Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Generative AI and interrupted peripheral participation: evidence of structural closure from Stack Exchange

Junlei Du
Beijing Normal University
Shuang Li
lilybnu@bnu.edu.cn

Beijing Normal University
Yinpei Yu
Beijing Normal University
Xinyi He
Beijing Normal University
Jingxi Chen
Beijing Normal University

Article

Keywords:

Posted Date: August 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10221893/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Generative AI and interrupted peripheral participation: evidence of structural closure from Stack Exchange

Anonymous

Abstract

Generative AI can satisfy routine information needs without creating the public interactions through which newcomers enter and become embedded in online knowledge communities. This study examines whether that substitution interrupts legitimate peripheral participation and changes community social structure differently across knowledge domains. Seven Stack Exchange communities from 2016 to 2025 are compared as algorithmic, disciplinary, or practice-based. ChatGPT's release is treated as a quasi-experimental shock in a difference-in-differences design, supplemented by event studies, wild-cluster bootstrap inference, synthetic control, leave-one-community-out estimates, and placebo-year tests. Algorithmic communities show increased modularity and longer average paths. Newcomer one-year retention falls by nearly half, with the decline concentrated among low-activity peripheral members while retention among high-activity members remains stable. Code density, rather than formula density, follows the gradient of structural impact, and practice-based communities remain structurally stable. These findings extend legitimate peripheral participation by showing how an external information substitute can interrupt the entry-stage interactions through which communities renew themselves. Procedural codifiability bounds this interruption, while the resulting loss of peripheral ties produces structural closure: bridging weakens as bonding among incumbents remains comparatively intact. The evidence identifies community-level structural effects; individual decisions to use an LLM are not directly observed.

1 Introduction

Generative AI (GenAI) is unavoidably reshaping online knowledge communities. Its world-model competence—knowledge accumulated across web-scale text—lets it absorb the routine, codifiable questions that knowledge platforms once channelled into human Q&A, and the volume effects are by now well documented. Burtch, Lee and Chen (2024) found that question volume on

Stack Overflow fell sharply after ChatGPT's public release in November 2022, with the steepest declines among newer, less embedded contributors. del Rio-Chanona, Laurentsyeva and Wachs (2024) measured a 25% drop in posting within six months of ChatGPT's release, with adopters disengaging from the platform. Shan and Qiu (2025) found that the answers that remained grew shorter and less detailed. These studies agree on the empirical pattern: *aggregate volume* is contracting. What they leave open is the shape of that contraction. Is it uniform across members and across domains, or selective along an axis that prior work has not characterized? The distinction matters because a community is not a pile of independent contributions but a network of ties that sustains distributed knowledge production (Wasko and Faraj, 2005; Faraj, Jarvenpaa and Majchrzak, 2011)—and a literature focused on aggregate volume cannot tell us whether the social-capital structure has been left intact, hollowed out, or reorganised along a community's codifiability profile.

Stack Exchange, as a reputation-governed platform with both an embedded expert core and a transient, network-affiliated periphery, makes the case for unpacking that shape concrete. It exhibits the cohesive features Wenger and colleagues call *communities of practice* (CoP) (Wenger, 1998; Wenger, McDermott and Snyder, 2002), and at the same time the self-organized features Wasko and Faraj (2005) calls *networks of practice* (NoP), with many loosely affiliated participants exchanging knowledge through indirect, pay-it-forward reciprocity. In this hybrid, routine questions do more than add content. They provide low-stakes occasions for *legitimate peripheral participation* (LPP), through which newcomers begin situated engagement and may develop into fuller members (Lave and Wenger, 1991). They also create weak, cross-topic ties between a practice-community core and a networked periphery (Putnam, 2000; Granovetter, 1973; Zhang and Watts, 2008; Faraj et al., 2011).

LLMs alter this entry arrangement because they can satisfy an information need before a community interaction is created. If a newcomer takes a routine question to an LLM, the answer may be obtained, but the public question, reply, and social tie through which peripheral participation would have begun do not occur. This suggests an interruption of LPP at its entry stage rather than a uniform reduction in participation. The interruption should be strongest where knowledge is *procedurally codifiable*—expressible as code that can be generated, executed, and checked—and weaker for multi-step formal derivations or situated practical knowledge (Polanyi, 1966; Lundvall and Johnson, 1994). At the community level, the corresponding structural consequence should be asymmetric: the cross-topic ties associated with peripheral participation

should erode faster than bonding among embedded incumbents. Prior studies show post-ChatGPT declines concentrated among newer contributors and changes in adjacent forms of digital knowledge production (Burtch et al., 2024; Yeverechyahu, Mayya and Oestreicher-Singer, 2024), but they do not connect entry-stage substitution to the renewal and social-capital structure of online knowledge communities.

We address this gap with a longitudinal comparison of seven Stack Exchange communities (2016–2025), grouped by knowledge codifiability into algorithmic (AI, Data Science), disciplinary (Statistics, Physics, Chemistry), and practice-based (Cooking, Travel) tiers. All seven share the same reputation system, moderation, and interface, which holds platform design constant and isolates the codifiability axis. We treat ChatGPT's public release as a quasi-experimental shock and analyse the effect across three convergent levels of evidence—network structure, user retention, and content composition—then pose three research questions:

RQ1: How does the bonding and bridging social capital network structure of online knowledge communities with different levels of knowledge codifiability evolve across 2016–2025?

RQ2: Did the emergence of LLMs significantly restructure these communities' social capital, and does the effect differ across communities by knowledge codifiability?

RQ3: How do structural changes relate to functional outcomes (response latency, knowledge content), and what user-level and content-level mechanisms accompany them?

The study makes one central theoretical contribution: it extends LPP by specifying how a platform-external information technology can interrupt the apprenticeship trajectory before its first community interaction. It then identifies procedural codifiability as a boundary condition on that proposed interruption and asymmetric social-capital reorganisation as its community-level consequence. Evidence from network structure, user retention, and content composition is consistent with this account: algorithmic communities lose peripheral renewal and cross-topic connectivity while their incumbent core remains comparatively stable. Because we do not observe users' choice to consult an LLM directly, we treat interruption of LPP as the mechanism most consistent with the convergent evidence, not as a directly observed individual transition.

2 Literature Review

2.1 Online Knowledge Communities and Their Social Capital Network Structure

Online Q&A platforms such as Stack Exchange are a distinctive kind of community: participation is voluntary, knowledge exchange runs on self-organized question-answering, and reputation systems coordinate contributions without formal authority (Mamykina, Manoim, Mittal, Hripcsak and Hartmann, 2011; Wasko and Faraj, 2005). Knowledge flow depends on the underlying *network structure*, not aggregate volume alone. These communities typically take a core-periphery shape—a small expert elite produces most highly rated content, surrounded by a much larger periphery of transient askers (Anderson, Huttenlocher, Kleinberg and Leskovec, 2012; Movshovitz-Attias, Movshovitz-Attias, Steenkiste and Faloutsos, 2013)—the expected outcome of cumulative advantage, in which preferential attachment drives a few nodes to accumulate most ties (Barabási and Albert, 1999) and the Matthew effect compounds early success into lasting advantage (Merton, 1968). Vranić, Tomašević, Alorić and Mitrović Dankulov (2023) showed that an early, stable core predicts long-term survival, yet the same concentration turns fragile once core members leave. The periphery is easy to dismiss as occasional askers, but it generates the cross-cutting interactions that connect topic clusters (Preece and Shneiderman, 2009); without those boundary-crossing ties, a community is only a set of isolated subgroups.

Social capital theory links these structural conditions to community function and has long been applied to online settings (Blanchard and Horan, 1998). Lin (2001) defined social capital as resources embedded in networks and accessed through one's position; Putnam (2000) distinguished *bonding capital*—dense, homogeneous within-subgroup ties that reinforce solidarity—from *bridging capital*—weak, cross-cutting ties that carry novel information across clusters—and, building on Granovetter (1973), Burt (2004) showed that those who bridge disconnected clusters gain an outsized information advantage. In Q&A communities, bonding sustains specialized exchange within clusters and bridging sustains knowledge flow across them, deciding whether the community works as one integrated network or a set of insular enclaves. The empirical record agrees: Wasko and Faraj (2005) found network centrality predicts contribution volume and helpfulness; Best and Krueger (2006) found that interaction with ties first formed online relates positively to social capital; Ellison, Steinfield and Lampe (2007) found that online relationship formation predicts bridging-capital gains, especially for less integrated users; and Safadi, Johnson and Faraj

(2021) found that epistemically marginal participants contribute unusually valuable content when socially embedded.

This distinction makes network structure consequential for interpreting post-GenAI decline. A uniform contraction can leave the balance of bonding and bridging largely intact; a selective weakening of the periphery can reorganise that balance even when both core and peripheral activity decline in absolute terms.

2.2 Legitimate Peripheral Participation and Entry-Stage Interruption

The CoP framework (Wenger, 1998; Wenger et al., 2002) treats communities as small, cohesive groups whose members share an ongoing practice and build a common identity, language, and norms; knowledge sharing runs on *direct reciprocity* among members who know one another. Lave and Wenger (1991) added a generative mechanism, *legitimate peripheral participation* (LPP): newcomers begin with low-stakes routine activities at the edge of the practice, gain access to more central tasks, and through sustained situated engagement become full members. LPP is how a community renews itself—the periphery is not noise around the core but its apprentice pipeline. Wasko and Faraj (2005) extended this with *networks of practice* (NoP): larger, loosely coupled groups in which most participants are strangers and contributions follow *indirect*, pay-it-forward reciprocity. Zhang and Watts (2008) and Faraj et al. (2011) argued that many large knowledge platforms sit *between* these ideal types—a CoP-like core of repeat contributors and an NoP-like periphery of transient askers whose questions create the cross-topic ties the platform's broader knowledge function relies on. Recent comparative work confirms this two-layer structure (Li, Du and Yu, 2024; Bighash, Oh, Fulk and Monge, 2018; Vlegels and Huisman, 2021).

The LPP trajectory has a specific shape that matters for the GenAI era. Its earliest stage often begins with routine, low-stakes questions ("how do I import this package," "why does this error appear"). An external LLM can satisfy such an information need without generating the public exchange through which situated learning, recognition, and relationship formation would otherwise begin. Theoretically, this is not simply participant exit: it also includes *non-entry*, because some prospective interactions never appear in platform records. Among people who do enter, the same substitution opportunity may reduce repeated peripheral participation and therefore retention. Zhou, Cho and Terveen (2025) documents an analogous disruption on Wikipedia, where LLM-drafted contributions displace entry-level edits that previously helped socialize new editors. Linking external information substitution to these two margins—reduced

entry and weaker continuation—extends LPP from an account of how participation accumulates to one that specifies how the trajectory can be interrupted.

