Tin Cheuk Leung* and Kwok Ping Tsang

Migration Talk and Platform Overlap: Network Exposure in #EconTwitter

https://doi.org/10.1515/rne-2026-0039
Received June 5, 2026; accepted November 1, 2026; published online November 17, 2026

Abstract: How does discussion of an alternative platform relate to participation in an existing professional network? We study Mastodon-related discussion in #EconTwitter after Twitter's 2022 ownership change. We measure potential exposure through retweets of content from sources an economist does not directly follow. Economists who accumulated more exposure posted less on Twitter during October–December, controlling for economist and week fixed effects. We find no comparable decline in #EconTwitter-specific posting. The relationship is more negative among economists who were more central in the community's pre-episode interaction network. A later snapshot shows that alternative-platform profiles often coexist with Twitter/X profiles. The evidence links migration discussion to posting behavior and platform overlap, but does not establish causal peer effects or platform exit.

Keywords: platform competition; network exposure; professional communities; multihoming

JEL Classification: D83; L14; L86; Z13

1 Introduction

How does discussion of Mastodon relate to economists' participation in Twitter? Mastodon attracted attention in #EconTwitter around Twitter's ownership change and November 2022 layoffs. We examine whether economists exposed to more of

***Corresponding author: Tin Cheuk Leung**, Wake Forest University, Winston-Salem, USA, E-mail: leungtc@wfu.edu
Kwok Ping Tsang, Virginia Tech, Blacksburg, USA, E-mail: byront@vt.edu.
https://orcid.org/0000-0002-0045-6185

Open Access. © 2026 the author(s), published by De Gruyter. This work is licensed under the Creative Commons Attribution 4.0 International License.

that discussion changed their posting and whether they later maintained profiles on alternative platforms.

Professional communities provide a useful setting for studying platform competition. Users value access to colleagues as well as the features of a platform. An economist may therefore open a Mastodon account while retaining a Twitter account. Network effects and switching costs can explain why competing platforms coexist (Biglaiser et al. 2022; Farrell and Klemperer 2007). Profile overlap alone does not distinguish these frictions.

We combine economists' posting histories with community posts, following relationships, and a later cross-platform collection. The exposure measure follows a simple path. Economist A follows B, who retweets a Mastodon-related post written by C. A does not follow C and is not C. The retweet makes discussion from an otherwise unfollowed source potentially available to A. We observe the path, not whether A sees or reads the post.

The short-run analysis relates weekly posting to exposure accumulated through the previous week. Economist fixed effects account for stable differences across users, and week fixed effects account for common changes. The main analysis covers October–December, when migration discussion became prominent. Accumulating exposure preserves information about both how much discussion becomes available and when it reaches an economist's network.

Economists who accumulated more exposure were less likely to post in a given week and produced fewer posts. We do not find a comparable decline in posts carrying #EconTwitter. The negative general-posting association also appears when exposure is measured using posts from directly followed accounts. Received likes and replies are also lower, which may reflect reduced content production rather than less attention per post.

The posting association is more negative among economists who were more central in the community's pre-episode reply-and-retweet network. That relationship is not explained by follower counts alone. Discussion is concentrated among a small set of intermediaries and original authors, but we find no clear evidence that supplier centrality changes the association. In the later collection, many economists have alternative-platform profiles alongside Twitter/X profiles.

The paper contributes evidence on how discussion of a competing platform relates to participation within an established professional community. The sample drawn from Research Papers in Economics (RePEc) lets us connect identifiable economists to their network positions and subsequent platform profiles. The sample is not representative of all academics or social-media users. Following and retweeting are selected, and the following graph was collected after the episode. We therefore interpret the estimates as associations, not causal effects of exposure.

Section 2 discusses related work and the setting. Section 3 describes the data and exposure measure. Sections 4 and 5 examine posting and network position. Section 6 turns to later platform presence and overlap.

