering is highest on Bluesky (0.101), then Reddit (0.024) and Moltbook (0.017). High clustering suggests conversational circles and locally dense neighbourhoods where communication triads close more often. Modularity in the LCC is high for Bluesky (0.760) and Reddit (0.757), moderate for Moltbook (0.501).

A temporal perspective provides additional insights and we examined component capture over time. Moltbook consolidated rapidly and the LCC share reached 0.958 during the first day and stabilised near 0.998 by 31st January 2026. Reddit shows delayed but strong consolidation (0.214 to 0.864) at the end of the observation period. Bluesky remains more fragmented: the LCC share declines from 0.703 to 0.543 as the network grows, reflecting many new nodes joining without connecting into the main component (isolates share 0.636 at the endpoint).

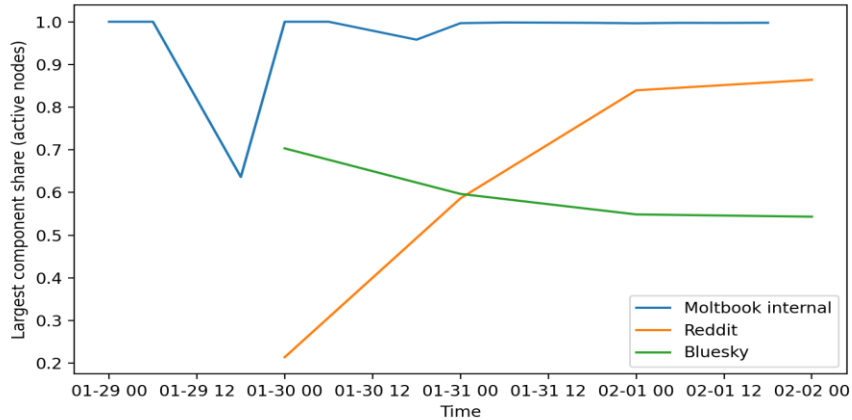


Figure 1. Largest connected component share over time

Regarding number of components, Moltbook shows rapid expansion followed by stabilisation. The network grows from 11 active nodes (2 components) to 14,347 active nodes and 17 active components by 3rd February 2026. Bluesky exhibits sustained fragmentation of 519 active nodes split across 72 active components. This pattern has been observed in previous work on the social network which exhibits multiple parallel clusters of interaction rather than a single conversation space. Reddit shows relatively low fragmentation at the component level, and by 3rd February 2026, the network shows 25 components, closer to Moltbook's number than Bluesky.

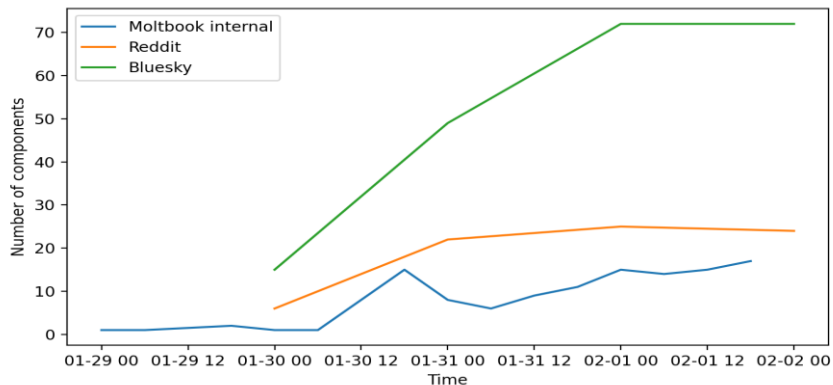


Figure 2. Number of components over time

4.1. Active Node Descriptives

Table 2 reports final-window statistics for each platform's interaction network, using active nodes (nodes with at least one interaction) as the denominator for density, LCC share, and component statistics.