The hybrid CoP/NoP framing translates this interruption into observable implications. The core runs on reputation, recurring exchange, and intrinsic motivation that a one-shot LLM interaction does not reproduce (Wenger et al., 2002; Faraj et al., 2011). The periphery more often brings routine, one-shot questions that an LLM can answer in a self-contained interaction (Kabir, Udo-Imeh, Kou and Zhang, 2024; Hou, Mettille, Man, Li, Zastudil and MacNeil, 2024). If this substitution selectively weakens entry and continuation at the periphery, missing or discontinued questions also remove weak, cross-topic ties, while recurring ties among incumbents remain comparatively intact. This yields the general structural hypothesis and its two observable components:

H1: GenAI disruption produces an *asymmetric* restructuring of social capital: bonding capital intensifies while bridging capital erodes.

H1a: Bonding capital indicators (e.g., modularity, within-cluster density) will rise more in algorithmic communities than in disciplinary or practice-based communities.

H1b: Bridging capital indicators (e.g., E-I index, average path length on the giant connected component) will deteriorate more in algorithmic communities than in disciplinary or practice-based communities.

2.3 GenAI Disruption and Selective Substitution

Empirical research on how AI affects online knowledge communities has grown fast, but most stops at *volume*. Kabir et al. (2024) found that 52% of ChatGPT's answers to sampled Stack Overflow questions were incorrect, yet users preferred them 35% of the time and overlooked the errors 39% of the time. del Rio-Chanona et al. (2024) measured a 25% drop in posting within six months of ChatGPT's release. Burtch et al. (2024) found the decline concentrated among newer users and absent in matched Reddit communities, suggesting that social fabric buffers against substitution. Shan and Qiu (2025) found surviving answers grew shorter and less detailed; Hou et al. (2024) found computing students increasingly take routine questions to LLMs while saving the community for complex problems; and Zhou et al. (2025) found LLMs opened an expertise-based participation divide on Wikipedia, helping experienced editors while alienating newcomers. The disruption reaches beyond Q&A: Yeverechyahu et al. (2024) found

GitHub Copilot changed how developers contribute to open-source projects, evidence that LLM assistants reshape collaborative knowledge production in adjacent settings.

The knowledge management literature specifies a boundary condition on this interruption: the substitutability of the entry task. Codifiable knowledge—facts, algorithms, and derivations that can be stated in explicit, rule-based form—is generally more reproducible by LLMs, whereas tacit, practice-embedded knowledge resists formalization (Polanyi, 1966; Lundvall and Johnson, 1994; Nonaka and Takeuchi, 1995). Ide and Talamàs (2024) models this contrast directly. Yet explicit knowledge takes at least two empirically distinct forms. *Procedural codifiability* is knowledge expressible as code and checkable through execution; *formal codifiability* is knowledge expressible as mathematical derivation but not necessarily easy for a non-expert to verify step by step. The proposed LPP interruption should therefore follow the execution-verifiable substitution surface more closely than codifiability in the abstract. This boundary condition gives us H2:

H2a: A difference-in-differences estimator using the public release of ChatGPT as a quasi-experimental shock will show *significant* effects on the H1a/H1b indicators for algorithmic communities relative to a low-exposure practice-based comparison group.

H2b: Treatment effects will be *ordered along the procedural codifiability subdimension*: largest for algorithmic, intermediate for disciplinary, smallest for practice-based communities. The formal codifiability subdimension (on which disciplinary leads) should not exhibit the same gradient.

2.4 From Structural Change to Functional and Content Consequences

If GenAI reshapes social capital asymmetrically, the next question is whether that hurts knowledge-sharing performance. Prior network research suggests it should: high modularity slows information diffusion by trapping flows inside subgroups (Centola, 2010), and longer average paths mean queries take more steps to reach someone who can answer (Kossinets and Watts, 2009). In Q&A communities, response latency—the time from a question to its first answer—is a key functional outcome that depends on how reachable expertise is (Mamykina et al., 2011), so we expect the structural changes in H1–H2 to show up as measurable functional decline. Knowledge content can fragment too: as cross-topic askers leave, tags that used to co-occur on bridging questions appear alone and the tag co-occurrence network thins (Xing and Du, 2023). This dual-level fragmentation—social and semantic ties dissolving together—sets GenAI-era decline apart from

earlier episodes of volume loss without content restructuring (Iriberri and Leroy, 2009; Butler, 2001). We capture this in H3:

H3: Structural changes in bonding and bridging capital are associated with degraded knowledge-sharing function (rising response latency) and with parallel fragmentation of knowledge content (declining tag co-occurrence density).

The predicted topology also requires a conceptual boundary. Dense subgroups connected by fewer cross-group ties resemble one network signature discussed in echo-chamber research (Cinelli, De Francisci Morales, Galeazzi, Quattrociocchi and Starnini, 2021; Quattrociocchi, 2017; Valensise, Cinelli and Quattrociocchi, 2023), but this study does not measure ideological homogeneity, selective exposure, or recommendation. We therefore use *structural closure* to denote the community-level consequence of weakened peripheral bridging. Unlike curation- and motivated-selection accounts, the proposed pathway operates through interactions that fail to form or continue, not through the content choices of users who remain active.

3 Research Design and Methodology

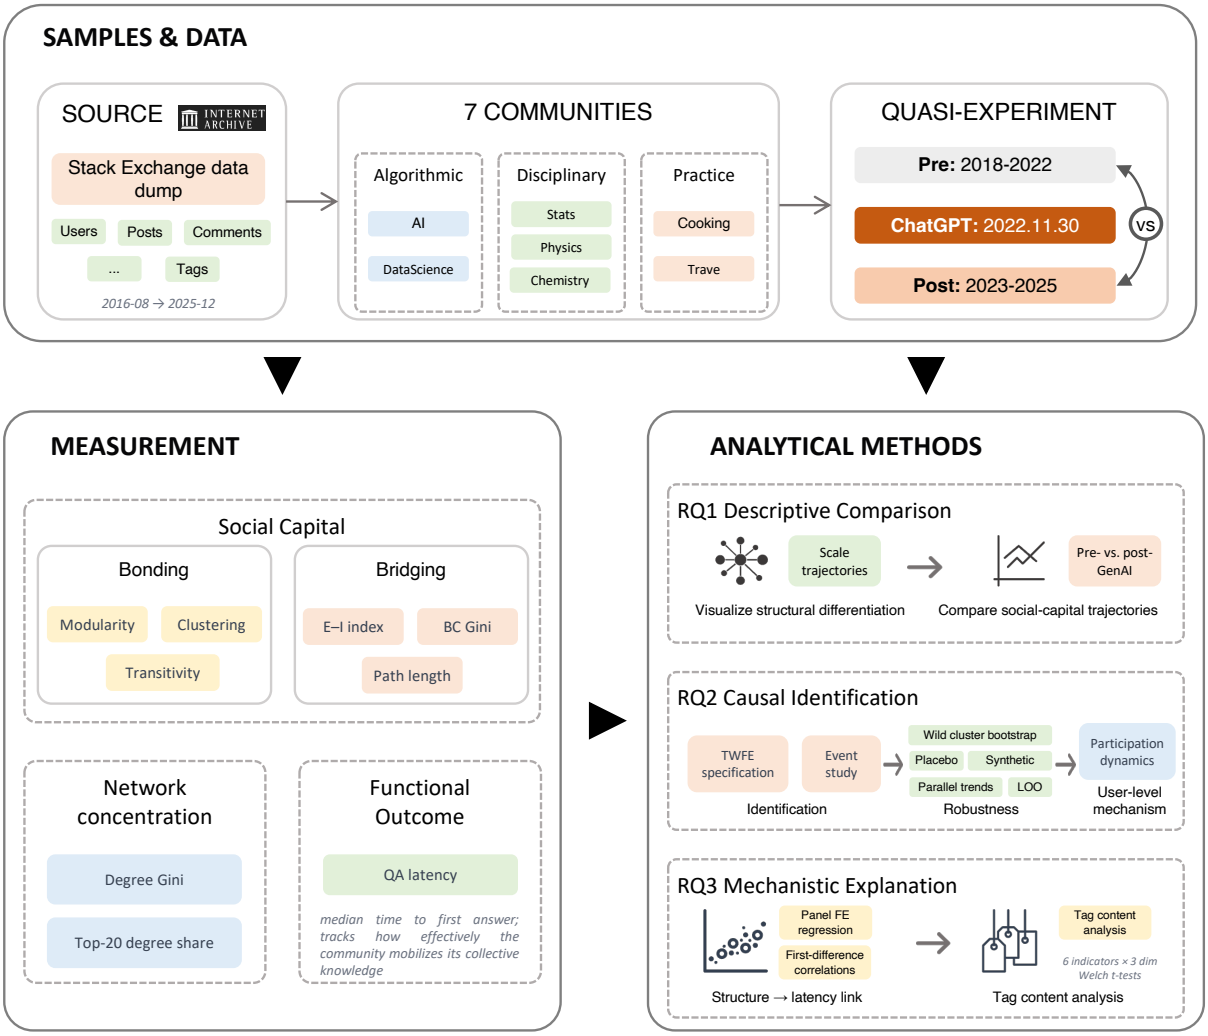


Figure 1: Research design and analysis flow. The diagram links the Stack Exchange data source and knowledge-domain grouping to the network, user-retention, and content analyses used to address the three research questions.

Figure 1 gives an overview of the research design and analysis flow—from the data dump and codifiability-based community grouping, through the measurement of structural features, to the three analytical tracks that address RQ1–RQ3.

3.1 Samples and Data

We use the complete Stack Exchange data dump through December 31, 2025 (Stack Exchange Inc., 2025),[1] which contains every question, answer, comment, and user-activity record since each community's inception. Our analysis window is the ten complete calendar years 2016–2025; we

[1]The complete data dump is publicly archived at `https://archive.org/details/stackexchange_20251231` under the CC BY-SA 4.0 license.

omit pre-2016 data because several of the smaller communities (notably AI.SE, founded in 2016) had not yet exited their early-bootstrap phase, so their year-1 statistics are not comparable to later years.

Our multi-community comparative design exploits natural variation in *knowledge codifiability*—how far a domain's knowledge can be stated in explicit, rule-based form (Polanyi, 1966; Nonaka and Takeuchi, 1995)—across seven communities that share the same platform, reputation system, and governance rules, holding institutional design constant. LLMs reproduce codifiable knowledge well but struggle with practice-based knowledge that needs situated judgment (Polanyi, 1966), a gap confirmed empirically across domains (Kabir et al., 2024). We sort the seven communities into three groups along this axis:

- **Algorithmic knowledge (High procedural codifiability):** AI and Data Science Stack Exchange—communities in which executable code is a common form of knowledge expression.
- **Disciplinary knowledge (High formal codifiability):** Statistics, Physics, and Chemistry Stack Exchange—communities combining formal derivation with domain-specific judgment.
- **Practice-based knowledge (Low procedural and formal codifiability):** Cooking and Travel Stack Exchange—communities relying more heavily on situated, experiential knowledge.

Grouping by codifiability sets up a dose-response analysis: if GenAI disruption reshapes a community's social capital, the effect should grow with how procedurally codifiable the community's knowledge is—that is, with the form of codifiability LLMs actually substitute.