2 Related Literature and Setting

The distinction between learning about an alternative and abandoning an incumbent connects this setting to platform competition. Switching costs and network effects can sustain incumbency even when alternatives are attractive (Biglaiser et al. 2022; Farrell and Klemperer 2007). Multihoming changes competition when users can maintain access to several platforms (Athey et al. 2018).

Studies of platform migration examine both social ties and public discussion. La Cava et al. (2023) relate Twitter-to-Mastodon migration to network structure, commitment, identity, and knowledge exchange. He et al. (2023) track migration to Mastodon, network reconstruction, and activity on both platforms. Jeong et al. (2024) compare migration and user attitudes across Mastodon, Bluesky, and Threads.

Research on academics is especially close to our setting. Quelle et al. (2026) link scholars to Bluesky profiles and study adoption, peer influence, and retention. Wang et al. (2026) examine early academic adopters and the difficulty of sustaining an alternative network. We focus on discussion arriving through sources an economist does not directly follow. Linking those paths to posting within #EconTwitter lets us study participation in the existing community as well as later alternative-platform presence.

The episode combines an ownership change, layoffs, and uncertainty about platform governance. Figure 1 marks October 27, 2022, when the acquisition closed, and November 4, when major layoffs occurred. Mastodon discussion continued beyond the layoffs. The dates help describe the episode, but do not provide an independent source of exposure variation.

3 Data and Potential Network Exposure

3.1 Sample, Sources, and Calendar

The tweet data were collected through the Twitter application programming interface (API), whose academic research access was available from January 2021 through June 2023. For each tweet, we collected its text, date and time, likes and retweets, and author information, including the author's name, username, and Twitter ID. We also collected reply and quote counters, reference types, and account

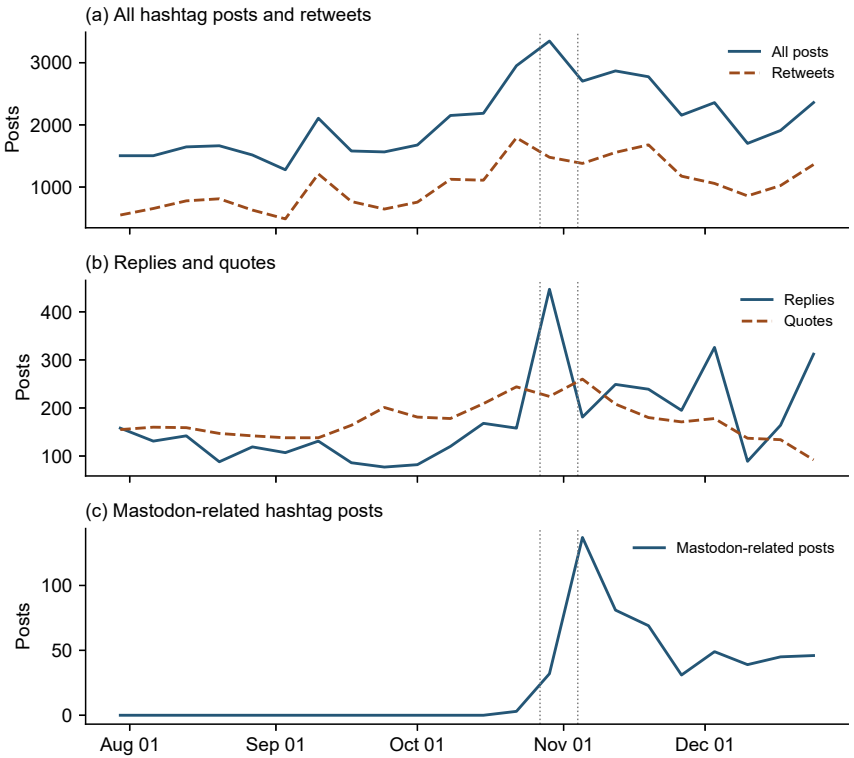

Figure 1: Community posting and Mastodon discussion. Notes: Posts in the #EconTwitter source archive, July 30 – December 31, 2022. Panels separate all posts and retweets, replies and quotes, and dictionary-matched Mastodon discussion. All posts includes reference types, not only original posts. Weeks follow Stata's within-year convention and are shown at their starting dates. Vertical lines mark October 27 and November 4. Counts describe the available archive.

metadata, including follower, following, and total-post counts. The API provided historical posts, so its access period does not define the archive's coverage or our analysis windows.

