

# The Failed Migration of Academic Twitter: A Case Study of Precocious Adopters

Xinyu Wang\*, Sai Koneru\*, Sarah Rajtmajer

The Pennsylvania State University  
University Park, PA, USA

xzw5184@psu.edu, sdk96@psu.edu, smr48@psu.edu

## Abstract

Following changes in Twitter’s ownership in 2022 and subsequent changes to content moderation policies, many in academia looked to move their discourse elsewhere and migration to Mastodon was pursued by some. Our study examines the behavior of a self-organized group of early academic adopters who joined Mastodon following changes in Twitter’s ownership. Utilizing publicly available user account data drawn from a voluntarily curated list of academics, we track the posting activity of these early adopters on Mastodon over a one-year period. We also study follower-followee and interaction relationships to map internal networks, finding that the subset of academics who migrated to Mastodon were well-connected. However, this strong internal connectivity was insufficient to prevent users from returning to Twitter/X. Our analyses show that early adopters struggled to maintain engagement, shaped by Mastodon’s decentralized design and competition from alternatives such as Bluesky and Threads. The migration effort lost momentum after an initial surge, as most early adopters reduced activity or returned to Twitter. Our survival analysis further reveals that retention is strongly linked to diverse cross-server engagement and topic-specific server communities. Users with large pre-existing Twitter presence face significantly higher attrition risk, highlighting the challenge of replicating established social connections in a decentralized ecosystem. By examining the coordinated migration attempt of early adopters, we find that even this highly motivated group faced substantial challenges, suggesting that later or less coordinated efforts would likely encounter even greater barriers.

## Introduction

Social media platforms have become essential outlets for communication, information dissemination, and professional networking. Among these platforms, Twitter<sup>1</sup> (now known as “X”) has gained significant popularity and has been likened to a virtual *town square* for its role in facilitating public discourse (Mascaro and Goggins 2012). Academics and journalists populate sub-communities on Twitter

and use the platform for professional communication (Jurkowitz and Gottfried 2022; Malik, Heyman-Schrum, and Johri 2019). In this work, we use the term *Academic Twitter* not to denote a formal organization, but as a shorthand for a diffuse network of scholars who use the platform for professional discourse. Academics have extensively used Twitter to share research findings, announce academic events such as conferences, and contribute to wider scholarly discourse (Veletsianos 2012; Mohammadi et al. 2018; Holmberg and Thelwall 2014). However, after Elon Musk’s acquisition of the company and controversial changes to content moderation policies, many on so-called Academic Twitter sought alternatives (Cava, Aiello, and Tagarelli 2023).

Among academics, Mastodon emerged as a popular alternative, with communities organized around shared interests. Unlike traditional platforms, Mastodon operates on a federated model in which independent servers communicate through a shared protocol, allowing users to join different instances while remaining connected across the network (Zignani, Gaito, and Rossi 2018). This structure aims to reduce centralized control over internet services by distributing data and services across multiple nodes, enhancing privacy, security, and resilience against censorship and corporate control (Raman et al. 2019). Yet despite these advantages, this decentralized structure also leads to lower user engagement compared to centralized platforms like Twitter, where network effects and a larger, more active user base drive higher engagement levels (Raman et al. 2019).

The migration of users between social media platforms has attracted research attention, especially in relation to changes in platform policies, user experiences, the factors shaping user decisions, and the resulting implications for online communities. Social media movements often occur in response to significant changes in platform policies, ownership, user dissatisfaction, or perceived deficiencies in a platform (Xu et al. 2014). Qualitative analyses, including interviews and surveys of users migrating from platforms like Fandom to alternatives such as Tumblr, reveal that policy and value-based aspects are primary catalysts of change (Fiesler and Dym 2020). Beyond platform-level triggers, prior studies highlight the role of social dynamics in shaping migration behavior. User migration has been linked to perceived platform attractiveness and peer effects (Hou and Shiao 2020). In the context of Twitter’s ownership change,

\*These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>Note on terminology: Twitter rebranded to X in July 2023, which occurred midway through our data collection. For simplicity, we use the terminology Twitter and tweets throughout the paper.

migration to Mastodon reflected social influence, network structure, commitment, and shared identity, while remaining characterized by low Twitter account deletion rates, concentration of users in large Mastodon instances, and divergent posting behaviors across platforms (Cava, Aiello, and Tagarelli 2023; He et al. 2023). Recent work raises questions about sustainability, reporting lower engagement and posting activity on Mastodon compared to Twitter, and pointing to challenges in retaining active users within a decentralized architecture (Jeong et al. 2023).

Despite prior work characterizing general trends across the population of platform users, there remains a lack of targeted research concerning migration of narrower and highly interconnected communities characterized by shared motivations and coordinated efforts. We address this gap by providing a targeted case study of early academic adopters responding to platform changes, examining their migration patterns and behaviors to offer a more nuanced understanding of how these users navigate platform transitions. Given the significant role that Twitter plays in many academics' professional lives (Knight and Kaye 2016; Klar et al. 2020), academics represent a valuable subpopulation for studying migratory dynamics. Our work addresses the following research questions:

**RQ1:** How did early academic adopters navigate the migration to Mastodon, and what do longitudinal, network, and cross-platform patterns reveal about their retention?

**RQ2:** Which measurable user and community characteristics are associated with long-term retention on Mastodon?

Our findings indicate that the majority of academic researchers who moved to Mastodon failed to maintain their posting activity and follower engagement over the long term. Although migrated academic accounts formed strongly-connected internal networks, these connections were insufficient to sustain user engagement over time. Of note, however, communities that established field-specific server instances had higher user retention rate compared to those on general-purpose servers. Cross-platform analysis of most recent Twitter activity showed that most users returned to Twitter. Survival analysis confirms that users' persistence is associated with a number of factors including diverse cross-server engagement, early platform activity, choice of server. Furthermore, cross-platform modeling reveals a divergent pattern, with a longer history of prior activity on Twitter associated with higher attrition risk. These results underscore the challenges and opportunities associated with building and maintaining communities on new platforms.

## Related Work

### Cross-Platform Migration

User movement between social networking services has been studied as a migration process shaped by contextual factors, including dissatisfaction with platform experience, social influence, and perceived migration costs (Xu et al. 2014). Early work analyzed it using Push Pull Mooring framework with Push effects (e.g., policy changes) drive users away, while Pull effects (e.g., better features) attract them to alternatives. However, Mooring factors such as

switching costs often prevent complete migration (Chang, Liu, and Chen 2014; Lin and Wang 2017). Evidence from user surveys indicates that platform migration, such as movements from Fandom to Tumblr, is often driven by policy changes and misalignment in platform values (Fiesler and Dym 2020). Additionally, research has shown that user migration is influenced by factors such as perceived platform attractiveness and peer influence (Hou and Shiau 2020). Study of users migrating from Twitter to Mastodon post-Elon Musk's acquisition indicated that social influence, network structure, commitment, and shared identity significantly drive collective behavior changes in online communities (Cava, Aiello, and Tagarelli 2023). Analysis showed that Twitter to Mastodon migration is characterized by low Twitter account deletion rates, concentration of users in large Mastodon instances, varied posting behaviors across platforms (He et al. 2023). A similar study on Twitter-to-Mastodon migration reveals lower user engagement and posting activity on Mastodon, highlighting challenges in retaining active users due to its decentralized architecture (Jeong et al. 2023). These results are consistent with evidence that the transition is often incomplete with users maintaining cross-platform presence rather than fully substituting one platform to the other (Peña-Fernández, Larrondo-Ureta, and Morales-i Gras 2025). More broadly, recent work on platform switching from Twitter shows that users disperse across multiple destination platforms rather than relocating to a single successor platform (Radivojevic et al. 2025).

Our work looks more deeply at the specific case of the relatively well-synchronized and coordinated movement of communities of academic users across platforms.

### Deplatforming as Involuntary Migration

Unlike voluntary platform migration, which is often motivated by user dissatisfaction or policy changes, deplatforming represents a form of forced migration where users are compelled to leave due to removal or suspension (Rogers 2020). A study of Reddit's 2015 ban shows that deplatforming can reduce the visibility of harmful content on mainstream platforms while simultaneously displacing communities to alternative or less moderated spaces (Chandrasekharan et al. 2017). Parler was deplatformed after the 2021 U.S. Capitol attack, while Gab expanded as a refuge for banned Twitter users; however, research shows that such platforms often attract smaller or more fragmented audiences and face challenges related to moderation and community stability (Aliapoulios et al. 2021; Zannettou et al. 2018). While distinct from voluntary migration, deplatforming likewise underscores how platform-level conditions shape user mobility and contextualizes the broader challenges communities face when establishing presence on alternative platforms.

### Characteristics of Mastodon

The decentralized structure of Mastodon and similar platforms leads to unique user engagement dynamics that are different from centralized platforms. Research suggests that users are drawn to Mastodon for the ease of protecting their information from data mining and the absence of algorithms

shaping their feeds (Lee and Wang 2023). However, this can also make content discovery more difficult (Zignani et al. 2019) and may contribute to network fragmentation (Kleppmann et al. 2024), shaping interactions and communication patterns that differ from centralized platforms (Brauweiler et al. 2025). The decentralized structure encourages strong, tight-knit communities within individual server instances but fewer connections across instances (Zignani, Gaito, and Rossi 2018). These characteristics pose challenges for migration as well. Our study examines how specialized communities—particularly academics—navigated these dynamics in their movement from Twitter to Mastodon.

## Data

Our analyses are primarily based on a public GitHub repository containing lists of academics on Mastodon<sup>2</sup>. These lists were curated by researchers from 50 distinct disciplines, who voluntarily added their own names and accounts. Each list corresponds to a research discipline, containing account details of researchers within that area, their Mastodon IDs, and a short description of their research. Some lists also included each user’s Twitter handle. For the purpose of longitudinal analysis, we focused only on scholars whose Mastodon accounts were already in the repository at the beginning of November 2022, noting that the Twitter-to-Mastodon migration began in late October 2022. That is, the users in our dataset were part of a *coordinated, synchronized migration effort* in November 2022.