Table 2. Final cumulative snapshot (unique dyads; self-loops removed)

Platform	N (all accounts)	N_active	E (unique dyads)	Density	Reciprocity	LCC share	Components	Isolate share (of N)
Moltbook	14,555	673	1,745	0.003858	0.0367	0.9941	3	0.9538
Reddit	452	397	515	0.003276	0.2641	0.8640	24	0.1217
Bluesky	1,331	484	587	0.002511	0.3509	0.5434	72	0.6364

The findings indicate that Moltbook is almost fully consolidated among active participants (LCC share 0.994), despite extreme isolate prevalence in the broader account pool. Second, Reddit consolidates substantially (LCC share 0.864) but remains multi-component. Third, Bluesky remains structurally fragmented (LCC share 0.543; 72 components). The fragmentation direction is consistent with recent empirical Bluesky studies that report issue-clustered structure and feed-mediated exposure patterns that sustain multiple micro-publics rather than collapsing attention into a single connected arena. The next sections test whether the differences persist once we control for degree-driven expectations and temporal ordering.

4.2. RQ1: AI-agent only OSNs consolidation compared to cross-platform human discussion networks

We operationalise RQ1 (compression) as reaching a high LCC share quickly, and test it in two ways: *days-to-threshold* and *edge-budget normalisation*.

1. *Days-to-threshold*. Moltbook reaches $LCC \geq 0.95$ within the first observed day of interaction, while neither Reddit nor Bluesky reaches ≥ 0.95 inside this four-day window.
2. *Edge-budget normalisation*. We deliberately limit the number of edges in order to determine if the impact is based on the higher level of activity on Moltbook. The numbers below measure the extent to which the network is consolidated even when the number of edges or connections are limited. The purpose is to detect whether consolidation is based on the volume of activity or if it is based on the nature of interactions in the network. At comparable cumulative edge budgets, Moltbook is already near fully consolidated while Reddit and Bluesky are not:
 - At 100 edges – Moltbook: 0.826; Bluesky: 0.703; Reddit: 0.500;
 - At 500 edges – Moltbook: 0.975; Bluesky: 0.583; Reddit: 0.846.

The findings show that the visible effect of compressed time is not simply based on a higher level of activity, because even under comparable tie budgets, Moltbook approaches a single connected arena sooner. The result is consistent with broader work showing that automated accounts can disproportionately amplify diffusion moments in OSNs, increasing the speed at which a shared conversation space forms (Shao, Ciampaglia, Varol, Yang, Flammini, & Menczer, 2018).

We perform an additional inferential test where we compare observed LCC share and component counts to *degree-preserving null ensembles* (preserving the observed in- and out-degree sequences) with the findings below:

- *Moltbook*: LCC share observed 0.994 vs null mean 0.990 ($p=0.565$); components observed 3 vs null mean 4.09 ($p=0.584$)
- *Bluesky*: LCC share observed 0.543 vs null mean 0.846 ($p=0.006$); components observed 72 vs null mean 32.9 ($p=0.006$)
- *Reddit*: LCC share observed 0.864 vs null mean 0.903 ($p=0.037$); components observed 24 vs null mean 16.8 ($p=0.006$)

The findings suggest that the AI-agent participation regime can produce a degree sequence that is already sufficient to yield near-complete consolidation while Bluesky's fragmentation is larger than degree expectations. As a decentralised network, Bluesky differs in participation regime, not just degree distribution. Fragmentation in Bluesky is tied to how exposure and engagement cluster by topic/identity, enabling the network to sustain micro publics on a given issue. Bluesky feeds shape encounter opportunities, which can sustain many components even when degree alone would predict a larger connected core.

Reddit sits between Moltbook and Bluesky as it consolidates strongly, but remains more fragmented than degree alone predicts. This aligns with a thread-structured participation regime where conversation is organised into many parallel subthreads that do not necessarily cross-link, an effect documented in prior Reddit thread studies.

4.3. RQ2 (reciprocity as a conversational signature): reciprocity beyond degree expectations

Observed reciprocity (Table 3) is substantially higher than degree-preserving expectations for all three platforms (all $p=0.006$ in the null ensemble comparison). In practice, this means directed ties are not simply one-way broadcasts; they contain a strong back-and-forth component even after controlling for in/out-degree sequences.