The community archive contains posts carrying #EconTwitter and the referenced posts needed to identify their authors. Economists' posting histories were collected using Twitter handles from RePEc profiles and a supplemental RePEc-associated Twitter list. We link the accounts to RePEc author profiles using documented account-profile matches. Appendix F reports the linkage coverage and available academic-history information.

The history archive contains 2,321 economists. We retain the 2,239 with observed archive activity after 2021 and keep their zero-post weeks in the balanced

panel. Appendix Table A.1 reports the sample counts. The archive records posts available when collected. It cannot recover deleted or uncollected content.

Following relationships were collected in March 2023. Some relationships may therefore differ from those in place during the discussion episode. The later platform profiles were collected in May 2025, with collection dates varying by platform. We use that collection to describe platform presence, not to reconstruct continuous activity histories.

We use Stata's within-year weekly calendar for both posting and exposure. Week 1 starts on January 1, ordinary weeks contain seven days, and the final week includes the remaining days. In 2022, week 44 is October 29–November 4 and week 52 is December 24–31. All source timestamps are mapped to this common calendar before aggregation.

3.2 Mastodon-Related Network Exposure

To distinguish diffusion beyond an economist's direct source choices, we define an indirect event by the path in Figure 2. An archived post by B must be an actual retweet of C's post and satisfy the Mastodon dictionary described below. The March following graph must show A following B and not following C. We also exclude A = C. Replies and quotes do not qualify as indirect retweet events.

We recover C by joining the referenced-post identifier to its author whenever possible. A documented handle-based fallback is used only when the identifier join does not resolve the source, and the event file retains a resolution-quality flag. Intermediary B and source C remain separate identifiers. Source-count alternatives count C, not the intermediary who retransmitted the message. Multiple distinct B retweets of the same C post can generate distinct recipient-events, so counts measure available paths rather than unique items read.

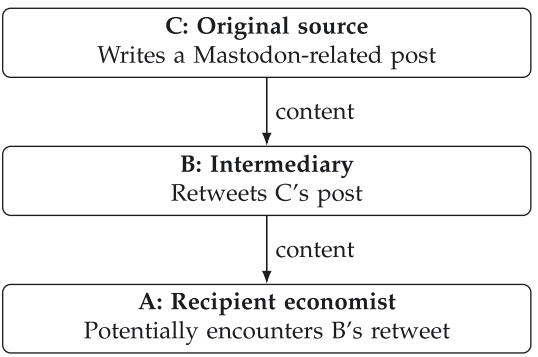

Figure 2: The indirect-exposure path. Notes: Arrows show content transmission, not following direction. A follows B, A does not follow C, and A is not C. B's post must be an actual retweet. Following is observed in March 2023. A potential path does not establish a feed impression or reading.

We identify Mastodon-related content using platform names, common spellings, and a limited set of contextual terms. Examples include `mastodon`, `mastadon`, `econtwitter.net`, `econtoot`, and `fediverse`. Appendix A gives the complete rule and word boundaries. Text validation determined the dictionary before regression comparisons.

Direct-follow exposure is an important alternative, but it measures a different information environment. It counts dictionary-matched posts of any reference type from accounts A follows, using the same archive, recipients, and calendar. An indirect B retweet can also be available as a post by directly followed B. The two counts consequently overlap and should not be added as independent messages. Both measures describe potential network availability, not ranked feed delivery or attention.

3.3 Posting and Platform Outcomes

Weekly total posts include original posts, retweets made, replies written, and quotes posted. We also examine each component separately. Binary outcomes indicate whether an economist posted at all and whether any post carried #EconTwitter. Count outcomes enter as log of one plus the count, retaining zero-post weeks. Appendix H gives the exact hashtag strings, variables, and missing-value rules.