In total, we collected the Mastodon profiles and corresponding public metrics, e.g., number of followers, following, and timestamps of most recent status for **7,542 academic user accounts** using a customized Mastodon API<sup>3</sup>. We collected data weekly from November 2022 to October 2023 allowing for temporal analysis of user activity and engagement over the span of a year. In addition, we collected the most recent Mastodon data related to follower and following accounts for each scholar to conduct network analyses based on fields of study and associated servers. Additionally, to capture dynamics beyond follower and followee connections, we collected the user interactions (replies, boosts, mentions, and favorites) covering the full observation window from November 2022 to October 2023. We collected activity for users in our sample to extract both outgoing interactions and incoming interactions. Note that this interaction dataset was collected retrospectively in December 2025.

For cross-platform analyses, we used self-reported Twitter handles to gather scholars’ profile information. Of 7,542 Mastodon scholars, a subset of 3,131 scholars accurately provided Twitter handles suitable for data collection. Using the Twitter (“X”) Basic API, we retrieved profile information, most recent tweet IDs, and engagement metrics. Additionally, we used the Twitter ID collection endpoint to extract the posting time of their most recent tweet.

Finally, to understand discussions surrounding Twitter migration to Mastodon on both platforms, we collected posts from both platforms during the one year window.

<sup>2</sup><https://github.com/nathanlesage/academics-on-mastodon>

<sup>3</sup><https://github.com/halcy/Mastodon.py.git>

Specifically, we collected tweets via Zeeschuimer<sup>4</sup> using the boolean search “mastodon (researcher OR researchers OR academic OR academics OR professor OR professors OR scholar OR scholars) lang:en until:2023-11-01 since:2022-10-01”. The set of academic-related terms was developed through a preliminary scan of tweets, ensuring coverage of both individual (e.g., “professor,” “scholar”) and collective (e.g., “academics,” “researchers”) identifiers. Since many tweets were included because the display name contained the word “mastodon” (as some users use their Mastodon account name as their Twitter display name) we further filtered out these tweets. We retained only those that contained relevant keywords in the main body of the tweet, resulting in a total of 1,497 tweets. We also collected posts through the Mastodon API using migration-related hashtags adapted from prior work (Jeong et al. 2024). We excluded the tag “MastodonSocial” because it is commonly used for general content on Mastodon rather than migration. From a total pool of 12,654 migration related posts, we then applied a second stage filter using the same set of academic keywords described earlier to identify content relevant to academic migration. This process yielded 129 posts in the final dataset.

## RQ1: Migration Dynamics and User Retention

Following, we present findings from multiple analyses conducted to understand the dynamics of the migration. First, we used the longitudinal data from Mastodon to understand user attrition and engagement trends over the year. Then, we used network analysis to understand the connectivity and distribution of academics across Mastodon servers. Next, we conducted a cross-platform analysis using the subset of users who provided Twitter handles comparing their activity patterns on Twitter and Mastodon. Taken together, these analyses identify patterns and potential factors associated with user retention, which in turn motivate the selection of features included in the survival analysis for RQ2.

### Longitudinal Patterns of Adoption and Retention

**Attrition of Academic Accounts on Mastodon** Figure 1 (Top) gives a bar plot illustrating the change in number of active users aggregated by month and a line plot showing the monthly retention rate based on the previous month. In this context, we considered a user to be active if they made at least a post in a given month. In the initial month of data collection (November 2022), there were 7,505 active scholars who had posted within that same month.<sup>5</sup> Over the next seven months, the number of active scholars experienced a decline, showing a minor resurgence in July 2023 before continuing its descent, ultimately reaching 2,398 by October 2023. We observe a substantial rate of attrition; each month, approximately 10% to 20% of scholars in our set stopped actively using the platform. The surge in July 2023 coincided

<sup>4</sup><https://github.com/digitalmethodsinitiative/zeeschuimer>

<sup>5</sup>1 scholar (April 2022), 4 (May 2022), and 32 (October 2022) had their last activity outside our collection window. We omit them from the bar plot but include them in subsequent analyses.




Figure 1: **Top:** Histogram showing number of monthly active users (active=at least one post per month). **Bottom:** Histogram showing distribution of user active spans (number of days they are on Mastodon aggregated by week).

with significant shifts in Twitter’s policy on limiting number of posts user can view and officially rebranded to “X” which may have prompted a renewed but temporary interest in Mastodon as alternative.

### Engagement with Academic Accounts on Mastodon

Given the central role of engagement in user activity, we look at the engagement with migrated academic accounts on Mastodon. Figure 1 (Bottom) presents the distribution of user active spans, which we use to characterize accounts into activity patterns for comparison across dimensions in later sections.

*Active users.* We define active users as scholars who engaged on Mastodon during the initial Twitter-to-Mastodon migration, which began in late October 2022. For our purposes, we’ve defined this initial period between October 15 and November 30, 2022. We categorize active users into four engagement groups based on the empirical distribution of active spans using the 25th, 50th, and 75th percentiles: one-time visitors; short-term, long-term, and persistent adopters. These data-driven cutoffs reflect the underlying heterogeneity in the duration of their activity on Mastodon following the migration. The days cutoffs are rounded to the nearest integer.

*One-time visitors.* These academics engaged only briefly at the beginning of the Twitter-to-Mastodon migration, with active spans falling below the 25th percentile of the overall distribution ( $n_{days} \leq 39$ ). Their minimal participation suggests exploratory behavior during the migration wave without continued investment in the platform.

*Short-term adopters.* This group falls between the 25th and 50th percentiles of active spans ( $39 < n_{days} \leq 222$ ). These users engaged beyond initial exploration but did not sustain participation long enough to become stable contributors, potentially reflecting unmet expectations or competing commitments on other platforms.

User Category	Subcategory	Count
Active	One-time visitors	1895
	Short-term adopters	1876
	Long-term adopters	1921
	Persistent adopters	1845
Inactive	One-time visitors	4
	Returned Users	1

Table 1: Number of users, by activity.

*Long-term adopters.* These users fall between the 50th and 75th percentiles of the distribution ( $222 < n_{days} \leq 329$ ). They demonstrated extended engagement following the migration but did not remain active through the full analytic period.

*Persistent adopters.* These users fall above the 75th percentile ( $n_{days} > 329$ ), corresponding to the longest active spans through the end of our analyses. Their sustained engagement suggests successful integration into Mastodon’s academic community and continued participation over time.

*Inactive users.* Academics in this category were already using Mastodon before the migration wave from Twitter. Despite their early presence on the platform, they remained inactive during the migration period. One user was inactive during the first month but engaged briefly afterward. Inactive users are not included in subsequent analyses. The distribution of users is given in Table 1.

We calculate the average following, follower, and post counts for each user category, aggregated by month as shown in Figure A1. Our analysis reveals a sharp increase in all three public metrics at the onset of the migration period, specifically from November 2022 to December 2022. However, beyond this period, the rate of increase for following and follower metrics slows considerably. For one-time, short-term, and long-term users, growth in following and follower counts stagnated beginning in early 2023. As these users were not regularly contributing content, their visibility and potential to attract followers were likewise limited. In contrast, persistent users continued to post regularly. Despite their consistent activity, however, growth in their connections remained slow. This slow growth highlights the challenges faced in establishing and expanding social ties and communities on Mastodon. The persistent users’ continued efforts to engage and post suggest a commitment to the platform, yet the limited increase in followers indicates potential barriers in user retention and community building.

### Field-specific Activity Differences among Academic Accounts on Mastodon

As discussed in Section Data, users self-identified their research disciplines in public lists. We aggregate these disciplines into 9 broad fields, consistent with existing taxonomies of academic disciplines (bepress, 2019). We adopt this taxonomy because its granular disciplinary structure aligns more closely with the self-reported research areas in our dataset, and it is widely deployed

in institutional repository metadata systems (e.g., <sup>6,7,8</sup>). Scholars identifying with multiple disciplines and open science researchers were categorized as multi-disciplinary researchers. Inactive users were excluded from subsequent analyses. Disciplinary statistics of the 7,537 academics in our dataset are provided in Table A1 in the Appendix.

We identify the proportion of one-time, short-term, long-term, and persistent users by field (see Table A2 in the Appendix). The highest proportion of persistent users is observed among multidisciplinary researchers. One possible explanation is that interdisciplinary scholars engage with discussions spanning multiple fields, which may provide a broader set of conversational opportunities and connections that support continued participation. This observation motivates the inclusion of a binary indicator for interdisciplinary research in our subsequent modeling of user retention.

### Network Structure and Community Formation

To understand the connectivity within and between different fields, we collected the complete follower-followee network and multiple types of interaction networks of all academic accounts on Mastodon, and constructed an internal connection network within each academic group, as well as a network aggregated by field of study.

**Internal Networks Aggregated by Field** Figure 2 shows the network connections across different fields. In this network, node size represents the number of academics in each field, while edge width indicates the frequency of connections between the user accounts from different fields. Given the nature of the field, certain fields such as *Social and Behavioral Sciences* are strongly connected to *Multi-disciplinary* field. Our analysis reveals that the academic migration has led to the formation of a well-connected internal network, suggesting that researchers actively attempted to establish a new community on Mastodon. However, despite these coordinated efforts, the migration still faced significant challenges, as reflected through the low retention rates reported in the previous section.

**Internal Networks Aggregated by Server** Prior studies have shown that the Mastodon’s server instance-based structure makes it easier for users to discover and connect with others within the same instance than outside the instance. This leads to stronger within instance ties compared to weaker links across instances (Zignani et al. 2019). Therefore, we analyzed the follower-followee network aggregated by Mastodon instances and the corresponding interaction networks among top servers, shown in Figure 3.