Table 3. Reciprocity compared across networks

Platform	Reciprocity observed	Reciprocity Null mean	Reciprocity Null low	Reciprocity Null high	p_reciprocity
Moltbook	0.036	0.001	0	0.003	0.006
Bluesky	0.350	0.011	0	0.024	0.006
Reddit	0.264	0.007	0	0.015	0.006

Reciprocity indicates interaction norms that support mutual exchange rather than unilateral attention capture. In all of these networks, reciprocity is systematically above degree expectations across platforms which suggests that these platforms differ in participation regimes, not only in scale or activity.

5. Discussion and Conclusion

Moltbook reaches near-complete connectivity among active agents within a single day, and within minutes once participation volume is non-trivial (≥ 100 active agents). Research on networks using a temporal perspective treats the time to consolidation of the network as a property of the *timing and ordering of interaction events*, not only the final graph. When edges arrive in bursts, connectivity can change discontinuously, and early edge placement can dominate later structure (Holme & Saramäki, 2012). In research using an activity perspective, heterogeneous activation rates can generate rapid hub formation and fast connectivity growth even without any change in platform architecture. In this paper, we tested the time to consolidation as against time- and degree-informed null expectations rather than as a deterministic platform property (Perra et al., 2012). Related work examining conventional OSNs indicate that automated “bot” accounts are disproportionately active in the early diffusion stage of the network and can use posting behaviour to achieve accelerated component stitching (Shao et al., 2018). Recent work on AI agent-only online communities in a simulated environment reports that LLM-based autonomous agents can form structured, centralised interaction networks (Chen et al., 2025). This research extends this finding by empirically validating it on an OSN that is entirely populated by AI agents.

We also find that participation regimes differ across Moltbook, Bluesky and Reddit. Platform effects can be expressed through participation regimes or how affordances are taken up in practice, not merely through declared technological features. Ethnographic work on affordances emphasises that observed outcomes are produced by how users and agents realise affordances in context (Costa, 2018). In this paper, fragmentation/consolidation trajectories may be behavioural signatures of platform users rather than determined by design. Our results indicate that Bluesky discussion remains fragmented and becomes more fragmented as more accounts participate, while Reddit discussion consolidates over days. This finding is supported by empirical analyses of Bluesky’s interaction networks which report pronounced structural clustering and indicates that increased participation does not necessarily result into convergence on a single connected arena (Quelle & Bovet, 2025; Failla & Rossetti, 2024). For Reddit, research evidence shows that thread growth dynamics and interaction conventions can yield distinctive structural evolution over time; this is consistent with our finding of consolidation across days while retaining multiple components (Goglia & Vega, 2024).

These differences are not reducible to degree alone on Bluesky/Reddit. Our inferential analysis using degree-preserving nulls show Bluesky’s fragmentation (low LCC share; high component count) significantly exceeds degree expectations. This indicates that additional structure is attributable to the participation regime rather than the degree sequence (Fosdick et al., 2018). Bluesky’s fragmentation is also consistent with related research

that network topology and curation relate to clustering and separation (Quelle & Bovet, 2025). Reciprocity above degree expectations further confirms that participation regimes vary across platforms (Chowdhary et al., 2023).

As social media increasingly will adopt AI models, including platforms populated entirely by AI agents like Moltbook, understanding these structures is crucial for predicting how influence will be achieved with implications for Marketing, Communication and Politics. When a platform's interaction graph consolidates quickly (a large connected component captures most active accounts in a short period), influence becomes easier to concentrate and harder to interrupt because fewer distribution accounts need to be controlled for a message to reach the active audience. In an AI-agent setting like Moltbook, the practical effect is a far shorter monitoring window than traditional social media as by the time a risky narrative is detectable, it may already have spread across the dominant component via high-frequency agent posting and reciprocal engagement loops (Pillay, 2026; Taylor, 2026). For marketing and strategic communications, campaign management may have to shift from week scale optimisation to hour scale anomaly detection and rapid counter-messaging. In these areas, companies are vulnerable because attention can be "manufactured" by automated accounts and then stabilised through rapid repeated amplification (Ferrara et al., 2016; Lazer et al., 2018). For public communication, faster consolidation increases the risk that unfounded rumors and misinformation outpaces institutional response cycles which require approvals and press-office, creating avoidable public harms if correction arrives after the narrative has saturated the main component (Lazer et al., 2018). For politics, the same structure reduces the cost of coordinated persuasion, existing work on computational propaganda via bots already frames mass persuasion as a production and amplification system, and rapid consolidation makes that system more efficient because a small set of orchestrated agents can reach a large share of active participants with fewer intermediaries (Mustafa et al., 2025; Woolley & Howard, 2016). The implication for governance is that platforms, regulators, and campaign teams need network level early warning (including consolidation speed) alongside content signals, because structural acceleration of narratives using AI changes time to respond from a preference of internal communication teams into a measurable risk parameter for organisations, institutions and governments.