Received engagement is distinct from posts made. We sum collection-snapshot likes, retweets, replies, and quotes attached to authored original, reply, and quote posts created in a week. Plain retweets are excluded because their counters can describe the redistributed content's popularity. These are neither reactions arriving during that week nor engagement per post. Weeks without authored content contribute zeros, and the available archived metrics are checked for missingness before aggregation.

For the later collection, we observe whether a profile was located and its recorded posting counter. A positive counter indicates posting in the account's recorded history, not necessarily recent activity. We call a located Twitter/X profile together with an alternative-platform profile *profile overlap*. That is the observable form of multihoming in these data. An unlocated profile does not establish exit.

4 Migration Discussion and Short-Run Posting

4.1 Accumulated Exposure and Weekly Posting

We ask whether economists post less as Mastodon-related discussion accumulates in their networks. Let $C_{i,t-1}$ count indirect exposure events from January 1, 2022 through the week before t. We estimate

$$Y_{it} = \alpha_i + \lambda_t + \beta \log(1 + C_{i,t-1}) + \varepsilon_{it}. \quad (1)$$

Here α_i denotes economist fixed effects and λ_t denotes week fixed effects. The main window is October 1–December 31, 2022. All analysis economists remain in the balanced sample, and standard errors are clustered by economist. Binary outcomes stay in levels and count outcomes use $\log(1 + Y)$.

The measure accumulates exposure as new Mastodon-related content becomes available through the network. It therefore preserves variation in both the amount and timing of discussion. The October–December window covers the ownership change and the ensuing discussion. Using exposure through the prior week separates it from current-week posting.

Table 1 shows that economists who accumulated more exposure posted less. The coefficient for any post is −0.0154, with a standard error of 0.0075. A one-unit increase in log of one plus exposure is therefore associated with a 1.54 percentage-point lower weekly posting probability. The coefficient for logged total posts is −0.0588, with a standard error of 0.0203. In contrast, the estimates for #EconTwitter participation and volume are small relative to their standard errors.

The association is strongest over October–December, when migration discussion became prominent. Appendix Tables A.2–A.3 extend the outcome window backward to September, August, and late July without changing exposure accumulation. Both general-posting coefficients remain negative and become steadily smaller as earlier weeks enter the comparison. The longer-window estimates are also less precise.

Table 1: Network exposure and weekly Twitter activity.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.0154**	−0.0057	−0.0588***	−0.0022
	(0.0075)	(0.0044)	(0.0203)	(0.0043)
Constant	0.6783***	0.0580***	1.4874***	0.0481***
	(0.0033)	(0.0020)	(0.0090)	(0.0019)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. ** $p < 0.05$, *** $p < 0.01$.

4.2 Interpreting the Posting Association

The estimates describe changes in posting associated with changes in potential exposure. They do not establish that reading Mastodon discussion causes an economist to withdraw from Twitter. Economists choose whom to follow, and intermediaries choose which posts to retweet. The March following graph may also reflect relationships formed after the episode. Fixed effects remove stable individual differences and common weekly changes, but not every time-varying influence on exposure and posting.

Potential exposure can rise even when an economist is inactive. Appendix Table A.4 stops exposure accumulation after each economist's last observed post in the window. The posting-probability estimate becomes weaker, while the logged-volume estimate remains negative. Because that restriction uses subsequent posting, it is a diagnostic rather than an alternative main specification.

The evidence is consistent with reduced general posting during migration discussion, but neither the timing of the episode nor a last observed post establishes departure. Shared source messages also create dependence across economists that person-clustered standard errors may not fully capture.

4.3 Direct Availability and Outcome Robustness

The negative general-posting association is not confined to indirect exposure. Appendix C uses Mastodon-related posts from directly followed accounts instead, holding the outcome observations and model fixed. The general-posting coefficients are again negative. When both exposure measures enter together, their separate coefficients are imprecise. The comparison shows that the two measures capture related information environments, not independent causal channels.