The academic users were spread across a total of 342 servers, with 9 server instances being isolates without any connecting edges to other instances, which were removed for clarity. Table A3 in the Appendix shows network centrality metrics for each server and the distribution of the users across different servers. It also includes the official description of the server instances, indicating whether they are

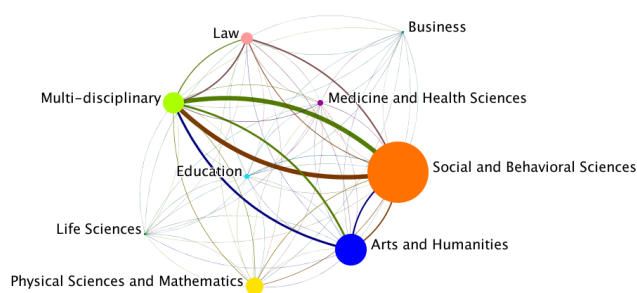


Figure 2: Internal field network (size: number of scholars in each field; edge weights: frequency of connections; edge color: same as the source; self-loops are omitted for better visualization.)

customized to a specific group of scholars. During the migration, different disciplines used different strategies: some used general-purpose servers like *mastodon.social* which host a wide variety of content beyond just academic content, while others took a coordinated approach on setting up their field-specific server instances such as *infosec.exchange* and *journa.host*. Figure 3 (Top) shows the follower-followee connections across different servers showing certain general purpose instances such as *mastodon.social* are popular among the tracked academic users (containing 21.3% of the total) than other general purpose servers like *mas.to*. We observe a central hub pattern in which the majority of follower–followee connections involve the *mastodon.social* server and other servers, forming a radiating structure. To examine interactions beyond follower-followee relationships, we further plot interaction networks among top servers with more than 100 scholars (highlighted in red in Figure 3 (Top)). Figure 3 (Bottom) shows the internal comment interaction network (boost, favorite, and mention networks are presented in Figure A2 in the Appendix). Node size is scaled by the total number of interactions for each server, and edge width is scaled by the number of interactions between server pairs. We find that even when servers are not directly connected through follower–followee relationships, they are connected through interactions such as comments, mentions, and boosts. This motivates the inclusion of both within-server and inter-server interaction in subsequent modeling, as these interaction-based connections capture federated engagement that is not visible in the follower–followee network alone. In particular, interaction networks reveal cross-server communication pathways that operate independently of stable follower ties.

**Comparison Across Communities.** To explore how a discipline specific server could influence user retention and engagement, we isolated the internal follower-followee network of accounts within *Information Security*. Figure A3 shows the internal network of the discipline, where most scholars congregated on a single server: *infosec.exchange* (highlighted in red). We compare the retention roles of scholars from the discipline with those from a discipline

<sup>6</sup><https://guides.lib.odu.edu/c.php?g=502940&p=5839940>

<sup>7</sup><https://stars.library.ucf.edu/faq.html>

<sup>8</sup><https://digitalcommons.lib.uconn.edu/about.html>

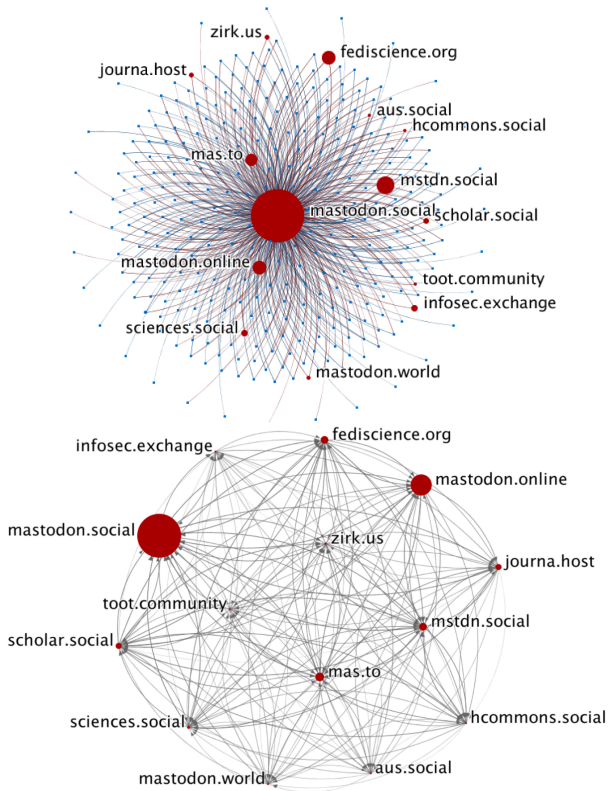


Figure 3: **(Top)** Server network constructed from internal follower–followee relationships. Isolated servers are removed for clarity. Node size is scaled by the number of scholars within each server. Servers with more than 100 scholars are highlighted in red. **(Bottom)** Internal comment interaction network among servers with more than 100 scholars. Node size is scaled by the total number of interactions for each server, and edge width is scaled by the number of interactions between servers. Internal boost, favorite, and mention networks are reported in Figure A2.

where users were concentrated in a general-purpose server and other sub-fields of similar size (Education and Communication & Media Studies), as shown in Figure 4. We see a significantly higher percentage of persistent users amongst information security scholars (40.9%) in comparison to other comparable disciplines. This suggests that establishing field- or discipline-specific communities or servers could foster a stronger sense of belonging, which may help retain scholars who have migrated, reducing the likelihood of their return to Twitter or seeking alternative platforms.

### Cross-Platform Patterns and Usage Shifts

We conducted a cross-platform analysis to analyze the state of the early adopters’ accounts on Twitter. Approximately 50% of the scholars voluntarily provided their Twitter usernames, which were accurate and traceable to their Twitter profile pages. Using the Official Twitter Basic API, we collected the profile information of these accounts. Despite some account deletions, we successfully retrieved 3,131

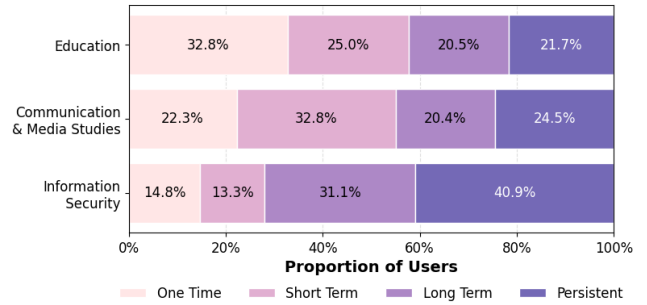


Figure 4: User retention breakdown for Information Security scholars compared to the sub-fields that are similar in size (Info. Sec.: n=264; Comm. & Media Studies: n=265; Education: n=244.)

Twitter accounts that could be matched to Mastodon accounts. Next, we retrieved the most recent posting times of these Twitter accounts to assess their activity status.

Based on their posting behavior, we categorized the Twitter account activity into the following statuses:

**Locked.** Scholar’s account is locked down, meaning their posts are private and cannot be accessed.

**Inactive.** Scholar’s latest tweet was posted before December 31, 2022, and their posts were made before this date, indicating no recent activity since the start of the year.

**Semi-active.** Scholar’s latest tweet was posted between December 31st, 2022 and November 1, 2023. This indicates some activity during the year but no recent posts after November, 2023, the end of our Mastodon tracking.

**Active.** Scholar’s latest tweet was posted after November 1st, 2023, indicating continuous activity and engagement on the platform.

Using the above categorization, we calculated the count and connection statistics for each category, and all statistics are presented in Table 2. We observe that 79.72% of scholars are still active on Twitter, indicating no signs of migration or intention to fully migrate. These active Twitter users have accumulated approximately 10 times more followers, which makes it more difficult to give up the platform. Some active users still make their Mastodon information accessible to the community. For instance, one scholar noted:

*“Hi ya’ll for those moving over to #Mastodon you can currently find me <Scholar’s Mastodon Account>.”*

	Count	Following Count	Follower Count
<b>Active</b>	2496	$\mu:1989.07, \sigma:3087.68$	$\mu:13208.76, \sigma:108495.54$
<b>Semi-active</b>	220	$\mu:1137.56, \sigma:1263.87$	$\mu:4467.28, \sigma:19152.90$
<b>Inactive</b>	91	$\mu:814.56, \sigma:1062.39$	$\mu:2401.46, \sigma:5203.50$
<b>Locked</b>	324	$\mu:1859.72, \sigma:2300.91$	$\mu:6347.10, \sigma:31583.98$

Table 2: Engagement statistics of 3,131 scholars (number of users per category with mean and standard deviation of following and follower counts).




Figure 5: Cross-platform engagement mapping of Twitter and Mastodon; counts of scholars classified by Twitter status (rows) and Mastodon usage (columns).

This group of scholars does not personally intend to migrate but aims to maintain connections with those who do intend to fully migrate to Mastodon. Furthermore, 7.03% of the scholars are semi-active users on Twitter. They did not stop using Twitter during the migration, but have reduced their presence gradually during a one-year grace period, potentially transitioning away from Twitter entirely. In contrast, 2.91% of the scholars still have accessible but inactive accounts since November 2022, indicating they will not return to Twitter. For example, one scholar’s last post reads,

*“I’m leaving Twitter but I’ve been advised not to delete my account in order to avoid possible misuse - so I’ll leave it here as a placeholder. See you at Mastodon!”*

Approximately 10.35% of the scholars chose to lock their Twitter accounts, often signaling their departure from the platform by leaving a message in their profile directing followers to their Mastodon accounts.

We cross-mapped Twitter activity levels with Mastodon role categories, as shown in Figure 5. Our analysis reveals that 39.95% of the sampled scholars are relatively active on both Mastodon (long-term or persistent users) and Twitter (active users). Among those who are inactive on Twitter or have locked their accounts, 56.14% are long-term or persistent users of Mastodon. However, this group of dedicated Mastodon users represents only 7.44% of the total scholars included in the cross-platform analysis. This suggests that while there is a core group of loyal users who have fully migrated to Mastodon, they make up a relatively small portion of the broader group of early academic adopters.

We also plotted the follower and following counts for the Mastodon and Twitter accounts of the same 3,131 scholars and the absolute differences in the metrics, as shown in Figure A4. The difference in connection scale is clear: while scholars who migrated to Mastodon made efforts to establish a new community, these networks remain far smaller than the connections many of them had built on Twitter over the years. Additionally, we plotted the Twitter account creation date against the total number of tweets, displayed in

Figure A5. The color of the points represents the total number of likes received by each account. A great proportion of scholars created their Twitter accounts between 2008 and 2012, indicating that the deep, long-standing networks, high levels of engagement, and rich posting histories developed on Twitter are not easily replicable or discarded. This entrenched presence highlights the difficulty of transitioning to a new platform like Mastodon without sacrificing years of accumulated influence and community ties.