As a limitation, this study uses a short, early-life observation window of Moltbook and relies on publicly observable interaction traces (reply ties) that may omit other salient relations (e.g., follows, reposts/quotes) and can be affected by data-access constraints, deletions, and scraper/API sampling artefacts, all of which can bias network estimates and their interpretation. In addition, AI agent participation cannot be assumed to be fully autonomous as online automation can range from scripted bots to human-in-the-loop systems, and attribution/detection remains imperfect, so time-to-consolidation should be read as a property of observed participation rather than solely agent autonomy. Future work should estimate consolidation at finer temporal resolutions (e.g., hourly) over longer horizons and link consolidation to downstream outcomes such as workload by human and non-human OSN moderators by integrating network dynamics with content measures.

Declaration on the use of generative AI

We acknowledge the use of generative AI in the preparation and formatting of this article. Consensus.app and Google Scholar lab search were used for literature search and providing an overview of existing literature. ChatGPT 5.2 Thinking was used for developing code in Python. Claude Opus 4.5 and Grammarly was used for proofreading and editing assistance.

References

- Antonakaki, D., Fragopoulou, P., & Ioannidis, S. (2021). A survey of Twitter research: Data model, graph structure, sentiment analysis and attacks. *Expert Systems with Applications*, 164, 114006. <https://doi.org/10.1016/j.eswa.2020.114006>
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. <https://doi.org/10.1126/science.286.5439.509>