Raw posting counts give the same negative direction but less precise estimates, as shown in Appendix D. Logs reduce the influence of unusually high-volume weeks without changing the sample. Since raw and logged coefficients have different units, we compare their signs and precision rather than their numerical magnitudes.

Logged likes and replies received by authored posts are also negatively associated with exposure. Received retweets and quotes are imprecise, as are the raw engagement estimates. Lower engagement can accompany lower production without implying less attention per post.

5 Content Supply and Network Position

Who supplies the discussion, and which economists respond more strongly? Table 2 shows that exposure is concentrated among a small set of intermediaries B and

Table 2: Concentration of indirect exposure.

Role	Accounts	Recipient-events	Top 1 (%)	Top 5 (%)	HHI
B	79	5,859	14.6	43.9	0.058
C	35	5,859	20.4	67.1	0.109

Indirect events, January 2022–February 2023 archive. B is the followed intermediary and C the original author. Shares weight recipient-events, not unique messages. HHI is the Herfindahl-Hirschman index, the sum of squared account shares.

original authors C. Many economists can thus receive content originating in the same message. A large count of recipient-events need not represent a large number of independent messages.

To distinguish community position from audience size, we construct a directed network of replies and retweets observed before the episode, from January 1 through October 26, 2022. Edges point from the interacting account to the referenced author. Repeated interactions form one binary edge, without frequency weighting. PageRank uses damping 0.85. We standardize log PageRank over mapped sample economists, retain mapped isolates, and distinguish unmapped observations from zero centrality. Appendix E reports in-degree and March-follow-network alternatives.

The relationship between exposure and posting is more negative among economists who were more central in the pre-episode interaction network. Table 3 interacts exposure with PageRank and log follower counts. The PageRank interaction remains informative after accounting for audience size, including for #EconTwitter participation. Economist effects absorb the level of each time-invariant characteristic. The interactions describe differences across economists, not the effect of changing someone's network position.

We find no clear evidence that the relationship depends on supplier centrality. The supplier estimates are imprecise, and coverage is limited, especially for original sources C. We keep exposure from unmapped suppliers separate rather than assigning them low centrality. Appendix E gives the construction and results.

Appendix Table A.15 also reports how predicted later platform overlap changes when each economist's observed exposure is increased by 10 percent. This is a descriptive comparison based on the estimated model. It is not the effect of supplying more tweets, which could reach different recipients.

Bibliography completeness and reliable pre-2022 dates vary across economists, limiting comparisons by academic standing. Appendix F reports coverage, source quality, and descriptive statistics for available academic histories.

Table 3: Recipient position beyond follower size.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.1866***	−0.0499***	−0.5202***	−0.0544***
	(0.0340)	(0.0192)	(0.1196)	(0.0181)
Exposure x log followers	0.0231***	0.0064**	0.0647***	0.0074***
	(0.0044)	(0.0025)	(0.0162)	(0.0024)
Exposure x pre-episode PageRank	−0.0272***	−0.0221***	−0.1641***	−0.0205***
	(0.0064)	(0.0044)	(0.0229)	(0.0043)
Constant	0.6806***	0.0583***	1.4927***	0.0484***
	(0.0034)	(0.0020)	(0.0089)	(0.0019)
Observations	29,029	29,029	29,029	29,029
Economists	2,233	2,233	2,233	2,233
R^2	0.544	0.299	0.792	0.398

October 1–December 31, 2022. Balanced mapped sample where position is observed, with sample sizes shown. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Position is standardized log PageRank on the January 1–October 26 reply/retweet graph unless labeled otherwise. Follower interactions use uncentered log(1+followers), so the exposure main term in those models is evaluated at zero followers. March follower/follow-graph information is post-episode. See text for coverage and supplier weighting. ** $p < 0.05$, *** $p < 0.01$.

6 Later Platform Profiles and Overlap

The May 2025 collection shows whether alternative-platform profiles coexist with Twitter/X profiles. Figures 3 and 4 report platform combinations at collection for lower- and higher-exposure economists. The bands describe a snapshot, not migration paths.