In summary, we observe that the majority of scholars are either solely using Twitter or using both Twitter and Mastodon simultaneously. Comparing the follower and following counts across these platforms, Mastodon hosts a smaller community and Twitter continues to exert a strong pull on these early adopters.

## RQ2: Modeling Long-Term Retention

In previous sections, we have identified multiple factors that could potentially influence persistence on Mastodon. To understand the effect of these factors on the persistence of Mastodon usage, we model them using a Cox proportional hazards model (Cox 1972). We selected this model to account for right censored data (users remaining active beyond the observation window) and to estimate the hazard ratios of covariates without assuming baseline hazard distribution.

### Mastodon Survival Analysis

We used a Mastodon-centric model as our primary analysis and a cross-platform extension to understand the roles of Mastodon- and Twitter-derived features. The covariates used can be grouped as follows:

*Mastodon activity metrics.* These represent users’ platform activity level upon migration and social network metrics: the initial number of followers, following, and posts.

*Federated interaction diversity.* To capture how broadly a user engages across the Fediverse beyond their home server, we compute a network diversity metric from interactions. For each user, we extract the set of outgoing interactions and map each interaction to a target server. We then compute Shannon entropy (Lin 2002) over the distribution of target servers, yielding higher values when a user’s outgoing interactions are spread across many servers rather than concentrated within a single server. Our primary model uses diversity measured over the first 30 days of observed activity and we report sensitivity checks using alternative definitions (14 days, 60 days, and the first 50 outgoing interactions) in Figure A6 and A7.

*Account and server-level features.* We capture users’ profile settings, the network structure of their server, and the server type. Specifically, we consider whether the account is locked (binary), indicating if their posts are private and/or discoverable (binary), which indicates whether a user opted to being indexed to appear in search. At the server level, we identify the server type (binary) as either field-/topic-specific or general purpose by manually examining their public descriptions. A complete list of servers is provided in Appendix A4. At the server level, we define *in-degree ratio* as the proportion of inbound connections relative to total connections

and *out-degree ratio* as the proportion of outbound connections relative to total connections. Additionally, we categorize users based on their academic discipline, creating a binary indicator for Multidisciplinary scholars in Table A1 to test if these communities show a distinct retention patterns.

*Cross-platform metrics.* These capture users' premigration history on Twitter. We consider the followers and following counts on Twitter, the total number of Tweets since account creation, and the age of their Twitter account in days. In both models, for each account, the number of days between our first data collection (11-23-2022) and the account's last recorded posting activity within the study window was used as survival time. We treat a user having experienced the event of becoming inactive if their last posting occurs before the end of our data collection (10-27-2023). The counts covariates are log transformed and z-score normalized to handle the covariate distribution and comparable coefficients. We used a penalization ( $L_2 = 0.1$ ) to prevent overfitting.

### Mastodon-centric Model

Our first model used the complete dataset of user accounts and features extracted exclusively from Mastodon data. We filtered out accounts on isolated servers and those with incomplete data, reducing our sample from 7,542 to a final account count of 7,357 users. The model achieved a strong predictive capability (concordance of 0.69) with several factors significantly predicting the risk of becoming inactive on Mastodon. The initial number of a user's followers ( $HR = 0.72, p < 0.0005$ ) and number of posts ( $HR = 0.74, p < 0.001$ ) were strongest protective factors; one standard deviation increase in either was associated with around 30% lower likelihood of becoming inactive (refer to Appendix Figure A6). The initial number of accounts followed by a user did not have a significant effect ( $HR = 1.02, p = 0.262$ ).

For the server type and profile settings, being on a topic-specific server was a strong protective factor ( $HR = 0.84, p < 0.0005$ ), aligning with our observations from the information security server that tailored communities are associated with higher user retention. Notably, users identified as Multidisciplinary scholars exhibited higher retention rates with a hazard ratio of 0.59 ( $p < 0.001$ ). Furthermore, having a discoverable profile was also significantly linked to lower attrition ( $HR = 0.93, p = 0.009$ ), suggesting that being able to be found and connected is linked to persistent usage. In contrast, having the profile locked was not found to be a statistically significant predictor. The out-degree ratio was a risk factor ( $HR = 1.80, p < 0.001$ ) suggesting servers that are heavily outward looking are associated with lower user persistence. However, the proportion of inbound connections was not a significant predictor.

Federated interaction diversity is a robust protective predictor. In the 30-day diversity model, network diversity is associated with lower attrition ( $HR = 0.84, p < 0.001$ ), indicating that a one standard deviation increase in early cross-server interaction diversity corresponds to an approximately 16% lower hazard of becoming inactive. Sensitivity checks using alternative windows (14 days, 60 days, and first 50

outgoing interactions) yield consistent results (Figures A6), with effect sizes ranging from  $HR = 0.89$  (14 days) to  $HR = 0.75$  (first 50 outgoing interactions).

### Cross-platform Model

Our second model was fitted using the subset of 3,131 users for whom we had complete cross-platform features from both Mastodon and Twitter. This allowed us to understand how a user's Twitter activity related to their retention on Mastodon. The cross-platform model achieved concordance of 0.69 indicating that the addition of Twitter features does not add much overall predictive power. The addition of Twitter covariates provides a granular view of retention correlates. Notably, the association with interaction network diversity remains robust ( $HR = 0.86, p < 0.001$ ), indicating that cross-server engagement predicts retention independent of a user's prior status on Twitter.

Among covariates used, the association of metrics extracted from Mastodon remained largely consistent with the first model (see Figure 6). Initial followers on the platform ( $HR = 0.72, p < 0.0005$ ) and number of posts ( $HR = 0.77, p < 0.001$ ) were still strong factors linked to retention. Similarly, server type ( $HR = 0.79, p < 0.001$ ) and discoverability ( $HR = 0.86, p < 0.001$ ) remained significant and positively associated with persistence. However, unlike the Mastodon-centric model, server structural metrics and Multidisciplinary scholar indicator were no longer significant predictors of attrition in the cross-platform model.

Results of modeling covariates from Twitter revealed nuanced effects of some of the cross-platform user characteristics on persistence. A large number of Twitter following ( $HR = 1.13, p < 0.0005$ ) and followers ( $HR = 1.11, p < 0.0005$ ) were significant risk factors for becoming inactive on Mastodon. In contrast, a higher number of Twitter posts ( $HR = 0.92, p < 0.001$ ) and older Twitter account age ( $HR = 0.88, p < 0.0005$ ) were linked to higher retention.

The results from the cross-platform analysis shows a clear relationship with Twitter history. A large Twitter network is associated with higher attrition, while a long history of Twitter activity is linked to persistence showing divergent associations for these factors. Users with a well-established presence and high posting volume on Twitter were more committed to migration, but the size of their network may have been difficult to replicate on Mastodon, pulling them back to Twitter. This suggests that a successful migration of a community from a well-established platform like Twitter requires not only individual commitment but also a critical mass of connections that can migrate to the new platform.

### Contextual Factors and Platform Environment

The models identify measurable predictors of sustained use, yet migration and retention dynamics also unfold within a broader contextual environment that is not fully observable from activity traces alone. In this section, we examine three such contextual factors: the emergence of competing alternative platforms, public discourse surrounding the academic migration, and Mastodon's platform affordances. These factors help contextualize the conditions under which early adopters navigated migration and sustained participation.

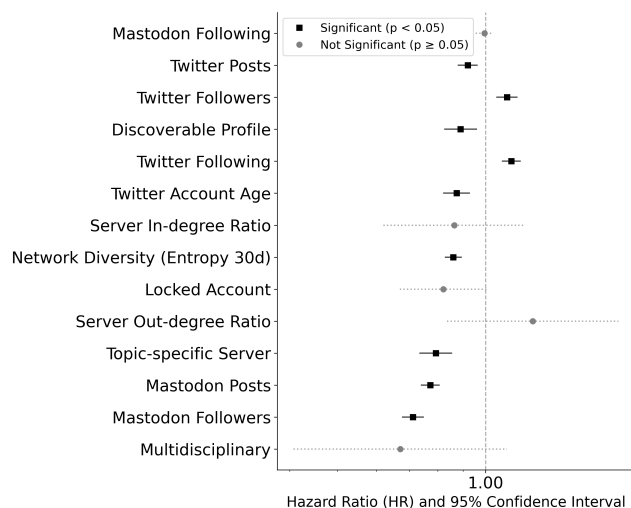


Figure 6: Plot showing hazard ratio and 95% CI for each covariate in the cross-platform model. Vertical dashed line at HR of 1.00 represents no effect.

### Alternatives to Mastodon

In addition to Mastodon, other centralized alternatives to Twitter have emerged in recent years, and many academics explored these venues. Unlike Mastodon, these platforms provided familiar interfaces and ease of transitioning. Those present on these alternative platforms often added their account information in their Twitter account descriptions to make it easier for others to discover and connect with them on these new platforms. To understand the extent to which academics were exploring the alternatives, we analyzed their usernames and user descriptions for mentions of Bluesky<sup>9</sup> and Threads<sup>10</sup> using keyword matching as well as user name pattern matching. For Mastodon and Bluesky, we also included scholars who use decentralized platform-style usernames as their Twitter usernames (e.g., <username>@<servername>).

Figure 7 shows an UpSet plot illustrating the number of users who explicitly mention each platform or a combination of platforms. In addition to 626 scholars mentioning Mastodon, 308 and 69 users also reference BlueSky and Threads, respectively, with some scholars providing account information for multiple platforms. This indicates that while these scholars have accounts on Mastodon, only about 20% actively direct their Twitter connections to it. The trend highlights that Mastodon is not the sole option for scholars migrating away from Twitter and faces significant competition from other emerging platforms. This competition further complicates Mastodon’s ability to establish itself as the primary platform for academic discourse.