To validate this theory-driven grouping against observable data, we compute two deterministic proxies from question content over the full corpus, treating codifiability as two sub-dimensions. The *procedural* sub-dimension, proxied by *code-block density* (the share of body characters in HTML `<code>` or `<pre>` tags), captures knowledge expressible as executable procedure; the *formal* sub-dimension, proxied by *formula density* (LaTeX math segments per 1,000 body characters),[2] captures knowledge expressed as formal derivation. Algorithmic knowledge is

[2]We count occurrences of all four inline and display LaTeX math delimiters: `$...$`, `$$...$$`, `\(...\)`, and `\[...\]`.

canonically procedure (code), disciplinary knowledge formal derivation (math), and practice-based knowledge resists both. The two are not redundant—they distinguish two forms in which codifiable knowledge manifests—and we report them separately.

From the dump we build, for each community and calendar year $t \in \{2016, \ldots, 2025\}$, a weighted undirected interaction graph $G_t = (V_t, E_t, w)$. Nodes are users who authored at least one question, answer, or comment in year t (system accounts excluded). We add an edge (u, v) if u answered a question by v, or if either commented on a post by the other; its weight $w(u, v)$ is the total count of such interactions in that year. This yields 70 community-year graph snapshots, the unit of analysis throughout the paper. Supplementary Table S1 summarises the seven communities.

Empirical validation of the grouping. The codifiability proxies in Supplementary Table S1 validate the three-tier grouping—but only when codifiability is decomposed into its two sub-dimensions. On the procedural sub-dimension, algorithmic communities lead by a wide margin (26.8% code-block density vs. 7.8% disciplinary, 0.4% practice-based). On the formal sub-dimension, the ranking inverts: disciplinary communities carry 4.1 LaTeX math segments per 1,000 characters against 1.1 in algorithmic and 0.01 in practice-based. Algorithmic and disciplinary knowledge are both highly codifiable, but in different forms.

3.2 Measurement

Working from the 70 community-year graphs constructed above, we measure four structural features of each network: the two canonical forms of social capital—*bonding* and *bridging* (Putnam, 2000), following Lin's (2001) structural-access framework—together with a complementary *network-concentration* property and a *functional* outcome.

Bonding capital measures the strength and cohesion of ties within subgroups—how far members interact inside specialized topic clusters rather than across the wider community. The primary indicator is the *modularity* Q of the Louvain community partition (Blondel, Guillaume, Lambiotte and Lefebvre, 2008); higher values mean sharper subgroup boundaries and more insular exchange. Two supporting indicators, the *average clustering coefficient* (Watts and Strogatz, 1998) and *transitivity* (Newman, 2003), capture local triadic closure as further facets of cohesion.

Bridging capital measures how cross-group connections are distributed and how reachable they are—the paths along which knowledge crosses topic boundaries. We use three indicators. The *E-I index* (Krackhardt and Stern, 1988) is $(E - I)/(E + I)$, where E and I count inter- and intra-community edges from the Louvain partition; values near $+1$ mean members interact mostly across boundaries (bridging dominance), values near -1 mean interaction stays within subgroups (bonding dominance). The *average shortest path length* on the giant connected component (GCC) measures reachability; longer paths mean knowledge needs more steps to travel between parts of the community. The *betweenness-centrality Gini* (G_{BC}) measures inequality in brokerage roles; higher values mean the capacity to connect different parts of the community rests with fewer intermediaries (Burt, 2004).

Network concentration captures how evenly interaction is spread across participants—a distributional property of the network that complements, rather than belongs to, the bonding/bridging social-capital measures. The primary indicator is the *degree Gini coefficient*, the overall inequality in interaction volume. As a supporting indicator we report the *Top-20 degree share*: the fraction of total weighted degree held by the 20 highest-degree nodes (Merton, 1968).

Functional outcome is proxied by the *median QA response latency* ($\tau_{0.5}$, hours), the median time from a question being posted to its first answer. It captures how effectively the community mobilizes its collective knowledge to answer incoming questions; a rising $\tau_{0.5}$ signals that the knowledge-sharing infrastructure is deteriorating.

3.3 Analytical Methods and Process

We tailor the analytical method to each empirical question, matching the granularity of the question to the identification strategy that can answer it. Throughout the empirical analysis we define the pre-GenAI period as 2018–2022 and the post-GenAI period as 2023–2025; 2018–2022 is preferred over 2016–2022 to avoid warm-up bias from the August 2016 start of the data dump.

3.3.1 Descriptive Comparison

We profile all seven communities over the full ten years (2016–2025), computing annual values of every structural indicator from Section 3.2 along with $\tau_{0.5}$ on each of the 70 community-year graphs. Group-level trends are within-group means, which surface systematic cross-community patterns; the resulting pre- versus post-GenAI summary is reported in Supplementary Table S2,

with the year-by-year trajectories in Supplementary Fig. S1. All network computations use NetworkX 3.2.1 (Hagberg, Schult and Swart, 2008) in Python 3.11, including the Louvain partition that anchors the modularity and E-I indicators.

For the visual comparison in Figure 2, the full weighted graphs are too dense, so we extract each graph's statistically significant backbone using the multiscale disparity filter of Serrano, Boguñá and Vespignani (2009) ($\alpha = 0.05$), which preserves edges significant at each node's local scale rather than collapsing to high-degree hubs. We render the giant connected component for readability and colour nodes by their k-shell position in the full graph to mark core versus periphery; we use k-shell rather than degree- or betweenness-based identification because it depends on connectivity structure rather than absolute scores and so stays comparable across very different community sizes. We pair the structural panel with two descriptive layers: aggregate scale trajectories (active users and edges) per community-year, which separate volume decline from structural reorganisation, and the two deterministic content proxies—code-block and formula density (Section 3.2)—computed over the full corpus to validate the grouping.

3.3.2 Difference-in-Differences

To distinguish a GenAI-induced structural shift from mere temporal correlation, we exploit ChatGPT's public release on November 30, 2022 as a quasi-experimental shock and adopt a difference-in-differences design (Angrist and Pischke, 2009) with two-way fixed effects (Wooldridge, 2010):

$$Y_{it} = \alpha_i + \gamma_t + \beta_1 \cdot \text{Post}_t \times \text{High}_i + \beta_2 \cdot \text{Post}_t \times \text{Medium}_i + \varepsilon_{it} \tag{1}$$

where α_i and γ_t are community and year fixed effects, $\text{Post}_t = \mathbb{1}(t \geq 2023)$, and High_i / Medium_i flag algorithmic and disciplinary communities respectively; β_1 and β_2 capture effects relative to the low-exposure practice-based comparison group. This specification absorbs both time-invariant community baselines (size, age, moderation culture) and SE-wide time trends, leaving cross-group differential variation in the post window as the identifying signal. We use a 3-tier treatment intensity rather than a binary split because the prediction is a codifiability-graded dose-response, not a single average effect; collapsing Medium and Low into one untreated bucket would mask the gradient. Because LLMs answer some practice-based questions, that group is low-exposure rather than untreated, so β_1 and β_2 identify *differential* rather than absolute effects. We add two checks: an event-study specification estimating year-specific effects relative to 2022 (letting

the timing carry identifying weight), and a formal parallel-trends test (Section 3.3.2).[3]

Robustness suite (five layers) With $N_{\text{cluster}} = 7$, standard cluster-robust inference is unreliable (Cameron, Gelbach and Miller, 2008), and the algorithmic treatment group holds only two communities. We therefore wrap the TWFE estimates in a five-layer robustness suite: a parallel-trends test, a Rademacher wild cluster bootstrap (Cameron et al., 2008), a synthetic control that builds a synthetic AI.SE from the five non-algorithmic communities (Abadie, Diamond and Hainmueller, 2010), leave-one-community-out re-estimation, and placebo treatment years (2020 and 2021). The layers do not change the point estimates from Equation 1; each supplies parallel inference under a different identifying assumption, and we report how well they agree as our robustness verdict. Full procedural detail is in the Supplementary Methods.

3.3.3 Structure-to-Function Analysis

RQ3 asks whether the structural reorganisation identified in RQ2 carries through to functional outcomes (response latency) and content composition. We estimate panel fixed-effects regressions of QA response latency on each structural indicator (Wooldridge, 2010):

$$\tau_{0.5,it} = \alpha_i + \gamma_t + \delta \cdot X_{it} + \varepsilon_{it} \tag{2}$$

We use panel fixed-effects rather than cross-sectional regression because community-level baselines—size, moderation density, response norms—confound any pooled estimate, while the panel-FE specification isolates the within-community over-time link between structure and latency. We add first-difference correlations ($\Delta\tau_{0.5}$ vs. ΔX) as a complementary check that imposes no functional form and detects year-over-year co-movement the FE level regression could miss. We also run a tag-based content analysis with six indicators per community-year, paralleling the multi-indicator structural design across three content dimensions: *breadth* (tag diversity, concentration, entropy), *dynamics* (tag turnover via size-equalized Jaccard, tags per question), and *inter-topic structure* (tag co-occurrence density). Pre- versus post-GenAI differences are tested with group-level Welch's t-tests (Welch, 1947), which do not assume equal variances.

[3]Event study: $Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq 2022} \beta_k \cdot \text{High}_i \times \mathbb{1}(t = k) + \varepsilon_{it}$. Pre-treatment coefficients ($k < 2022$) serve as a visual diagnostic for parallel trends.

4 Results

4.1 RQ1: Comparative Social Capital Evolution Across Community Types

4.1.1 Visualizing the Structural Differentiation across Scales

Figure 2 shows all seven communities in 2019 (pre-GenAI) and 2024 (post-GenAI), drawn on the Serrano–Boguñá–Vespignani disparity-filter backbone (Serrano et al., 2009) ($\alpha = 0.05$). Nodes are coloured by their k-shell position in the full interaction graph: dark blue marks the core (k-shell ≥ 3, i.e. within a subgraph where every member has at least three internal ties, so the cluster survives removal of any single bridging node), light orange the periphery. This lets us read the contraction on two axes—the *core* skeleton versus the *peripheral* fringe—and compare that split across community types.