- Boyd, D. M., & Ellison, N. B. (2008). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13(1), 210–230. <https://doi.org/10.1111/j.1083-6101.2007.00393.x>
- Bright, J. (2018). Explaining the emergence of political fragmentation on social media: The role of ideology and extremism. *Journal of Computer-Mediated Communication*, 23(1), 17–33. <https://doi.org/10.1093/jcmc/zmx002>
- Broido, A. D., & Clauset, A. (2019). Scale-free networks are rare. *Nature Communications*, 10, 1017. <https://doi.org/10.1038/s41467-019-08746-5>
- Charmanas, K., Mittas, N., & Angelis, L. (2024). Content and interaction-based mapping of Reddit posts related to information security. *Journal of Computational Social Science*, 7(2), 1187–1222. <https://doi.org/10.1007/s42001-024-00269-4>
- Chen, H., Wang, Z., & Ren, M. (2025). Unveiling the collective behaviors of large language model-based autonomous agents in an online community: A social network analysis perspective. *Data and Information Management*, 10(1), 100107. <https://doi.org/10.1016/j.dim.2025.100107>
- Chowdhary, S., Andres, E., Manna, A., Blagojević, L., Di Gaetano, L., & Iñiguez, G. (2023). Temporal patterns of reciprocity in communication networks. *EPJ Data Science*, 12, 38. <https://doi.org/10.1140/epjds/s13688-023-00382-w>
- Costa, E. (2018). Affordances-in-practice: An ethnographic critique of social media logic and context collapse. *New Media & Society*, 20(10), 3641–3656. <https://doi.org/10.1177/1461444818756290>
- De Angelis, I., & Alexander, R. (2025). What are Canadian Members of Parliament doing on Bluesky? *Canadian Journal of Political Science*, 1–22. <https://doi.org/10.1017/S0008423925100619>
- Dutta, U., Fosdick, B. K., & Clauset, A. (2025). Sampling random graphs with specified degree sequences. *Journal of Computational and Graphical Statistics*, 34(3), 934–947. <https://doi.org/10.1080/10618600.2024.2418817>
- Failla, A., & Rossetti, G. (2024). “I’m in the Bluesky tonight”: Insights from a year worth of social data. *PLOS ONE*, 19(11), e0310330. <https://doi.org/10.1371/journal.pone.0310330>
- Fosdick, B. K., Larremore, D. B., Nishimura, J., & Ugander, J. (2018). Configuring random graph models with fixed degree sequences. *SIAM Review*, 60(2), 315–355. <https://doi.org/10.1137/16M1087175>
- Gauvin, L., Génois, M., Karsai, M., Kivela, M., Takaguchi, T., Valdano, E., & Vestergaard, C. L. (2022). Randomized reference models for temporal networks. *SIAM Review*, 64(4), 763–830. <https://doi.org/10.1137/19M1242252>
- Gel, Y. R., Lyubchich, V., & Ramirez Ramirez, L. L. (2017). Bootstrap quantification of estimation uncertainties in network degree distributions. *Scientific Reports*, 7, 5805. <https://doi.org/10.1038/s41598-017-05885-x>
- Goglia, D., & Vega, D. (2024). Structure and dynamics of growing networks of Reddit threads. *Applied Network Science*, 9, 48. <https://doi.org/10.1007/s41109-024-00654-y>
- Himmelboim, I., Smith, M., Rainie, L., Shneiderman, B., & Espina, C. (2017). Classifying Twitter topic-networks using social network analysis. *Social Media + Society*, 3(1). <https://doi.org/10.1177/2056305117691545>
- Holme, P. (2015). Modern temporal network theory: A colloquium. *The European Physical Journal B*, 88, 234. <https://doi.org/10.1140/epjbe/e2015-60657-4>
- Holme, P., & Saramäki, J. (2012). Temporal networks. *Physics Reports*, 519(3), 97–125. <https://doi.org/10.1016/j.physrep.2012.03.001>
- Iacopini, I., Karsai, M., & Barrat, A. (2024). The temporal dynamics of group interactions in higher-order social networks. *Nature Communications*, 15, 7391. <https://doi.org/10.1038/s41467-024-50918-5>
- Karsai, M., Kivela, M., Pan, R. K., Kaski, K., Kertész, J., Barabási, A.-L., & Saramäki, J. (2011). Small but slow world: How network topology and burstiness slow down spreading. *Physical Review E*, 83(2), 025102. <https://doi.org/10.1103/PhysRevE.83.025102>
- Kim, H., Jo, H.-H., & Jeong, H. (2021). Impact of environmental changes on the dynamics of temporal networks. *PLOS ONE*, 16(4), e0250612. <https://doi.org/10.1371/journal.pone.0250612>
- Krings, G., Karsai, M., Bernhardsson, S., Blondel, V. D., & Saramäki, J. (2012). Effects of time window size and placement on the structure of an aggregated communication network. *EPJ Data Science*, 1, 4. <https://doi.org/10.1140/epjds4>
- Leifeld, P., Cranmer, S. J., & Desmarais, B. A. (2018). Temporal exponential random graph models with btergm: Estimation and bootstrap confidence intervals. *Journal of Statistical Software*, 83(6), 1–36. <https://doi.org/10.18637/jss.v083.i06>
- des Mesnards, N. G., Hunter, D. S., el Hjouji, Z., & Zaman, T. (2022). Detecting bots and assessing their impact in social networks. *Operations research*, 70(1), 1–22. <https://doi.org/10.1287/opre.2021.2118>
- Pillay, T. (2026, February 3). Moltbook is a social network for AI bots. Here’s how it works. *TIME*. <https://time.com/7364662/moltbook-ai-reddit-agents/>
- Perra, N., Gonçalves, B., Pastor-Satorras, R., & Vespignani, A. (2012). Activity driven modeling of time varying networks. *Scientific Reports*, 2, 469. <https://doi.org/10.1038/srep00469>
- Quelle, D., & Bovet, A. (2025). Bluesky: Network topology, polarization, and algorithmic curation. *PLOS ONE*, 20(2), e0318034. <https://doi.org/10.1371/journal.pone.0318034>
- Ribeiro, B., Perra, N., & Baronchelli, A. (2013). Quantifying the effect of temporal resolution on time-varying networks. *Scientific Reports*, 3, 3006. <https://doi.org/10.1038/srep03006>
- Rocha, L. E. C., Masuda, N., & Holme, P. (2017). Sampling of temporal networks: Methods and biases. *Physical Review E*, 96(5), 052302. <https://doi.org/10.1103/PhysRevE.96.052302>
- Safdari, H., Contisciani, M., & De Bacco, C. (2022). Reciprocity, community detection, and link prediction in dynamic networks. *Journal of Physics: Complexity*, 3(1), 015010. <https://doi.org/10.1088/2632-072X/ac52e6>
- Satter, R. (2026, February 2). ‘Moltbook’ social media site for AI agents had big security hole, cyber firm Wiz says. *Reuters*. <https://www.reuters.com/legal/litigation/moltbook-social-media-site-ai-agents-had-big-security-hole-cyber-firm-wiz-says-2026-02-02/>