<sup>9</sup><https://bsky.app/>

<sup>10</sup><https://www.threads.net/>




Figure 7: UpSet plot summarizing explicit cross-platform self-identification; intersection bars indicate counts of scholars mentioning Mastodon, Bluesky, and Threads (alone or jointly) in their profiles.

### Public Discourse on Academic Migration

To examine public discourse surrounding early academic migration, we analyze public posts from Twitter and Mastodon that explicitly reference academic migration, collected using a predefined keyword-based retrieval strategy (see Section Data). We plot the daily frequency of these posts and engagement counts over time aggregated by week in Figure 8. We observe similar patterns on both platforms: discussions about the migration in the academic context experienced a noticeable spike during the first two months of the movement on both platforms. After the migration period, conversations became less frequent on both platforms. However, they received more attention (favorites) on Twitter, indicating that the audience and connections of academics concerned about Mastodon largely remained active on Twitter.

Qualitative inspection of highly engaged posts during peak periods reveals recurring themes related to uncertainty about platform stability and long-term viability. For example, Mastodon struggled to accommodate the abrupt influx of users. On the peak day of discussion, November 18, 2022:

*“I have had a Mastodon presence for several months but haven’t used it. I still have a vain hope Twitter will sort itself out, but if not, I am <scholar’s Mastodon account>.”* –Twitter Post

However, this initial surge in discussion was short-lived. The frequency of related tweets sharply declined. Users expressed frustration over the platform’s lack of centralized theme and purpose, and its fragmented community structure.

*“My humanities folks have largely migrated to mastodon, but my science and public health feeds are still here on Twitter. More’s the pity.”* –Twitter Post

Despite these challenges, public efforts to encourage a shift toward Mastodon continued throughout 2023. For example, one scholar highlighted its advantages in a June 2023 post on Twitter:

*“Reasons to join us in the fediverse. 500 character limit. free edit button. Lots more #openaccess resources for #teachers, researchers, science communicators, lifelong learners. #STEM #education #open-science #K12 #HigherEd”* –Twitter Post

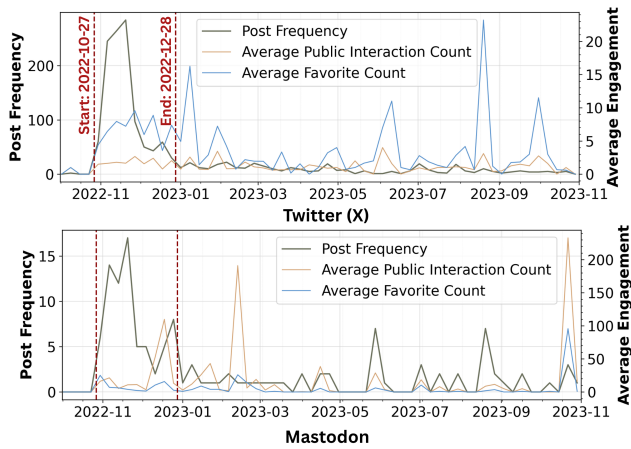


Figure 8: Dynamic frequency of Twitter and Mastodon conversations related to academics migration and corresponding engagement metrics (public interaction and favorite counts), aggregated by week.

Momentum was also visible within Mastodon itself. A widely circulated post from October 2023, which received 613 reblogs, signaled a clear departure from Twitter:

*Farewell, [bird emoji]! After 14 years on Twitter/X, we feel it's time to part ways. Recent developments there do not align with our core values. We stand for trustworthy information provision and are committed to open science. Follow us here for open science-related updates and news on our services and support for the academic community.[heart emoji] #OpenScience #AcademicLibrary #GoodbyeTwitter #ResearchSupport #EducationSupport –Masotdon Post*

Amongst the discussions on Twitter, we again observe concurrent mentions of other platforms such as Bluesky and Threads. These have diluted the migration as users explore various platforms for different purposes. For instance,

*“Hey if anyone’s still here on science Twitter and wondering where a lot of people have gone, they’re trying out bluesky and liking it. This time it’s real (not like mastodon). #academic” –Twitter Post*

Taken together, these posts illustrate a mix of optimism about Mastodon’s potential and recurring practical concerns users have faced during the migration.

### Mastodon’s Technical Affordances

Beyond the individual- and community-level factors identified above, structural characteristics of Mastodon, such as its moderation policies and discoverability mechanisms, can be considered additional factors contributing to the limited success of migration. In line with the network structure, moderation policy is also decentralized, as described<sup>11</sup>:

*Moderation in Mastodon is always applied locally, i.e. as seen from the particular server. An admin or*

*moderator on one server cannot affect a user on another server; they can only affect the local copy on their own server.*

This decentralized moderation stands in sharp contrast to Twitter’s centralized enforcement system. While local control empowers individual communities to define and enforce their own standards, it also creates inconsistencies across the network. For users accustomed to a uniform, centralized policy environment such as Twitter’s, these variations may generate uncertainty or limit trust in the platform.

Discoverability presents another technical challenge as shown in the models. Unlike Twitter, which supports platform-wide search and algorithmic amplification of content, Mastodon limits search to hashtags and profile information<sup>12</sup>. As a result, newcomers may struggle to discover communities, topics, or influential users, creating barriers to community formation and sustained engagement.

Taken together, these limiting affordances could help explain why Mastodon, despite its alignment with values of autonomy and community governance, has struggled to attract and retain large numbers of migrating users from Twitter. Importantly, these technical affordances are not experienced uniformly across the Fediverse. Because each Mastodon server operates with its own administrative history, governance structure, and moderation capacity, users’ experiences may vary substantially depending on the server they join. Servers established prior to the migration period may benefit from more mature moderation practices or greater administrative experience, which could influence user satisfaction and retention.

## Discussion and Conclusions

In this work, we explored the migratory dynamics of Academic Twitter communities to Mastodon. Through a series of analyses, we highlight that building a community on a new platform requires more than a coordinated effort and availability of alternatives. Results from longitudinal analysis showed that the initial surge in sign-ups did not translate into sustained long-term user engagement. This pattern is consistent with prior studies on Twitter to Mastodon migration in general user groups showing early migration waves often do not translate to stable long-term activity with many users maintaining cross-platform presence and continue using Twitter as a primary platform for engagement rather than fully switching (Jeong et al. 2023; He et al. 2023) and with broader findings that the transition away from Twitter remains incomplete for general audiences (Peña-Fernández, Larrondo-Ureta, and Morales-i Gras 2025). Most recent user activity on Twitter revealed that many academics in our dataset returned to using Twitter. Further analysis suggested that long-standing histories and strong communities on Twitter, in some cases spanning over a decade, proved too significant to overcome. While prior studies observed this general return to Twitter (Jeong et al. 2023; He et al. 2023), our cross-platform analysis with Twitter features suggests that, for established scholars, the difficulty of replicating existing networks is correlated with attrition. Without a

<sup>11</sup><https://docs.joinmastodon.org/admin/moderation/>

<sup>12</sup><https://docs.joinmastodon.org/user/network/#search>

complete migration of user history and established community, professionals and other specialized groups may naturally gravitate back to familiar platforms where their communities are already robust.

Likewise, several factors were linked to the failure of Mastodon to retain users. First, its decentralized structure creates a steeper learning curve compared to Twitter. Competition from more user-friendly platforms like Bluesky and Threads perhaps drew users away from Mastodon. Beyond direct Twitter substitutes, professional platforms such as LinkedIn continue to serve distinct functions centered around academic communication (Utz and Breuer 2019), which may further diffuse scholars' attention across platforms rather than concentrating migration on a single alternative. Our analysis adds to the understanding of challenges in migration of professional communities by identifying interaction diversity as a robust predictor of retention. Additionally, those who continued to post consistently on Mastodon did not enjoy the same level of engagement as on Twitter. That being said, select groups of academics stayed committed to Mastodon, making this migration less of a straightforward failure and more of a complex transition. Prior work suggests that social influence and a shared sense of identity or belonging can shape collective migration behavior (Cava, Aiello, and Tagarelli 2023). In our data, higher persistence among scholars on field or topic-specific servers is consistent with the possibility of community-specific servers supporting sustained engagement. This commitment is likely reinforced by self selection. For example, the most motivated early adapters may have actively sought out community-specific servers.

Future efforts that support migrations to decentralized platforms must address these core issues, foster community building by reducing friction, and consider network effects on user engagement. While the idea of decentralization of social media platforms holds significant promise, practical challenges persist that require more than just technical solutions. The Academic Twitter community's large-scale attempt to migrate from Twitter to Mastodon is significant—a strong professional network with a shared purpose positioned it as a hopeful candidate to succeed with such migration. Yet, as encouraging as these efforts and their partial successes are, evidence suggests that the process of transition is extremely challenging.

### Limitations

One limitation of this study is inconsistency in the timing of data collection across platforms. This discrepancy could potentially introduce biases in the analysis, as activity levels of scholars might vary depending on when the data was captured. Additionally, incorrect or incomplete social media handles provided by the scholars led to gaps in data that we could not manually resolve due to the scale of the dataset. While this subset provides substantial coverage of early adopters and enables examination of cross-platform engagement behaviors within this group, the analysis is descriptive in nature and not intended to support population-level claims. Moreover, changes to platform access (e.g., Twitter API rate limitations) restricted our ability to gather

complete data throughout the study. Similarly, full text keyword search is unavailable on many Mastodon instances unless it is enabled by the server administrator, so our reliance on hashtag-based search, which works without special server configuration, may have resulted in a smaller set of retrievable posts. In addition, the keywords used to collect public discourse may not capture all forms of academic self-description, which introduces the possibility of missing some relevant posts. However, we believe our findings are nevertheless substantial and provide valuable insights for cross-platform analyses. Notably, collected data reflects early groups of academic migrants, which are likely more indicative of collective migration efforts. However, we acknowledge the possibility of selection bias, as these patterns may not represent later migrants. This focus on early adopters implies that our findings primarily capture the initial, coordinated dynamics of migration rather than the full range of user experiences. Furthermore, our analysis treats a user's server as static which does not account for potential between-server migration within the Fediverse after the server choice at account creation. Finally, there may be case-specific factors at both the individual and community levels, including differences in local governance, historical development, or disciplinary norms, that cannot be captured holistically through platform traces alone. While such factors fall outside the scope of this study, we believe the observed patterns nevertheless provide a useful empirical account of early academic migration dynamics.