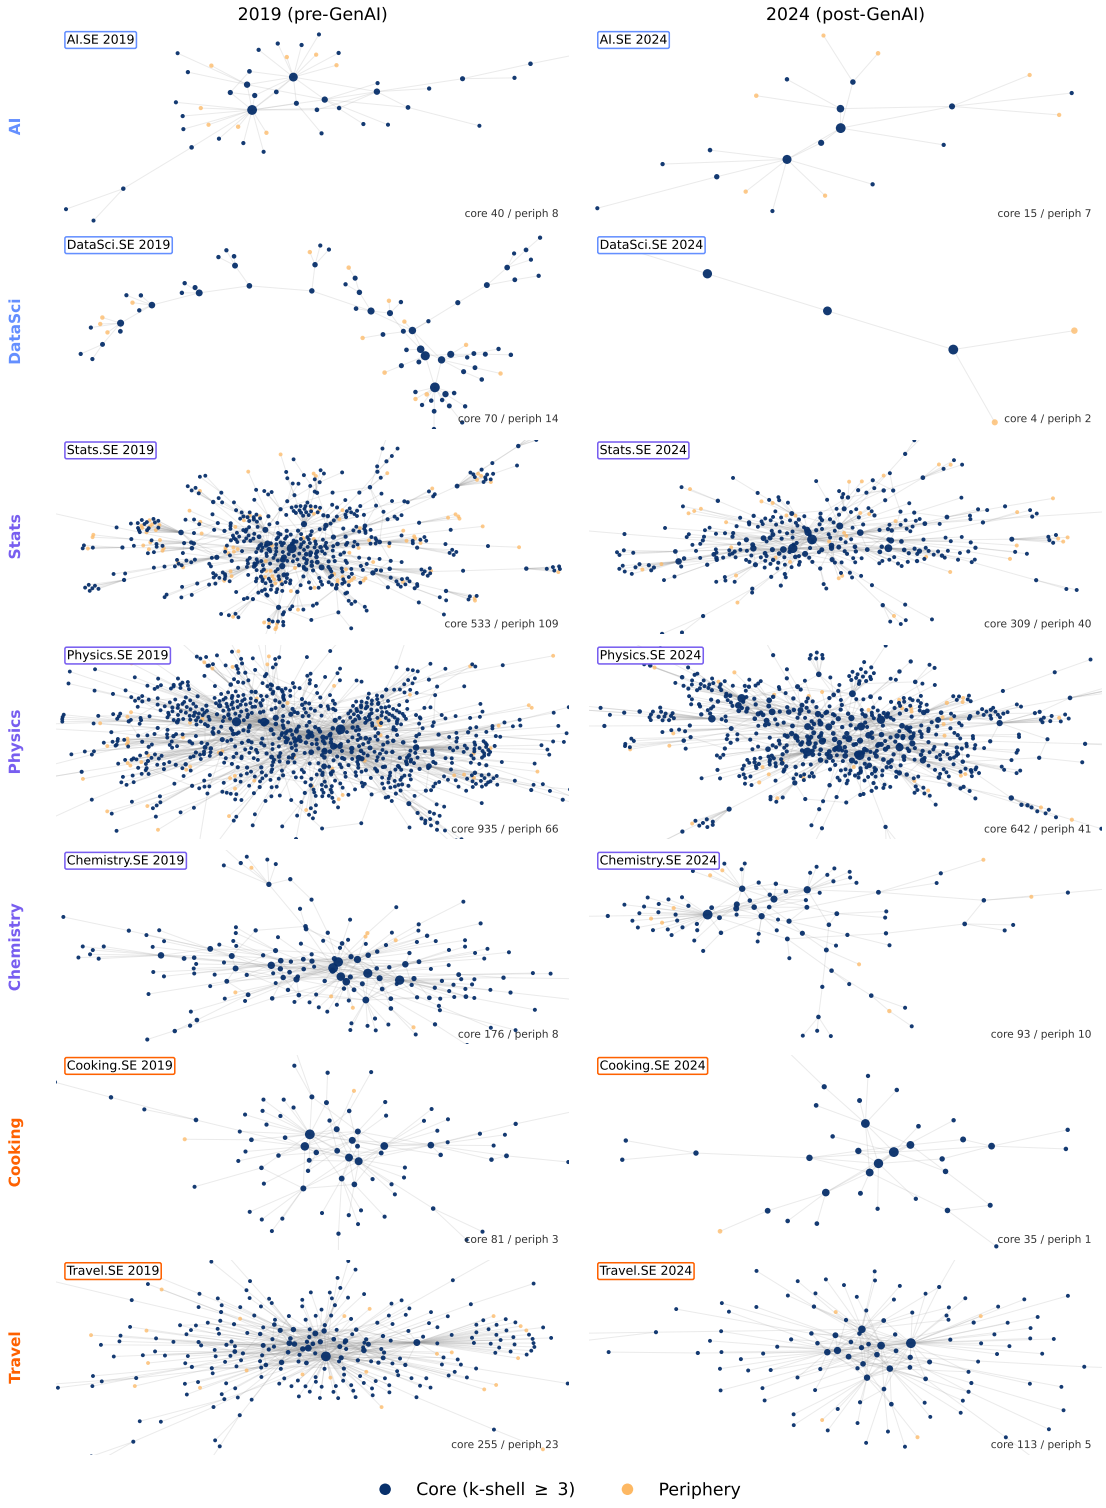


Figure 2: Network backbones before and after the release of ChatGPT. The statistically significant disparity-filter backbone ($\alpha = 0.05$) of each community is shown for 2019 and 2024. Dark blue nodes are core users (k-shell ≥ 3) and light orange nodes are peripheral users; labels report core- and peripheral-node counts.

Three patterns are visible. *Peripheral structure.* The fringe thins everywhere, but with different consequences: in practice-based (Cooking, Travel) and the larger disciplinary communities (Stats, Physics) the periphery shrinks faster than the core, so the surviving backbone becomes *proportionally* more core-dominated—the classic signature of selective peripheral exit that leaves the core skeleton intact. *Core structure.* The absolute core contracts everywhere as participation falls, but the contraction reaches *into* the core only in the algorithmic group, where the core share itself *falls* by roughly 15 percentage points and DataSci.SE's connected core collapses from 70 to 4 nodes. *Cross-community contrast.* The codifiability gradient thus shows up as network shape: practice-based and disciplinary communities shed their fringe but keep a denser core, whereas algorithmic communities lose the fringe *and* erode the core skeleton—two different shapes of contraction, not two intensities of one.

All seven communities also lost activity, but by very different amounts: algorithmic communities contracted most severely (around 80% from their 2020 peaks), disciplinary by 30–50%, and practice-based by about 40% (per-community trajectories in Supplementary Fig. S1). Because shrinking a network can move structural metrics on its own, the question is whether the divergences below reflect genuine reorganisation or mere size reduction. Two facts point to reorganisation: practice-based communities lost a similar user share ($\sim$40%) yet barely changed structurally, and the structural patterns track codifiability rather than the size of the decline. Algorithmic communities thus still look healthy on volume while their connection structure has already reorganised—isolating that volume-invisible loss is what the difference-in-differences design takes on next.

4.1.2 Bonding and Bridging Capital Trajectories

Supplementary Table S2 presents the pre- versus post-GenAI group-averaged values across the four structural dimensions from Section 3.2; year-by-year trajectories are in Supplementary Fig. S1.

Bonding rises in algorithmic communities by +0.104 in modularity, while disciplinary and practice-based stay within $\pm$0.02: interaction increasingly stays within tightly-knit topic clusters rather than crossing them—fragmentation along knowledge subdomains, not merely an activity decline. Bridging diverges most sharply on the E-I index, where algorithmic communities fall by -0.224, flipping from already-bonding-dominant to deeply inward, while the other groups barely move. Reachability tells the same story: algorithmic average path length rises by +0.276

against flat values elsewhere, so knowledge must travel through more intermediaries. Betweenness-centrality Gini has a scale-dependent baseline (Borgatti and Everett, 2000), so the reading is across groups: algorithmic BC Gini stays flat while the others drift down, meaning brokerage roles spread out elsewhere but not in algorithmic communities. The divergence is three-dimensional—direction (bonding up, bridging down), group (only algorithmic), and timing (acceleration after 2022)—and all three point to the same reading: procedurally codifiable knowledge, where LLMs substitute most reliably, takes the hit while the rest of the platform barely shifts.

4.2 RQ2: Impact of GenAI Disruption on Social Capital

4.2.1 Parallel Trends Validation

Supplementary Table S3 reports the parallel trends tests. For each indicator, we regress the outcome on a linear time trend interacted with treatment status, using pre-GenAI data (2016–2022).

Four indicators pass: modularity, average path length, Top-20 degree share, and degree Gini. The other six—including the E-I index and clustering—diverge before treatment. Throughout the rest of this section, "DiD-identified" refers to indicators that pass this parallel-trends check and can therefore carry a causal reading; estimates on the other six are read as suggestive, consistent with the identified pattern rather than as independent evidence.

4.2.2 Algorithmic Bonding Rose and Reachability Worsened After 2022

Table I presents the two-way fixed effects (TWFE; community and year fixed effects) estimates for all structural indicators. Each outcome has two coefficients: β_{High} (algorithmic vs. practice-based) and β_{Medium} (disciplinary vs. practice-based).

Table I: DiD estimates of GenAI's effect on social capital structural indicators.

	Algorithmic (β_{High})			**Disciplinary (β_{Medium})**	
Indicator	TWFE est. (SE)	Wild p	SCM est.	TWFE est. (SE)	Wild p
Bonding capital					
Modularity†	$+0.085^{***}$ (0.029)	< 0.001	$+0.158$	-0.006 (0.017)	0.808
Bridging capital					
Avg Path (GCC)†	$+0.294^{***}$ (0.111)	< 0.001	$+0.510$	-0.001 (0.071)	0.934
BC Gini	$+0.020^{**}$ (0.008)	< 0.001	—	$+0.012^{**}$ (0.005)	< 0.001
E-I Index	-0.136^{**} (0.055)	0.130	—	$+0.030$ (0.030)	0.684
Concentration					
Degree Gini†	-0.037^{***} (0.013)	< 0.001	—	$+0.013^{*}$ (0.007)	0.310

Note. †Passes parallel trends. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The two DiD-identified indicators, modularity and average path length, hold steady across all three identification frameworks; wild-cluster-bootstrap p-values for β_{High} are < 0.001 on both, more conservative than the cluster-robust p-values in the TWFE column.[4] The E-I index, BC Gini, and degree Gini appear in the TWFE column but did not pass parallel trends (the E-I index also fails wild-bootstrap inference), so we read these three as suggestive, not identified. Supplementary Table S4 summarises the full robustness suite: modularity passes every layer except a milder 2021-placebo effect (smaller than the main +0.085), consistent with the mid-2021 Copilot / GPT-3 rollout rather than a violation of the 2023 identification, while average path length passes every layer including both placebo years. The two effects also do not hinge on any single treated unit—they appear independently in both algorithmic communities (Figure 2), and the synthetic control reproduces them for AI.SE alone with DataSci.SE excluded from the donor pool.

For disciplinary communities, the suggestive associations are much weaker. Among the indicators in Table I, only BC Gini reaches conventional significance. Modularity and average path length—the two DiD-identified indicators—are not significant for this group. The causal evidence thus sits squarely on the algorithmic group, with disciplinary communities showing no echo of it on the main indicators.

4.2.3 Effects Concentrate at the 2023 ChatGPT Release Window

Figure 3 shows event-study coefficients for modularity and average path length. For both, pre-treatment coefficients hover near zero, visually confirming parallel trends. After treatment, modularity spikes in 2023 (90% CI excluding zero), then fades by 2025 as network contraction constrains partitioning; average path length shows significant effects in 2023 and 2024, with 2025 attenuation likely reflecting the very small network sizes by then. The effects thus appear as a discontinuity bound to the year ChatGPT became publicly available rather than a continuation of pre-existing drift, so the timing itself carries part of the identifying weight.