The snapshot shows substantial alternative-platform presence with continued Twitter/X profile overlap. Appendix Table A.20 gives the exact state counts behind the figures, so small categories remain visible without relying on band widths.

Table 4 relates the platform combinations to exposure over the source archive. Greater exposure is associated with both alternative-platform presence and overlap with Twitter/X. Requiring positive recorded counters gives the same qualitative pattern. The regressions control for March follower and following counts, verification, and the recorded post counter. They describe cross-sectional associations rather than causal migration effects.

Profile-creation dates record when an account was registered. Appendix G reports these dates for located Bluesky and Mastodon profiles and the incomplete Mastodon last-post dates. Registration need not coincide with first posting

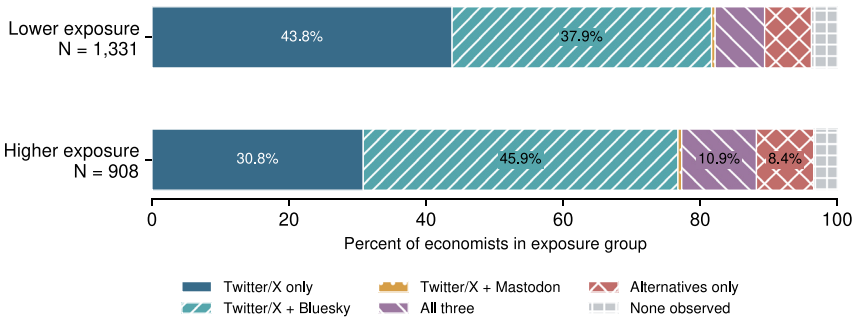

Figure 3: Platform-profile composition by network exposure. Notes: May 2025 collection. Higher exposure means at least three archive recipient-events, strictly above the full-sample median of two. Every economist belongs to one mutually exclusive endpoint state. The widths describe final composition, not migration trajectories. Twitter/X plus alternatives denotes located-profile overlap.

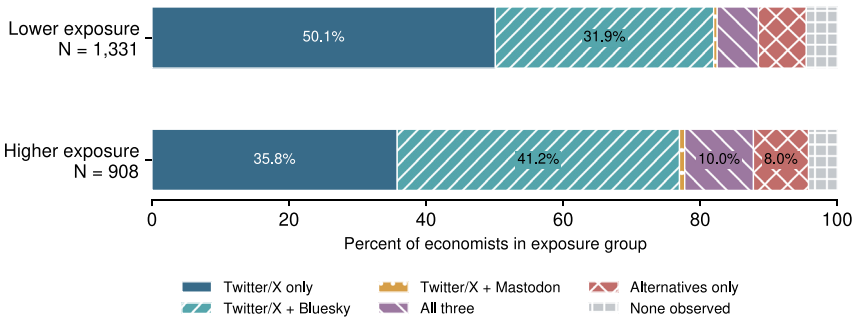

Figure 4: Composition requiring positive recorded counters. Notes: The same sample and exposure groups as Figure 3. A platform counts only when a profile is located and its recorded post counter is positive. Counters are account-history snapshots, not activity in a common recent period. An absent or zero-counter profile does not establish exit.

or primary-platform choice. A last-post date does not establish exit or provide a comparable recent-activity measure across the three platforms.

The snapshot documents the coexistence of alternative-platform and Twitter/X profiles. Opening or maintaining an alternative profile can be meaningful diversification even when a Twitter/X profile persists. Establishing coordinated exit or sustained simultaneous use would require comparable activity histories across platforms, which this collection does not provide.

Table 4: Network exposure and later platform overlap.