### Ethical Considerations

All account IDs in the original curated lists were voluntarily added by the users. Our data collection was restricted to publicly available interaction and profile metadata on Mastodon and Twitter at the time of collection. No private posts, direct messages, or other non-public information were accessed or retained. All identifiable information has been removed from the paper and associated research artifacts to further protect user privacy. Given the relatively narrow nature of the data collected in this study, we believe the risk of privacy violation and misuse to be minimal.

### References

- Aliapoulos, M.; Bevensee, E.; Blackburn, J.; Bradlyn, B.; De Cristofaro, E.; Stringhini, G.; and Zannettou, S. 2021. An early look at the parler online social network. *arXiv preprint arXiv:2101.03820*.
- bepress. 2019. Disciplines: Digital commons three-tiered taxonomy of academic disciplines.
- Brauweiler, M.; Pindiprolu, M. C.; Pfeffer, J.; and Brandes, U. 2025. Decentralized Discourse: Interaction Dynamics on Mastodon. In *Proceedings of the 17th ACM Web Science Conference 2025*, 358–368.
- Cava, L. L.; Aiello, L. M.; and Tagarelli, A. 2023. Drivers of social influence in the Twitter migration to Mastodon. *Scientific Reports*, 13(1): 21626.
- Chandrasekharan, E.; Pavalanathan, U.; Srinivasan, A.; Glynn, A.; Eisenstein, J.; and Gilbert, E. 2017. You can't

- stay here: The efficacy of reddit's 2015 ban examined through hate speech. *Proceedings of the ACM on human-computer interaction*, 1(CSCW): 1–22.
- Chang, I.-C.; Liu, C.-C.; and Chen, K. 2014. The push, pull and mooring effects in virtual migration for social networking sites. *Information Systems Journal*, 24(4): 323–346.
- Cox, D. R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2): 187–202.
- Fiesler, C.; and Dym, B. 2020. Moving across lands: Online platform migration in fandom communities. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1): 1–25.
- He, J.; Zia, H. B.; Castro, I.; Raman, A.; Sastry, N.; and Tyson, G. 2023. Flocking to mastodon: Tracking the great twitter migration. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, 111–123.
- Holmberg, K.; and Thelwall, M. 2014. Disciplinary differences in Twitter scholarly communication. *Scientometrics*, 101(2): 1027–1042.
- Hou, A. C.; and Shiau, W.-L. 2020. Understanding Facebook to Instagram migration: a push-pull migration model perspective. *Information Technology & People*, 33(1): 272–295.
- Jeong, U.; Sheth, P.; Tahir, A.; Alatawi, F.; Bernard, H. R.; and Liu, H. 2023. Exploring platform migration patterns between twitter and mastodon: A user behavior study. *arXiv preprint arXiv:2305.09196*.
- Jeong, U.; Sheth, P.; Tahir, A.; Alatawi, F.; Bernard, H. R.; and Liu, H. 2024. Exploring platform migration patterns between twitter and mastodon: A user behavior study. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, 738–750.
- Jurkowitz, M.; and Gottfried, J. 2022. Twitter is the go-to social media site for US journalists, but not for the public.
- Klar, S.; Krupnikov, Y.; Ryan, J. B.; Searles, K.; and Shmargad, Y. 2020. Using social media to promote academic research: Identifying the benefits of twitter for sharing academic work. *PloS one*, 15(4): e0229446.
- Kleppmann, M.; Frazee, P.; Gold, J.; Graber, J.; Holmgren, D.; Ivy, D.; Johnson, J.; Newbold, B.; and Volpert, J. 2024. Bluesky and the at protocol: Usable decentralized social media. In *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*, 1–7.
- Knight, C. G.; and Kaye, L. K. 2016. 'To tweet or not to tweet?' A comparison of academics' and students' usage of Twitter in academic contexts. *Innovations in education and teaching international*, 53(2): 145–155.
- Lee, K.; and Wang, M. 2023. Uses and Gratifications of Alternative Social Media: Why do people use Mastodon? *arXiv preprint arXiv:2303.01285*.
- Lin, C.-N.; and Wang, H.-Y. 2017. Understanding users' switching intentions and switching behavior on social networking sites. *Aslib Journal of Information Management*, 69(2): 201–214.
- Lin, J. 2002. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory*, 37(1): 145–151.
- Malik, A.; Heyman-Schrum, C.; and Johri, A. 2019. Use of Twitter across educational settings: a review of the literature. *International Journal of Educational Technology in Higher Education*, 16(1): 1–22.
- Mascaro, C.; and Goggins, S. P. 2012. Twitter as virtual town square: Citizen engagement during a nationally televised Republican primary debate. In *APSA 2012 Annual Meeting Paper*.
- Mohammadi, E.; Thelwall, M.; Kwasny, M.; and Holmes, K. L. 2018. Academic information on Twitter: A user survey. *PloS one*, 13(5): e0197265.
- Peña-Fernández, S.; Larrondo-Ureta, A.; and Morales-i Gras, J. 2025. Straddling Two Platforms: From Twitter to Mastodon, an Analysis of the Evolution of an Unfinished Social Media Migration. *Social Sciences*, 14(7): 402.
- Radivojevic, K.; Adams, D.; Laszlo, G.; Kery, F.; and Weninger, T. 2025. User migration in the Twitter diaspora. *EPJ Data Science*, 14(1): 36.
- Raman, A.; Joglekar, S.; Cristofaro, E. D.; Sastry, N.; and Tyson, G. 2019. Challenges in the decentralised web: The mastodon case. In *Proceedings of the internet measurement conference*, 217–229.
- Rogers, R. 2020. Deplatforming: Following extreme Internet celebrities to Telegram and alternative social media. *European Journal of Communication*, 35(3): 213–229.
- Utz, S.; and Breuer, J. 2019. The relationship between networking, LinkedIn use, and retrieving informational benefits. *Cyberpsychology, Behavior, and Social Networking*, 22(3): 180–185.
- Veletsianos, G. 2012. Higher education scholars' participation and practices on Twitter. *Journal of computer assisted learning*, 28(4): 336–349.
- Xu, Y. C.; Yang, Y.; Cheng, Z.; and Lim, J. 2014. Retaining and attracting users in social networking services: An empirical investigation of cyber migration. *The Journal of Strategic Information Systems*, 23(3): 239–253.
- Zannettou, S.; Bradlyn, B.; De Cristofaro, E.; Kwak, H.; Sirivianos, M.; Stringini, G.; and Blackburn, J. 2018. What is gab: A bastion of free speech or an alt-right echo chamber. In *Companion Proceedings of the The Web Conference 2018*, 1007–1014.
- Zignani, M.; Gaito, S.; and Rossi, G. P. 2018. Follow the "mastodon": Structure and evolution of a decentralized online social network. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12, 541–550.
- Zignani, M.; Quadri, C.; Gaito, S.; Cherifi, H.; and Rossi, G. P. 2019. The footprints of a "mastodon": How a decentralized architecture influences online social relationships. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 472–477. IEEE.

## Paper Checklist

### 1. General

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes, see Section Ethical Consideration.
- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes.
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes, we detail the approaches used in Research Methods section.
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? NA. In this work, we are focusing on a group of online users who participated within a coordinated migration. And we have no intention to generalize the claims to the overall population.
- (e) Did you describe the limitations of your work? Yes, see Limitation section.
- (f) Did you discuss any potential negative societal impacts of your work? Yes, please refer to the Ethical Considerations section of the article.
- (g) Did you discuss any potential misuse of your work? Yes, please refer to the Ethical Considerations section of the article.
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, Please refer to the Ethical Considerations section of the article.
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes.

### 2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? NA.
- (b) Have you provided justifications for all theoretical results? NA.
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA.
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? NA.
- (e) Did you address potential biases or limitations in your theoretical framework? NA.
- (f) Have you related your theoretical results to the existing literature in social science? NA.
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? NA.

### 3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? NA.
- (b) Did you include complete proofs of all theoretical results? NA.

### 4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA.
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA.
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA.
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA.
- (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA.

### 5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...

- (a) If your work uses existing assets, did you cite the creators? Yes.
- (b) Did you mention the license of the assets? NA.
- (c) Did you include any new assets in the supplemental material or as a URL? NA.
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? No, in this study we are not releasing any new data with personally identifiable information.
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? NA.
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA.
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? NA.

### 6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...

- (a) Did you include the full text of instructions given to participants and screenshots? NA.
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA.
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA.
- (d) Did you discuss how data is stored, shared, and de-identified? NA.

## Appendix

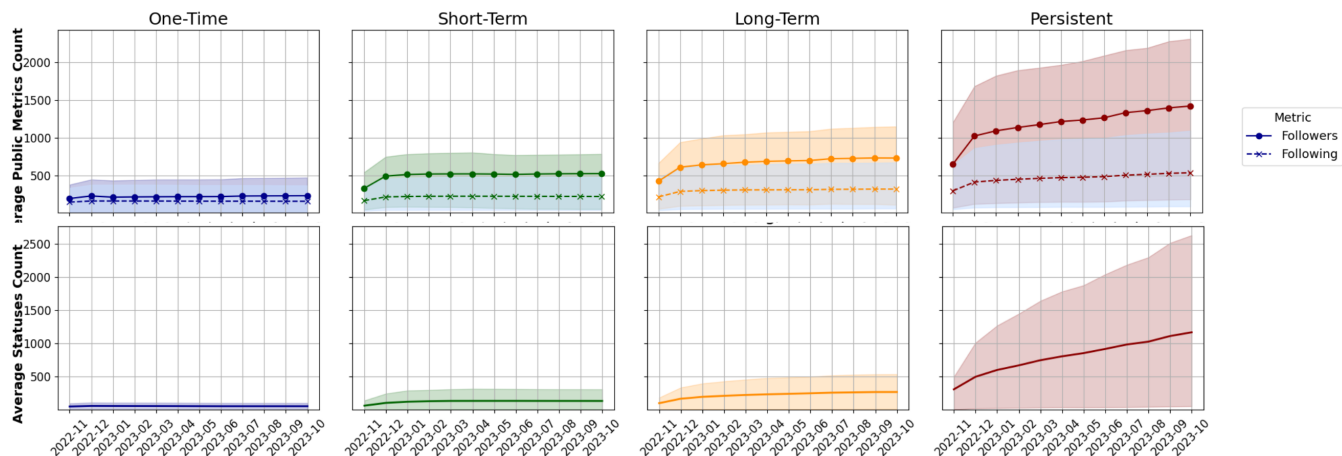


Figure A1: Monthly average total public metrics and statuses count for active users with 10–90% Range.