[4] SCM estimates exceed TWFE estimates (+0.158 vs. +0.085 for modularity; +0.510 vs. +0.294 for path length) because TWFE absorbs DataSci.SE's parallel treatment effect into the year fixed effects while SCM compares AI.SE against an untreated synthetic counterpart.

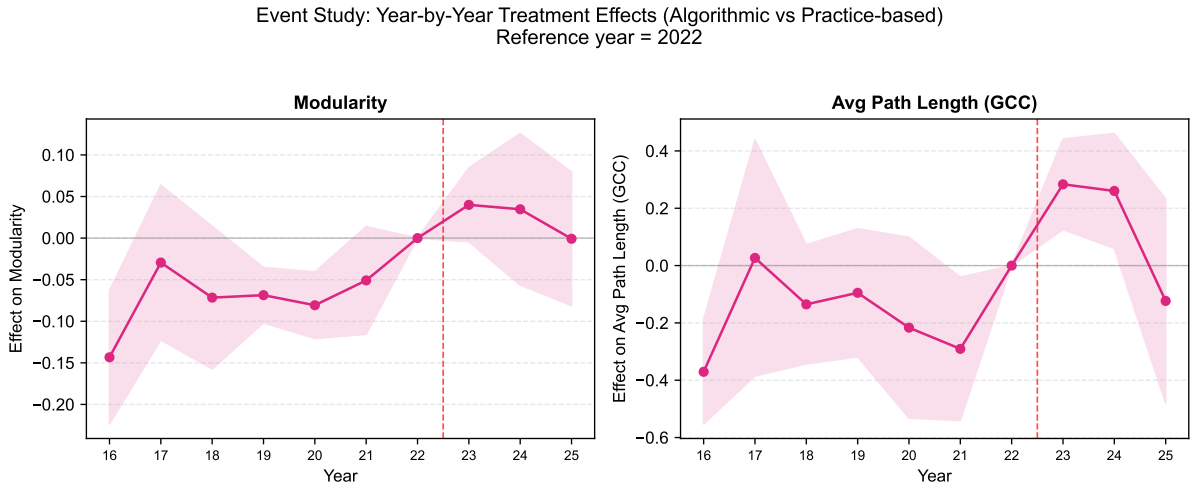


Figure 3: Event-study estimates for the two DiD-identified structural outcomes. Points show year-specific treatment coefficients for algorithmic communities relative to 2022; shaded bands show 90% confidence intervals. The dashed vertical line marks the transition to the post-GenAI period.

4.2.4 Selective Exit of Newcomers and Peripheral Askers

To test whether the community-level reorganisation reflects selective peripheral exit, we build a user-level panel from the same 70 community-year observations. We compute four participation-dynamics indicators: first-time asker share; newcomer 1-year retention; low-activity retention (users with 1–2 events, a peripheral proxy that avoids reputation-snapshot reverse causality); and high-activity retention (users with ≥ 10 events, a core benchmark). We restrict the pre-treatment window to 2018–2022 to avoid warm-up bias from the 2016-08 start of the data dump.

Table II: User-level participation dynamics: pre- versus post-GenAI group means.

Dimension	Indicator	Group	Pre	Post	Δ
Entry	First-time asker share	Algorithmic	0.859	0.806	-0.053^{**}
		Disciplinary	0.744	0.677	-0.067^{*}
		Practice-based	0.800	0.742	-0.058^{*}
Newcomer	1-yr retention	Algorithmic	0.115	0.060	-0.055^{***}
		Disciplinary	0.127	0.097	-0.031^{**}
		Practice-based	0.069	0.051	-0.018^{*}
Peripheral	Low-activity retention	Algorithmic	0.095	0.062	-0.033^{***}
		Disciplinary	0.125	0.102	-0.023
		Practice-based	0.112	0.105	-0.008
Core	High-activity retention	Algorithmic	0.577	0.541	-0.036
		Disciplinary	0.632	0.629	-0.003
		Practice-based	0.778	0.812	$+0.034$

Note. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Two findings stand out. Newcomer and peripheral retention drop selectively in algorithmic

communities: newcomer 1-year retention falls −5.5pp (−47% relative, $p < 0.001$) with a dose-response gradient (disciplinary −3.1pp, practice-based −1.8pp), and low-activity (1–2 event) retention falls significantly *only* in algorithmic communities ($p < 0.001$). Core ($\geq$ 10 event) retention is unchanged everywhere, confirming that the exit is structurally selective. The first-time asker share also declines, but more evenly across groups, suggesting LLM substitution shapes *who returns* more than *who first arrives*. Figure 4 plots the annual trajectories behind this pre–post comparison.

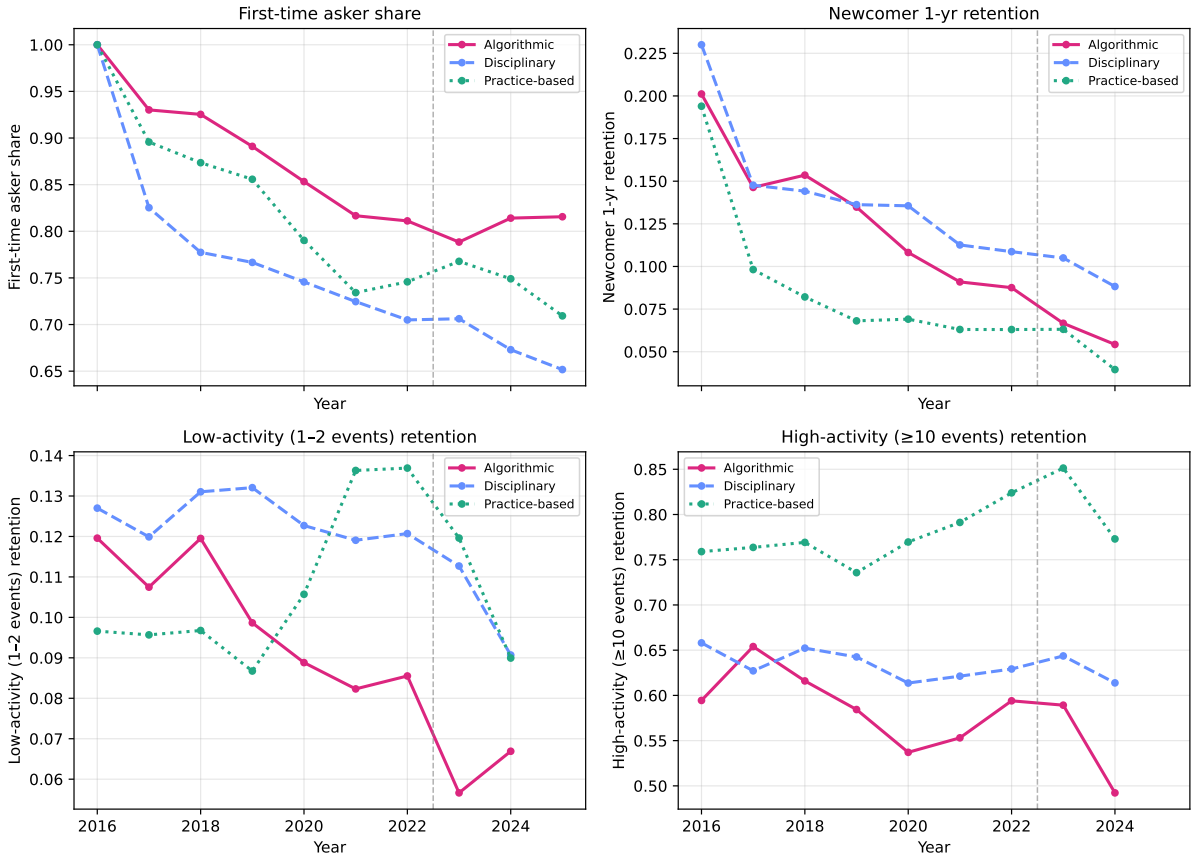


Figure 4: User-level participation dynamics by community group, 2016–2025. Lines show annual group means for first-time asker share and one-year retention among newcomers, low-activity users, and high-activity users. The vertical marker denotes the start of the post-GenAI period.

The user-level evidence thus directly supports the selective-peripheral-exit mechanism implied by the network-level estimates: the algorithmic core shrank in absolute count even though per-member core retention stayed stable, because newcomers no longer arrive at the rate needed to replenish it.

4.3 RQ3: Structure-to-Function Pathway

Descriptive statistics of the QA-latency change (the functional row $\tau_{0.5}$ of Supplementary Table S2) show all three groups saw latency rise after GenAI by 2–5 hours; the algorithmic increase is largest in absolute size but underpowered with only two communities, while the disciplinary increase reaches conventional significance. The rise is the functional fingerprint of the structural changes above: when modularity rises and average path length lengthens, a question reaches whoever can answer it through more intermediaries, so the wait grows even when the answerers are still present.

4.3.1 Panel Regressions: Structure Predicts Function

Table III reports panel fixed-effects regressions of QA latency on the structural indicators.

Table III: Panel fixed-effects regressions of QA latency on structural indicators.

Dimension	**Predictor**	β	**SE**	p
Bonding	Modularity	$+36.16^{**}$	15.15	0.017
	Avg Clustering	$+80.21$	134.67	0.551
	Transitivity	$+10.82$	51.41	0.833
Bridging	BC Gini	-132.48^{*}	75.55	0.080
	E-I Index	-19.56^{**}	8.09	0.016
	Avg Path (GCC)	$+6.95$	5.66	0.220
Concent.	Top-20 Degree Share	$+21.92$	26.41	0.407
	Degree Gini	-98.30^{***}	37.75	0.009

Note. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Two structural indicators significantly predict QA latency after community and year fixed effects. For modularity, a 0.1-unit increase is associated with about a 3.6-hour rise in latency, consistent with stronger bonding trapping knowledge inside subgroups; for the E-I index, a 0.1-unit decline (toward bonding dominance) is associated with about a 2-hour rise, pointing at the same mechanism from the bridging side.

A combined model with modularity, the E-I index, and Top-20 degree share leaves modularity and Top-20 degree share significant while the E-I index is subsumed, suggesting modularity captures the bonding side while concentration adds an independent effect. These regressions bridge the structural reorganisation and the latency rise: the structural shift translates into a measurable knowledge-sharing slowdown.

4.3.2 First-Difference Correlations

Year-over-year first-difference correlations corroborate the panel-FE result in the predicted direction across all three indicators (full set in Supplementary Table S5). The structure-to-latency link thus holds within communities over time, not only across community-year cells, lowering the chance that it is driven by cross-community heterogeneity rather than the structural change itself. Figure 5 visualises the corresponding structure-to-function relationships.