- Seckin, O. C., Silva, F. N., Truong, B. T., Kim, S., Huang, F., Liu, N., Flammini, A., & Menczer, F. (2025). The rise of Bluesky. *arXiv preprint arXiv: 2504.12902v1* <https://doi.org/10.48550/arXiv.2504.12902>
- Shao, C., Ciampaglia, G. L., Varol, O., Yang, K.-C., Flammini, A., & Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature Communications*, 9, 4787. <https://doi.org/10.1038/s41467-018-06930-7>
- Sheng, A., Su, Q., Li, A., Wang, L., & Plotkin, J. B. (2023). Constructing temporal networks with bursty activity patterns. *Nature Communications*, 14(1), 7311. <https://doi.org/10.1038/s41467-023-42868-1>
- Soliman, A., Hafer, J., & Lemmerich, F. (2019). A characterization of political communities on Reddit. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media* (pp. 259–263). Association for Computing Machinery. <https://doi.org/10.1145/3342220.3343662>
- Squartini, T., Picciolo, F., Ruzzenenti, F., & Garlaschelli, D. (2013). Reciprocity of weighted networks. *Scientific Reports*, 3, 2729. <https://doi.org/10.1038/srep02729>
- Taylor, J. (2026, February 2). What is Moltbook? The strange new social media site for AI bots. *The Guardian*. <https://www.theguardian.com/technology/2026/feb/02/moltbook-ai-agents-social-media-site-bots-artificial-intelligence>
- Tkacz, N., & Gehl, R. W. (2025). After Twitter: Fragmentation, platform politics and protective sociality. *Social Media + Society*, 11(3). <https://doi.org/10.1177/20563051251366907>
- Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. *Proceedings of the International AAAI Conference on Web and Social Media*, 11(1), 280–289. <https://doi.org/10.1609/icwsm.v11i1.14871>
- Weber, M. S., Fulk, J., & Monge, P. (2016). The emergence and evolution of social networking sites as an organizational form. *Management Communication Quarterly*, 30(3), 305–332. <https://doi.org/10.1177/0893318916629547>
- Weng, L., Menczer, F., & Ahn, Y.-Y. (2013). Virality prediction and community structure in social networks. *Scientific Reports*, 3, 2522. <https://doi.org/10.1038/srep02522>
- Zhang, Z.-K., Liu, C., Zhan, X.-X., Lu, X., Zhang, C.-X., & Zhang, Y.-C. (2016). Dynamics of information diffusion and its applications on complex networks. *Physics Reports*, 651, 1–34. <https://doi.org/10.1016/j.physrep.2016.07.002>
- Zhu, Y.-X., Zhang, X.-G., Sun, G.-Q., Tang, M., Zhou, T., & Zhang, Z.-K. (2014). Influence of reciprocal links in social networks. *PLOS ONE*, 9(7), e103007. <https://doi.org/10.1371/journal.pone.0103007>