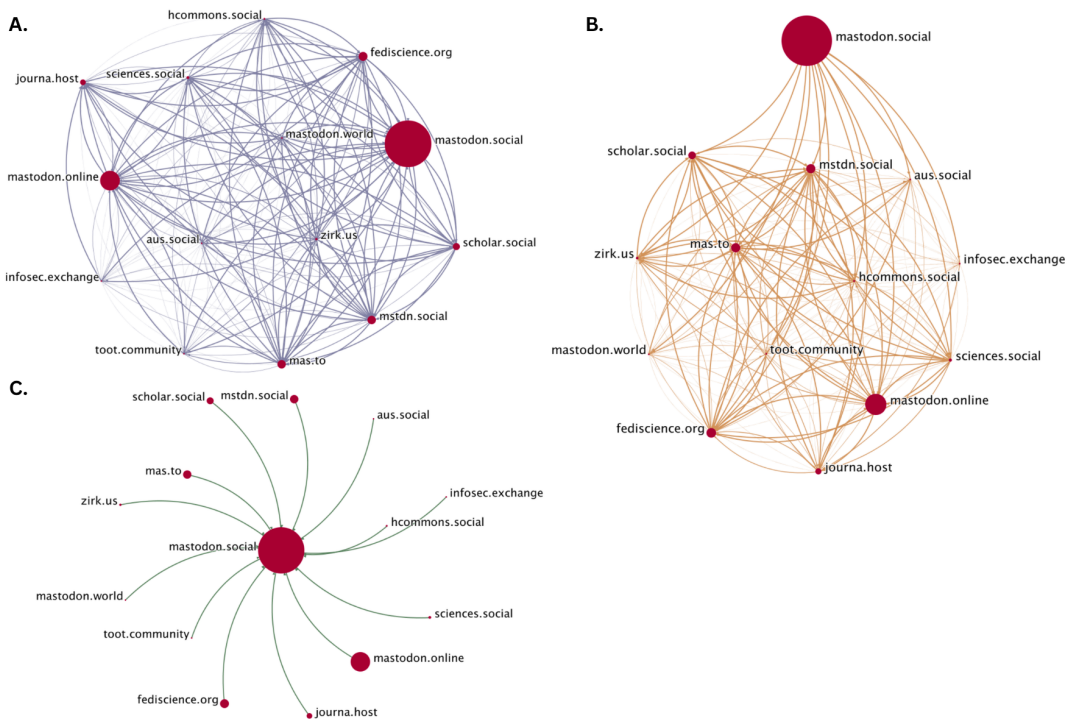


Figure A2: Internal interaction networks among servers with more than 100 scholars. A: boost; B: mention; C: favorite.




Figure A3: internal follower-followee network for Information Security community (removed isolates for better visualization).

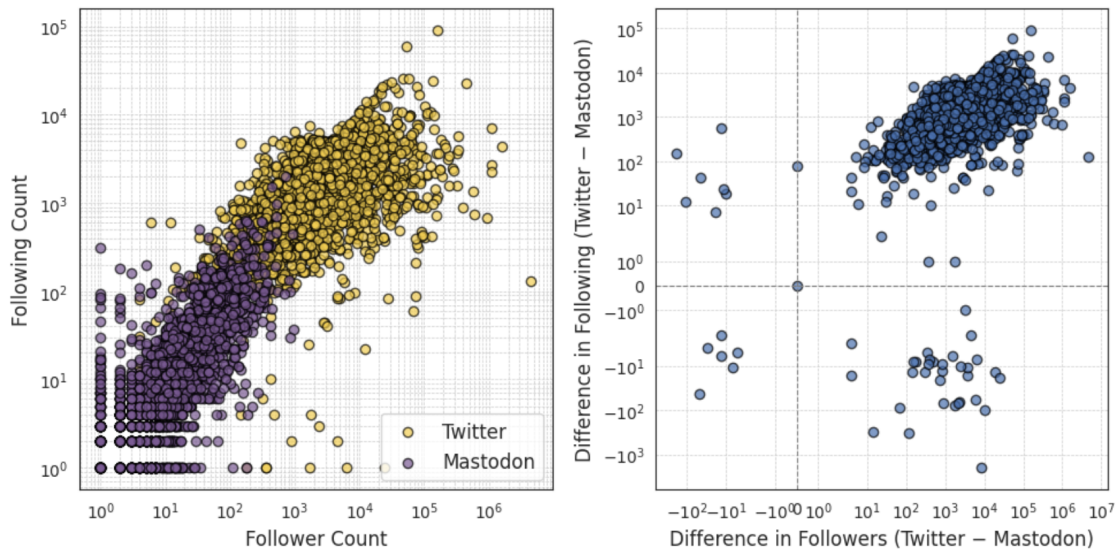


Figure A4: Connection metrics mapping and differences in metrics between mastodon and Twitter.

Field	Discipline	Count	Total
Arts and Humanities	Archaeology	126	1522
	Art and Design	8	
	Biblical Studies	4	
	Digital Humanities	208	
	Glams (Galleries, Libraries, Archives, Museums, Special Collections)	109	
	Literature Studies	142	
	Philosophy	220	
	Religion	56	
	Theology	27	
	History	554	
	History and philosophy of science	11	
	BookHistory	36	
Medieval studies	21		
Business	Marketing	4	4
Education	Education	244	244
Law	Law	587	587
Life Sciences	Bioinformatics	13	84
	Biophysics	8	
	Palaeogenomics	6	
	Plant Science	57	
Medicine and Health Sciences	Neuroscience	264	287
	Medical AI	23	
Physical Sciences and Mathematics	Chemistry	240	851
	Cheminformatics and Computational Chemistry	11	
	Geology	145	
	Meteorology	93	
	Nuclear fusion	1	
	Astrophysics	53	
	Atmospheric and air quality	41	
	Dendrochronology	3	
Information Security	264		
Social and Behavioral Sciences	Anthropology	50	2918
	Communication and Media Studies	265	
	Criminology	41	
	Genealogy	168	
	Geography	282	
	Linguistics	108	
	Political Science	364	
	Sociology	228	
	Asian Studies	26	
	Iran and middle eastern studies	5	
	Psychology	321	
	Public Policy	30	
	Technology Policy	152	
	Science and technology studies	41	
	Semantic Web	25	
Qualitative Research	35		
Journalism	777		
Multi-disciplinary	Multi-disciplinary	961	1040
	Open Science	79	

Table A1: Number of scholars by disciplines within major academic fields.

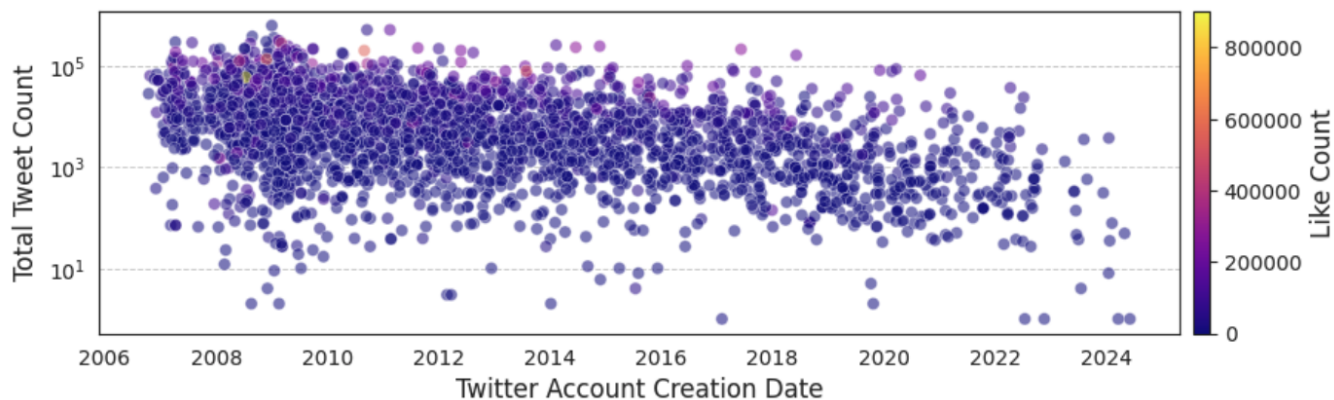


Figure A5: Scatter Plot of Scholar's Twitter Account Creation Date VS. Engagement Metrics.

Field	One Time(%)	Short Term(%)	Long Term(%)	Persistent(%)	Total	Percent	In-Degree	Out-Degree
Social & Behavioral Sciences	25.0	25.8	24.1	25.1	2918	38.72	98881	91958
Arts & Humanities	29.0	21.4	26.0	23.6	1522	20.19	38692	39986
Physical Sciences & Mathematics	22.1	23.4	27.6	26.9	851	11.29	11453	11980
Life Sciences	20.2	22.6	35.7	21.4	84	1.11	338	648
Education	32.8	25.0	20.5	21.7	244	3.24	2249	4171
Law	29.6	28.3	25.4	16.7	587	7.79	21664	27450
Medicine & Health Sciences	24.7	31.0	26.8	17.4	287	3.81	4292	4923
Multidisciplinary	18.6	25.1	27.1	29.2	1040	13.80	56371	52818
Business	25.0	50.0	0.0	25.0	4	0.05	60	66

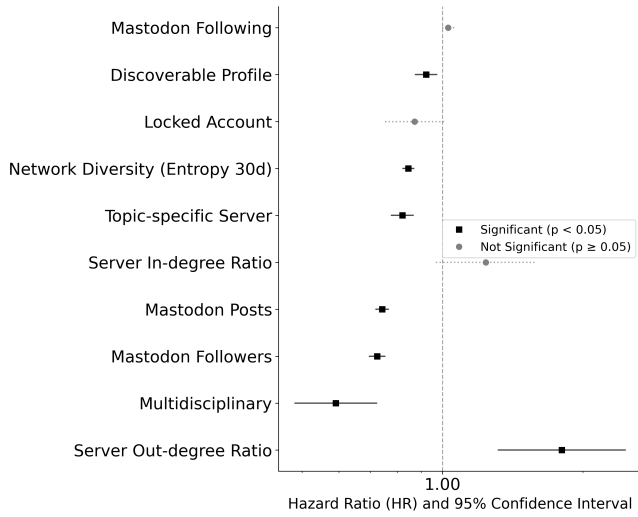
Table A2: User activity patterns by academic field (one-time, short-term, long-term, persistent) and field-level statistics (total users, percentage, and network centrality indicators).