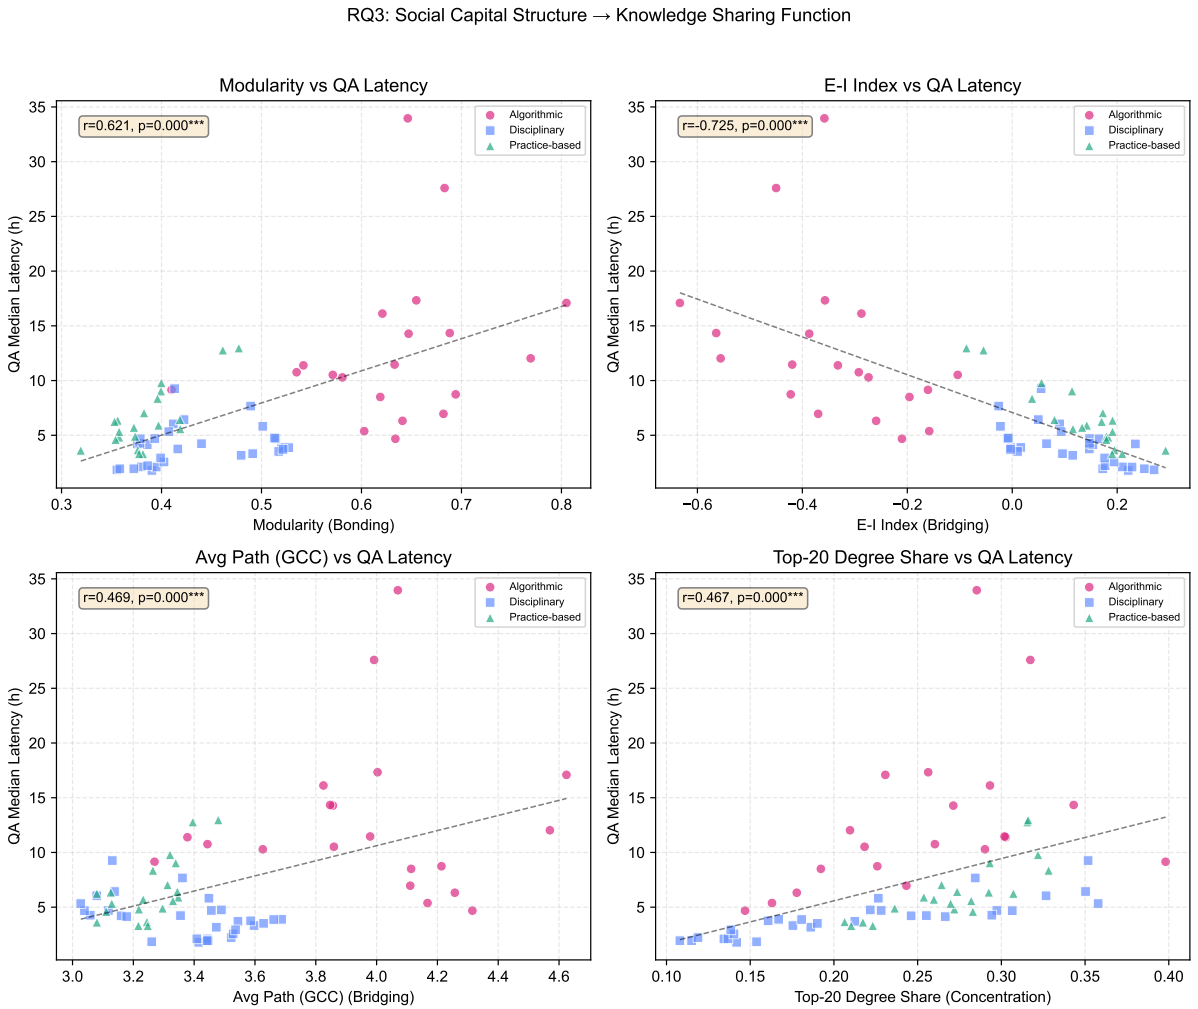


Figure 5: Associations between network structure and question-answering latency. Panels show the relationships of modularity, the E-I index, average path length, and Top-20 degree share with median time to first answer across community-year observations; fitted lines summarise the structure-to-function associations.

4.3.3 Tag Co-occurrence Density Fell Most in Algorithmic Communities

Tag-based content indicators show whether the structural reorganisation carried into knowledge content. Table IV summarises the pre- vs. post-GenAI comparison, and Figure 6 shows the group-averaged trends.

Table IV: Pre- versus post-GenAI group-averaged tag-based content indicators.

Dimension	Indicator	Group	Pre	Post	Δ
Breadth	Tag diversity	Algorithmic	0.269	0.466	$+0.197^{**}$
		Disciplinary	0.063	0.153	$+0.090^{***}$
		Practice-based	0.312	0.614	$+0.302^{**}$
	Tag concentration	Algorithmic	0.731	0.680	-0.051^{***}
		Disciplinary	0.591	0.605	$+0.014$
		Practice-based	0.516	0.499	-0.017
	Tag entropy	Algorithmic	6.97	7.29	$+0.32^{**}$
		Disciplinary	7.79	7.88	$+0.09$
		Practice-based	7.92	7.91	-0.01
Dynamics	Tag turnover	Algorithmic	0.349	0.411	$+0.062$
		Disciplinary	0.108	0.152	$+0.044$
		Practice-based	0.353	0.458	$+0.105^{***}$
	Tags/question	Algorithmic	3.18	3.11	-0.07
		Disciplinary	3.01	3.24	$+0.23$
		Practice-based	2.78	2.90	$+0.13$
Structure	Co-occur. density	Algorithmic	0.039	0.024	-0.015^{***}
		Disciplinary	0.059	0.044	-0.015^{**}
		Practice-based	0.0165	0.0150	-0.0015^{**}

Note. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Three patterns emerge. First, tag co-occurrence density declined most sharply in algorithmic communities (-37%), so knowledge domains grew more disconnected—a fragmentation of the inter-topic links through which integrative knowledge work happens—with disciplinary comparable and practice-based least (magnitudes and p-values in Table IV). Second, algorithmic communities showed falling tag concentration and rising entropy, a "residual question" effect: as LLMs absorbed the core, repetitive queries, the questions that remained spread more thinly across the topic space. Third, size-normalized tag turnover rose most in practice-based communities, reflecting organic renewal through new destinations and seasonal questions, while algorithmic communities showed only a non-significant rise, suggesting their narrowing landscape limited genuine renewal.

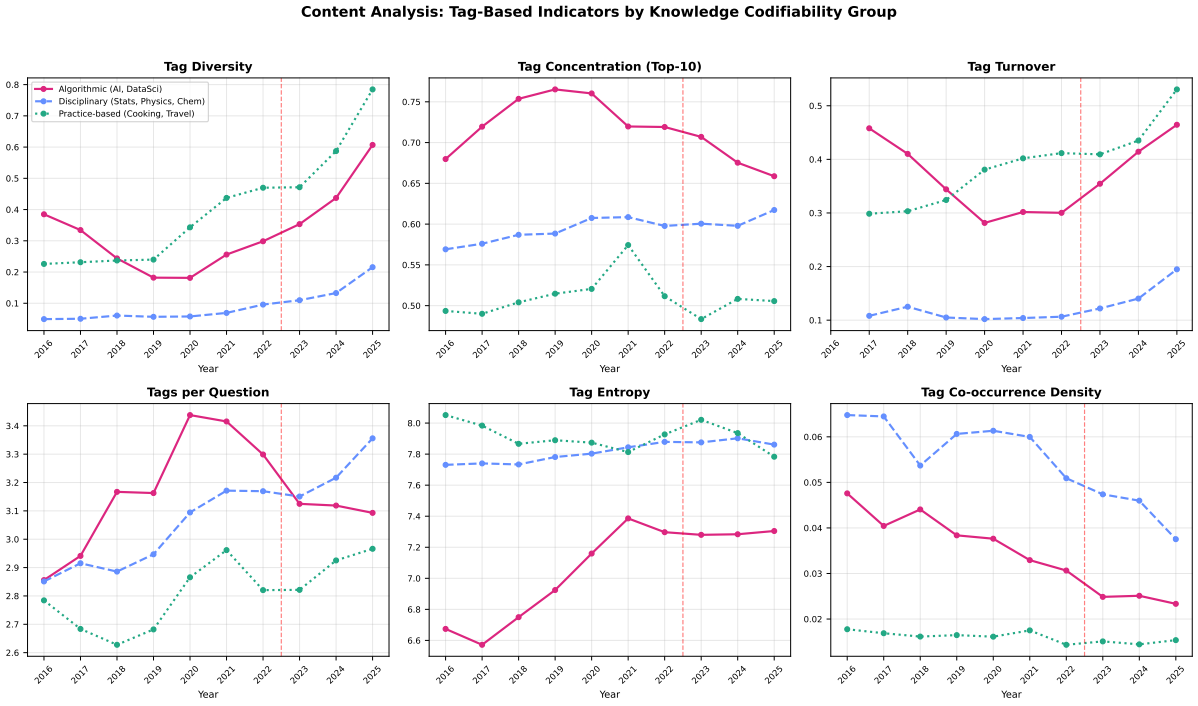


Figure 6: Tag-based content indicators by community group, 2016–2025. Lines show annual group means for tag diversity, concentration, entropy, turnover, tags per question, and tag co-occurrence density. The vertical marker denotes the start of the post-GenAI period.

This dual-level erosion—social ties and semantic ties thinning together—shows that the structural reorganisation reached into knowledge content, not just participation volume.

4.4 Hypothesis Verdicts

Across the three RQs, **H1, H1a, and H1b** (asymmetric restructuring proportional to codifiability) are DiD-identified through modularity and average path length on the algorithmic group, both surviving all five robustness layers, so **H2a** is supported. **H2b** (dose-response) is partially supported: disciplinary effects on the two DiD-identified indicators are indistinguishable from zero, so the gradient is clearer between algorithmic and the other groups than between disciplinary and practice-based, consistent with the procedural—not formal—subdimension driving the response. **H3** holds as an association rather than a DiD-identified effect, evidenced by the modularity-to-latency regression, the first-difference corroboration, and the collapse in algorithmic tag co-occurrence density.

5 Discussion

5.1 Interrupting Legitimate Peripheral Participation

The primary theoretical contribution is to extend LPP from an account of how peripheral participation accumulates into membership to an account of how that trajectory can be interrupted by a platform-external information substitute. LPP ordinarily begins when newcomers undertake legitimate, low-stakes activities at the edge of a practice (Lave and Wenger, 1991). In an online Q&A community, asking a routine question can perform this entry function: it exposes a newcomer to public reasoning, establishes an initial interaction, and creates the possibility of return. An LLM separates the informational outcome from this social process. A newcomer may receive an answer without creating the exchange through which peripheral participation would otherwise begin.

The data reveal both weaker renewal and weaker continuation at the periphery. In algorithmic communities, the share of first-time askers and newcomer one-year retention declined; low-activity-member retention fell in parallel, while high-activity-member retention remained stable (Table II). The first pattern is consistent with fewer entry-stage interactions, and the second with peripheral participation being less likely to continue. These are distinct margins: people who never post are absent from platform records, whereas retention is observable only among those who did enter. We also do not observe whether any individual consulted an LLM. The evidence therefore supports LPP interruption as a coherent explanation connecting the margins, but does not directly identify an individual's migration from community search to LLM use.

This extension matters because the interrupted trajectory leaves a network trace. Algorithmic modularity rose and average path length increased under the full identification suite, while the E-I index moved toward bonding dominance and high-activity retention stayed stable. The pattern is consistent with missing or discontinued peripheral interactions reducing bridging capital while recurring ties among incumbents preserve, and relatively intensify, bonding capital. In Lin's (2001) terms, the resources accessible through peripheral structural positions are lost when those positions fail to form or persist. LPP interruption thus links an individual information-seeking alternative to a community-level outcome: exit-driven structural closure.

This account differs from cumulative-advantage and Matthew-effect explanations, which locate peripheral weakening in competition for ties inside the community (Barabási and Albert, 1999; Merton, 1968). Here degree inequality falls rather than rises in algorithmic communities,

which is inconsistent with a handful of super-users simply absorbing the contraction. It also differs from echo-chamber mechanisms based on recommendation or motivated selection among users who remain active (Cinelli et al., 2021; Quattrociocchi, 2017; Valensise et al., 2023). Our claim is limited to structural closure: cross-topic ties become scarcer as peripheral interactions fail to form or continue; the study makes no claim about ideological polarization.

5.2 Procedural Codifiability as the Boundary Condition

The comparison across community types specifies when the proposed interruption is most consequential. The clearest divide is not between codifiable and tacit knowledge in general, but between tasks that current LLM outputs make readily executable and checkable and tasks that remain difficult to verify or depend on situated judgment. In this sense, procedural codifiability is a boundary condition on LPP interruption, not a parallel explanation of community decline.

Algorithmic communities. The two DiD-identified indicators show full structural reconfiguration, and the estimates survive every robustness layer. These communities also have by far the highest code density. Code provides an execution-verifiable output surface: a user can often run, inspect, and revise an answer without initiating a public exchange. Prior evidence that users accept plausible LLM answers to programming questions and that coding assistants reshape open-source participation makes entry-stage substitution especially credible in this setting (Kabir et al., 2024; Yeverechyahu et al., 2024). The observed alignment between code density, peripheral weakening, and network change is therefore consistent with the procedural-codifiability boundary.

Disciplinary communities. Formula density is highest here, yet the two primary structural effects are indistinguishable from zero. This is the critical contrast: formal expression alone does not reproduce the algorithmic pattern. A multi-step derivation can be explicit while remaining difficult for a non-expert to verify, so an LLM answer does not necessarily eliminate the value of community scrutiny. The data do not establish that disciplinary users return for verification, but the structural stability is consistent with a smaller execution-verifiable substitution surface (Ide and Talamàs, 2024; Nonaka and Takeuchi, 1995).

Practice-based communities. Activity also declined after 2022, but the bonding–bridging structure remained stable. Situated questions about cooking or travel depend more heavily on context and experience that resist complete codification (Polanyi, 1966); they therefore offer a low-exposure comparison rather than an untreated group. Their proportional contraction shows why volume decline alone is insufficient evidence of LPP interruption: the theoretical signature is

selective peripheral weakening accompanied by structural reorganisation. Overall, the comparison refines the tacit–codifiable axis by locating vulnerability in the execution-verifiable substitution surface rather than in explicit representation alone (Lundvall and Johnson, 1994; Nonaka and Takeuchi, 1995).

5.3 Structural Change Drives Knowledge-Sharing Function Decline

We next ask whether this structural reorganisation is associated with a functional cost. On response time, panel fixed-effects regressions show that higher modularity is associated with longer median QA response latency, and the within-community first-difference correlation points in the same direction. On knowledge integration, tag co-occurrence density declines steepest in algorithmic communities and least in practice-based ones, so the social and semantic fabrics fragment together. The residual algorithmic question stream also becomes more debugging-heavy, while code-block density rises as the stream shrinks. This composition is consistent with routine, self-contained questions being more substitutable than debugging that requires inspection of a specific artefact. These links are associational rather than DiD-identified, so we treat them as corroboration rather than causal estimates. They fit Centola (2010), where high modularity slows diffusion by confining flows within clusters, and Burt (2004)'s structural-hole logic, in which losing inter-cluster brokerage degrades a network's capacity to aggregate information.

Dual-level fragmentation sets this episode apart from prior community-decline research, whose landmark accounts document volume loss without parallel content restructuring. Iriberri and Leroy (2009) trace participation contraction to membership churn, weakening incentives, and platform competition, but their framework has no content-fragmentation term, assuming surviving members still cover the original domain at its original breadth; Butler (2001) likewise modeled membership and participation churn through benefit-cost and resource-load equilibria, again with no content shift. Later network-evolution work (Kossinets and Watts, 2009; Vlegels and Huisman, 2021) tracked how tie formation changes over time but did not link structural to semantic change. Our data shows what these models miss: the social fabric (bridging ties between subgroups) and the semantic fabric (tag co-occurrences between domains) erode *together*, far more strongly in algorithmic than practice-based communities. This dual-level erosion is the diagnostic signature separating GenAI-induced decline from generic life-cycle or churn decline, and implies that future models of community trajectories under technological shocks should carry a content-fragmentation term alongside the usual volume and tie-density terms.

Shan and Qiu (2025) found that after GenAI, answers on Stack Overflow grew shorter and easier to read even as their number rose, reading this as an individual-level efficiency effect—contributors learning from LLMs to answer more questions in more digestible form. We suggest a complementary structural explanation: once bridging capital erodes, answerers have less network access to diverse expertise, so shorter, shallower answers may partly reflect impoverished knowledge networks rather than individual behaviour alone. This is Granovetter (1973)'s weak-ties argument in reverse—when the cross-topic ties that brought several domains to bear on a question dissolve, answers written from inside a single cluster have less integrative material to draw on. The falling tag co-occurrence density we document is the content-level counterpart: as cross-topic ties dissolve, the questions answerers receive are themselves less integrative. The individual-level (Shan and Qiu, 2025) and structural accounts together likely explain more of the answer-shallowing than either alone. *This extends Lin's (2001) access-to-social-capital argument by tracing how network access converts into instrumental returns: structural reorganisation shows up as measurable functional decline in how quickly and how deeply the community answers an incoming question.*

5.4 Implications for Online Communities of Practice and Platform Governance

Taken together, our three findings—asymmetric restructuring, differential impact across community types, and a structural-to-functional cost—carry direct implications for governing online communities of practice in the GenAI era, organised around a single priority: AI assistance should protect rather than absorb the periphery's entry path, keeping the human community the canonical knowledge source.

The selective-exit finding implies that AI assistance should be anchored in community-verified knowledge and keep the periphery-to-core apprenticeship path visible, so routine help routes newcomers toward the community rather than away from it. This recasts the symbiotic-AI stance of Kasneci, Seßler, Küchemann, Bannert, Dementieva, Fischer, Gasser, Groh, Günnemann, Hüllermeier et al. (2023) and Lau and Guo (2023) as a community-grounded design principle rather than a generic add-on. The cost of failing falls on both ends: when peripheral askers route routine questions to an external LLM, the novice–expert bridging tie never forms, novices lose the social-cognitive learning of watching peers reason in public (Lave and Wenger, 1991), and experts lose the dialogue that prompts them to re-examine settled assumptions. Rather than let

an external chatbot intercept newcomers, platforms can host an in-community assistant that answers routine questions but grounds them in the community's verified knowledge base and routes the asker toward the people who produced it—turning a one-shot answer into an entry point for participation.

The differential-impact finding argues against uniform intervention. Because the substitution surface is procedural rather than codifiable-in-the-abstract, safeguards should concentrate on algorithmic communities, take a lighter, verification-oriented form in disciplinary ones, and, for practice-based communities, actively reinforce the situated, tacit exchange that made them resilient (Polanyi, 1966; Wenger et al., 2002)—turning tacitness into a durable moat rather than treating these communities as merely unaffected. As LLMs absorb routine Q&A, the community's distinctive function shifts toward mining, structuring, and co-constructing this practical knowledge, and platform design should follow that shift—surfacing dispersed practical knowledge and, to counter bridging erosion, actively connecting people across separated circles, for instance by recommending cross-topic collaborators on content proximity rather than existing ties (Granovetter, 1973).

Finally, because the structural shift emerged before activity reached its trough, structural indicators such as modularity and average path length can serve as leading early-warning signals alongside volume dashboards—the structural analogue of the lifecycle warnings Iriberri and Leroy (2009) and Butler (2001) proposed, and the inverse of the within-cluster amplification seen in echo-chamber recommenders (Cinelli et al., 2021; Valensise et al., 2023). A common logic ties these together: keeping the human community canonical makes AI assistance additive rather than substitutive. Because online Q&A is itself a major source of LLM training data (Kabir et al., 2024), sustained bridging erosion thins the very base future AI depends on—so the community-grounded stance is the configuration under which AI assistance and human knowledge production can co-evolve.

6 Conclusion

This study asked whether GenAI contracts online knowledge communities uniformly or selectively and what such selectivity means for community renewal. Across seven Stack Exchange communities, three levels of evidence support a common account. In algorithmic communities, newcomer entry and low-activity-member retention weaken while high-activity-member retention

remains stable; at the same time, modularity and average path length increase under the full DiD identification suite. Disciplinary and practice-based communities remain structurally stable on these primary indicators. These patterns are consistent with an interruption of legitimate peripheral participation: an external information substitute can satisfy routine needs without creating or sustaining the community interactions through which newcomers develop into members. Procedural codifiability bounds this interruption, and the resulting loss of peripheral ties reorganises social capital asymmetrically by weakening bridging while leaving incumbent bonding comparatively intact. The associated rise in response latency and decline in tag co-occurrence suggest that this structural closure reaches both knowledge-sharing function and content integration, although these latter relationships are associational.

A few limitations frame how the findings should be read, and each points to a productive next step. We cannot observe prospective participants who never post or users' decisions to consult an LLM, so LPP interruption is a theoretically grounded interpretation of convergent traces rather than a directly identified individual mechanism. Because LLMs can already answer some practice-based questions, our comparison group is best understood as low-exposure rather than untreated; the difference-in-differences estimates therefore identify differential effects across knowledge types rather than the absolute magnitude of GenAI's impact. Our social-capital measures are structural proxies drawn from interaction networks, not direct evidence on members' motivation, perceived belonging, or subjective ties (Li, He and Chen, 2022; Li, Yu and Du, 2025). The deterministic codifiability proxies capture two informative subdimensions but miss aspects of substantive answerability that human-expert ratings could reveal. Finally, all seven communities share the Stack Exchange platform, so replication elsewhere is needed. As AI capabilities change, the execution-verifiable substitution surface will also move. For platform governance, the practical implication is to keep the human community visible as the source of verified knowledge—grounding AI assistance in community contributions and protecting the periphery's entry path—so that obtaining a routine answer can remain a route into participation rather than a substitute for it.

References

Abadie, A., Diamond, A. and Hainmueller, J. (2010), 'Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program', *Journal of the*

American Statistical Association **105**(490), 493–505.