Server	Count	Percentage	In-Degree	Out-Degree	Description	Scope
mastodon.social	1606	21.31%	147295	128525	The original server operated by the Mastodon GmbH non-profit.	Universal
mstdn.social	538	7.14%	9456	11947	A general-purpose Mastodon server with a 500 character limit. All languages are welcome.	Universal
mastodon.online	429	5.69%	6600	7172	A newer server operated by the Mastodon GmbH non-profit.	Universal
fediscience.org	415	5.51%	8220	9628	Fediscience is the social network for scientists.	Universal
mas.to	381	5.06%	5591	6534	Hello! mas.to is a fast, up-to-date and fun Mastodon server.	Universal
sciences.social	218	2.89%	5535	7460	Non-profit, ad-free social media for social scientists. Join thousands of social scientists here and across the fediverse.	Field-level
infosec.exchange	200	2.65%	1001	1362	A Mastodon instance for info/cyber security-minded people.	Discipline-level
scholar.social	193	2.56%	2842	3899	Microblogging for researchers, grad students, librarians, archivists, undergrads, high schoolers, educators, research assistants, profs—anyone involved in learning who engages with others respectfully.	Universal
journa.host	165	2.19%	3086	2717	The server for working journalists and news outlets on Mastodon. Home to active & retired journalists, media scholars, and a variety of news and journalism adjacent professionals.	Discipline-level
zirk.us	153	2.03%	2276	3087	#Newstodon Literature, philosophy, film, music, culture, politics, history, architecture: join the circus of the arts and humanities! For readers, writers, academics or anyone wanting to follow the conversation.	Field-level
mastodon.world	144	1.91%	1655	2213	Generic Mastodon server for anyone to use.	Universal
toot.community	107	1.42%	805	1202	A worldwide Mastodon instance from The Netherlands. Run by digital enthusiasts, inviting everyone, everywhere to join us in the #fediverse.	Universal
hcommons.social	121	1.61%	2110	2363	hcommons.social is a microblogging network supporting scholars and practitioners across the humanities and around the world.	Universal
aus.social	101	1.34%	1156	1255	Welcome to thundertoot! A Mastodon Instance for the People.	Universal


Table A3: Centrality metrics for Mastodon servers with more than 100 scholars, ordered by count.

Field/Topic specific servers
infosec.exchange, journa.host, scholar.social, fediscience.org, sciences.social, zirk.us, hcommons.social, astrodon.social, drosophila.social, fosstodon.org, qoto.org, masto.ai, genomic.social, hci.social, mathstodon.xyz, med-mastodon.com, ecoevo.social, bayes.club, dair-community.social, mapstodon.space, sigmoid.social, universeodon.com, sciencemastodon.com, neuroscience-mastodon.com, datasci.social, meteo.social, mstdn.science, ioc.exchange, wxcloud.social, skew-t.com, eupolicy.social, law.builders, mastodon.lawprofs.org, esq.social, legal.social, publiclaw.space, openbiblio.social, ausglam.space, lingo.lol, archaeo.social, medievalist.masto.host, econtwitter.net, digipres.club, glammr.us, c18.masto.host, religion.masto.host, fedihum.org, bildung.social, book_historia, coastalhistory, fediphilosophy.org, newsie.social, writing.exchange, urbanists.social, mastodon.gamedev.place, people-making.games, photog.social, mastodon.education, educator.social, mastodon.oeru.org


Table A4: List of field-specific servers.



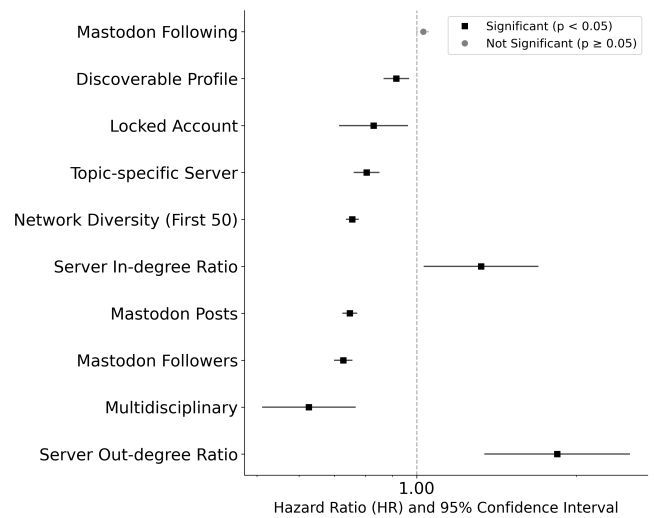
(a) Primary Model: 30-Day Entropy Window (Concordance: 0.69)



(b) 14-Day Entropy Window (Concordance: 0.68)

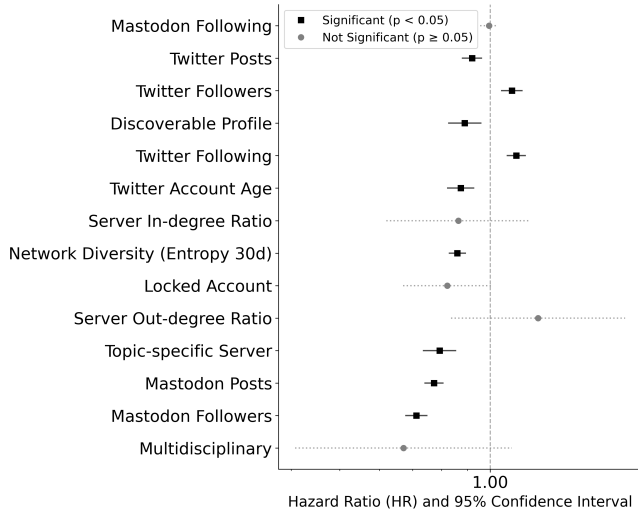


(c) 60-Day Entropy Window (Concordance: 0.69)




(d) First 50 Interactions (Concordance: 0.70)

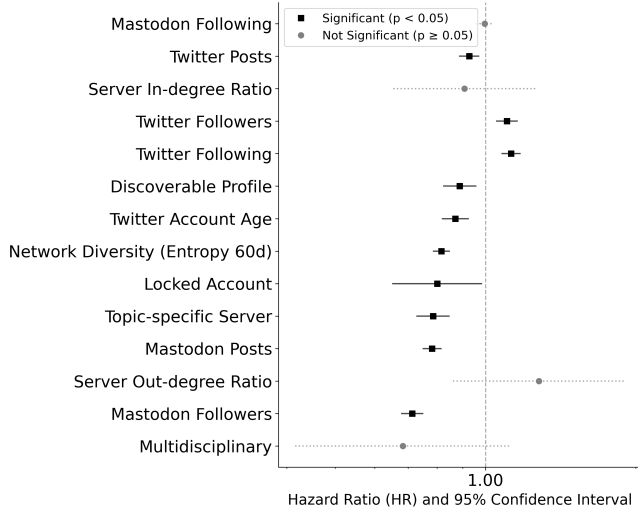
**Figure A6: Mastodon-Centric Model Sensitivity Analysis.** Forest plots displaying hazard ratios for predictors of inactivity across different specifications of the the federation interaction network metric. The protective effect of network diversity and multidisciplinary status remains robust across all time windows. Note that *Locked Account* status becomes a significant risk factor only in the 60-day window ( $p < 0.05$ ), suggesting a cumulative effect.



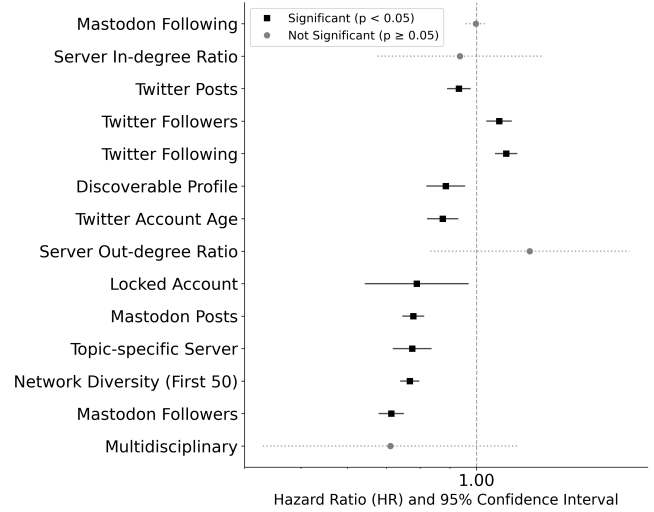
(a) Primary Model: 30-Day Entropy Window (Concordance: 0.69)



(b) 14-Day Entropy Window (Concordance: 0.70)



(c) 60-Day Entropy Window (Concordance: 0.69)



(d) First 50 Interactions (Concordance: 0.70)

**Figure A7: Cross-Platform Model Sensitivity Analysis.** Forest plots for the subset of users with Twitter history. The plots confirm that Federated Interaction Diversity predicts retention independent of Twitter metadata. Unlike the Mastodon-centric model, the Multidisciplinary indicator is not statistically significant in this subset, while Twitter-derived metrics (Followers, Account Age) show consistent divergent associations with retention.