Anderson, A., Huttenlocher, D., Kleinberg, J. and Leskovec, J. (2012), Discovering value from community activity on focused question answering sites: A case study of Stack Overflow, *in* 'Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining', ACM, pp. 850–858.

Angrist, J. D. and Pischke, J.-S. (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, Princeton, NJ.

Barabási, A.-L. and Albert, R. (1999), 'Emergence of scaling in random networks', *Science* **286**(5439), 509–512.

Best, S. J. and Krueger, B. S. (2006), 'Online interactions and social capital: Distinguishing between new and existing ties', *Social Science Computer Review* **24**(4), 395–410.

Bighash, L., Oh, P., Fulk, J. and Monge, P. (2018), 'The value of questions in organizing: Reconceptualizing contributions to online public information goods', *Communication Theory* **28**(1), 1–21.

Blanchard, A. and Horan, T. (1998), 'Virtual communities and social capital', *Social Science Computer Review* **16**(3), 293–307.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R. and Lefebvre, E. (2008), 'Fast unfolding of communities in large networks', *Journal of Statistical Mechanics: Theory and Experiment* **2008**(10), P10008.

Borgatti, S. P. and Everett, M. G. (2000), 'Models of core/periphery structures', *Social Networks* **21**(4), 375–395.

Burt, R. S. (2004), 'Structural holes and good ideas', *American Journal of Sociology* **110**(2), 349–399.

Burtch, G., Lee, D. and Chen, Z. (2024), 'The consequences of generative AI for online knowledge communities', *Scientific Reports* **14**, 10413.

Butler, B. S. (2001), 'Membership size, communication activity, and sustainability: A resource-based model of online social structures', *Information Systems Research* **12**(4), 346–362.

Cameron, A. C., Gelbach, J. B. and Miller, D. L. (2008), 'Bootstrap-based improvements for inference with clustered errors', *Review of Economics and Statistics* **90**(3), 414–427.

Centola, D. (2010), 'The spread of behavior in an online social network experiment', *Science* **329**(5996), 1194–1197.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. and Starnini, M. (2021), 'The echo chamber effect on social media', *Proceedings of the National Academy of Sciences* **118**(9), e2023301118.

del Rio-Chanona, R. M., Laurentsyeva, N. and Wachs, J. (2024), 'Large language models reduce public knowledge sharing on online Q&A platforms', *PNAS Nexus* **3**(9), pgae400.

Ellison, N. B., Steinfield, C. and Lampe, C. (2007), 'The benefits of Facebook "friends": Social capital and college students' use of online social network sites', *Journal of Computer-Mediated Communication* **12**(4), 1143–1168.

Faraj, S., Jarvenpaa, S. L. and Majchrzak, A. (2011), 'Knowledge collaboration in online communities', *Organization Science* **22**(5), 1224–1239.

Granovetter, M. S. (1973), 'The strength of weak ties', *American Journal of Sociology* **78**(6), 1360–1380.

Hagberg, A. A., Schult, D. A. and Swart, P. J. (2008), Exploring network structure, dynamics, and function using NetworkX, *in* 'Proceedings of the 7th Python in Science Conference (SciPy 2008)', pp. 11–15.

Hou, I., Mettille, S., Man, O., Li, Z., Zastudil, C. and MacNeil, S. (2024), The effects of generative AI on computing students' help-seeking preferences, *in* 'Proceedings of the 26th Australasian Computing Education Conference', ACE '24, ACM, New York, NY, USA, pp. 39–48.

Ide, E. and Talamàs, E. (2024), 'Artificial intelligence in the knowledge economy', *Journal of Political Economy* . Articles in advance.

Iriberri, A. and Leroy, G. (2009), 'A life-cycle perspective on online community success', *ACM Computing Surveys* **41**(2), 1–29.

Kabir, S., Udo-Imeh, D. N., Kou, B. and Zhang, T. (2024), Is Stack Overflow obsolete? an empirical study of the characteristics of ChatGPT answers to Stack Overflow questions, *in* 'Proceedings of the CHI Conference on Human Factors in Computing Systems', ACM, pp. 1–17.

Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E. et al. (2023), 'ChatGPT for good? on opportunities and challenges of large language models for education', *Learning and Individual Differences* **103**, 102274.

Kossinets, G. and Watts, D. J. (2009), 'Origins of homophily in an evolving social network', *American Journal of Sociology* **115**(2), 405–450.

Krackhardt, D. and Stern, R. N. (1988), 'Informal networks and organizational crises: An experimental simulation', *Social Psychology Quarterly* **51**(2), 123–140.

Lau, S. and Guo, P. J. (2023), From "Ban It Till We Understand It" to "Resistance Is Futile": How university programming instructors plan to adapt as more students use AI code generation and explanation tools such as ChatGPT and GitHub Copilot, *in* 'Proceedings of the 2023 ACM Conference on International Computing Education Research', Vol. 1 of *ICER '23*, ACM, pp. 1–23.

Lave, J. and Wenger, E. (1991), *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge.

Li, S., Du, J. and Yu, X. (2024), 'Dynamics of cMOOC learner interactions in different social media: A longitudinal network analysis', *Interactive Learning Environments* **33**(3), 2460–2478.

Li, S., He, X. and Chen, J. (2022), 'Exploring the relationship between interaction patterns and social capital accumulation in connectivist learning', *Interactive Learning Environments* **32**(6), 2691–2712.

Li, S., Yu, X. and Du, J. (2025), 'Diversified social capital accumulation strategies for cMOOC learners: A longitudinal analysis of interaction patterns', *Interactive Learning Environments* . Online ahead of print.

Lin, N. (2001), *Social Capital: A Theory of Social Structure and Action*, Cambridge University Press, Cambridge.

Lundvall, B.-Å. and Johnson, B. (1994), 'The learning economy', *Journal of Industry Studies* **1**(2), 23–42.

Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G. and Hartmann, B. (2011), Design lessons from the fastest Q&A site in the west, *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', ACM, pp. 2857–2866.

Merton, R. K. (1968), 'The Matthew effect in science', *Science* **159**(3810), 56–63.

Movshovitz-Attias, D., Movshovitz-Attias, Y., Steenkiste, P. and Faloutsos, C. (2013), Analysis of the reputation system and user contributions on a Q&A website: Stack Overflow, *in* 'Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)', ACM, pp. 886–893.

Newman, M. E. J. (2003), 'The structure and function of complex networks', *SIAM Review* **45**(2), 167–256.

Nonaka, I. and Takeuchi, H. (1995), *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*, Oxford University Press, New York.

Polanyi, M. (1966), *The Tacit Dimension*, Doubleday, Garden City, NY.

Preece, J. and Shneiderman, B. (2009), 'The reader-to-leader framework: Motivating technology-mediated social participation', *AIS Transactions on Human-Computer Interaction* **1**(1), 13–32.

Putnam, R. D. (2000), *Bowling Alone: The Collapse and Revival of American Community*, Simon & Schuster, New York.

Quattrociocchi, W. (2017), 'Inside the echo chamber', *Scientific American* **316**(4), 60–63.

Safadi, H., Johnson, S. L. and Faraj, S. (2021), 'Who contributes knowledge? core-periphery tension in online innovation communities', *Organization Science* **32**(3), 752–775.

Serrano, M. Á., Boguñá, M. and Vespignani, A. (2009), 'Extracting the multiscale backbone of complex weighted networks', *Proceedings of the National Academy of Sciences* **106**(16), 6483–6488.

Shan, G. and Qiu, L. (2025), 'Examining the impact of generative AI on users' voluntary knowledge contribution: Evidence from a natural experiment on Stack Overflow', *Information Systems Research* . Published online 24 September 2025.

Stack Exchange Inc. (2025), 'Stack Exchange data dump', Internet Archive. Community-maintained archive; snapshot dated December 31, 2025. https://archive.org/details/stackexchange_20251231

Valensise, C. M., Cinelli, M. and Quattrociocchi, W. (2023), 'The drivers of online polarization: Fitting models to data', *Information Sciences* **642**, 119152.

Vlegels, J. and Huisman, J. (2021), 'The emergence of the higher education research field (1976–2018): Preferential attachment, smallworldness and fragmentation in its collaboration networks', *Higher Education* **81**(5), 1079–1095.

Vranić, A., Tomašević, A., Alorić, A. and Mitrović Dankulov, M. (2023), 'Sustainability of Stack Exchange Q&A communities: The role of trust', *EPJ Data Science* **12**(1), 4.

Wasko, M. M. and Faraj, S. (2005), 'Why should i share? examining social capital and knowledge contribution in electronic networks of practice', *MIS Quarterly* **29**(1), 35–57.

Watts, D. J. and Strogatz, S. H. (1998), 'Collective dynamics of 'small-world' networks', *Nature* **393**(6684), 440–442.

Welch, B. L. (1947), 'The generalization of 'Student's' problem when several different population variances are involved', *Biometrika* **34**(1–2), 28–35.

Wenger, E. (1998), *Communities of Practice: Learning, Meaning, and Identity*, Cambridge University Press, Cambridge.

Wenger, E., McDermott, R. and Snyder, W. M. (2002), *Cultivating Communities of Practice: A Guide to Managing Knowledge*, Harvard Business School Press, Boston, MA.

Wooldridge, J. M. (2010), *Econometric Analysis of Cross Section and Panel Data*, 2 edn, MIT Press, Cambridge, MA.

Xing, W. and Du, H. (2023), 'Mining large open online learning networks: Exploring community dynamics and communities of performance', *Journal of Educational Computing Research* **61**(1), 3–29.

Yeverechyahu, D., Mayya, R. and Oestreicher-Singer, G. (2024), 'The impact of large language models on open-source innovation: Evidence from GitHub Copilot', *arXiv preprint arXiv:2409.08379* .

Zhang, W. and Watts, S. (2008), 'Online communities as communities of practice: A case study', *Journal of Knowledge Management* **12**(4), 55–71.

Zhou, M., Cho, S. and Terveen, L. (2025), 'LLMs in Wikipedia: Investigating how LLMs impact participation in knowledge communities', *arXiv preprint arXiv:2509.07819* .

Data availability

The data analysed in this study are publicly available from the Stack Exchange Data Dump (`https://archive.org/details/stackexchange_20251231`). Derived datasets and analysis code are available from the corresponding author upon reasonable request.

Acknowledgements

The authors thank the Stack Exchange community for maintaining the public data dumps that made this research possible. During manuscript preparation, Claude (Anthropic) was used to support English-language and grammar refinement. The authors reviewed and edited all output and take full responsibility for the final text.

Author contributions

The author-contribution statement has been redacted for double-anonymous review. The complete statement is provided in the cover letter.

Competing interests

The authors declare no competing interests.

Ethical approval

Ethical approval was not required for this secondary analysis of publicly available, pseudonymised Stack Exchange data. The study involved no recruitment, intervention, interaction with users, or collection of identifiable private information.

Informed consent

Informed consent was not required. The study used only public, pseudonymised archival data and involved no direct participation.

Funding

Funding information has been redacted for double-anonymous review and is provided in the cover letter.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- 02supplementaryinformation.pdf