	Alt. profile	Alt. positive	Overlap profile	Overlap positive
Log(1 + exposure)	0.1128***	0.1209***	0.0963***	0.1075***
	(0.0144)	(0.0145)	(0.0146)	(0.0144)
Verified	0.1240**	0.1248**	0.1551***	0.1568***
	(0.0503)	(0.0524)	(0.0516)	(0.0533)
Followers (000s)	0.0001**	0.0001**	0.0001**	0.0002**
	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Following (000s)	0.0027	0.0045	0.0028	0.0045
	(0.0036)	(0.0036)	(0.0036)	(0.0036)
March post counter (000s)	0.0005	0.0011*	0.0009	0.0015**
	(0.0006)	(0.0007)	(0.0007)	(0.0007)
Constant	0.4492***	0.3682***	0.3870***	0.3035***
	(0.0185)	(0.0181)	(0.0183)	(0.0175)
Observations	2,239	2,239	2,239	2,239
Economists	2,239	2,239	2,239	2,239
R^2	0.034	0.042	0.029	0.041

May 2025 collection, one observation per economist. Linear probability models with heteroskedasticity-robust HC1 standard errors. Exposure covers the source archive through February 2023. Controls are March 2023 snapshots. No person or week fixed effects. Positive means a positive recorded account-history counter, not common-period recent activity. Overlap requires Twitter/X and at least one alternative. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7 Conclusions

Our findings connect discussion of an alternative platform with participation in an established professional network and later profile overlap. In our preferred October–December specification, economists who accumulated more Mastodon-related network exposure were less likely to post and produced fewer posts on Twitter. We find no comparable decline in #EconTwitter-specific posting. The general-posting relationship is more negative among economists who were more central in the pre-episode interaction network. Later alternative-platform profiles frequently coexist with Twitter/X profiles.

The negative cumulative-exposure association is strongest over October–December, when Mastodon discussion became prominent. Extending the outcome window backward to September, August, and late July reduces the magnitude of both general-posting coefficients and weakens the statistical evidence. The point estimates remain negative, but the longer-window estimates are less precise (Appendix Tables A.2–A.3).

The short-run interpretation also depends on how exposure is summarized. Alternative fixed-exposure specifications do not show a statistically significant decline in general posting. Those models assign each economist one episode-level exposure measure and compare average pre/post changes across economists. The preferred cumulative specification instead retains variation in when exposure accumulates within economists. The difference does not establish an absence of any relationship, but it limits the generality of the negative finding. Taken together, the comparisons support a qualified association between accumulated discussion and reduced general posting during the main episode, not a specification-invariant or causal effect on Twitter withdrawal.

For platform competition, the findings highlight adoption of alternative profiles and changes in content production as margins worth examining alongside complete switching. Incumbent exit alone may miss adjustments that leave users connected to several services. Policies or platform designs that make identities portable or help users find existing professional contacts across platforms could lower the cost of maintaining those connections. Competition might then operate through multihoming as well as complete switching, although the paper does not estimate the effects of such policies. The centrality results also suggest that community averages can conceal differences associated with accumulated network position.

Several data limitations further qualify the interpretation. Exposure records content potentially available through network paths, not feed impressions, reading, or attention. Economists choose whom to follow, and intermediaries choose what to retweet. The following graph was observed in March 2023, after the central episode, so it need not represent the network in place when discussion occurred. The May 2025 collection records profile presence and account-history posting counters. It does not provide continuous activity histories or identify primary-platform choice or verified exit. Finally, the RePEc-linked economist sample need not represent all academics, general Twitter users, or other online communities.

Future research could combine feed-impression data with changing following networks to distinguish potential availability from consumption and establish when network ties formed. Continuous posting histories across Twitter/X, Mastodon, and Bluesky would help separate profile adoption, active multihoming, substitution in posting, and exit. Combining these records with plausibly exogenous variation in exposure or platform access could test whether migration discussion changes incumbent participation. Such work would move beyond a comparison of staying and leaving to examine partial shifts across services. Understanding competition in professional networks requires following both information flows and participation across platforms.

Acknowledgments: We thank Rubaiyat Alam and seminar participants at CREED and SEA for helpful comments. We thank Munkhchimeg Sukhee and Eric Uehling for research assistance. The usual disclaimer applies.
Research ethics: Not applicable. This study uses publicly available online information and did not require institutional review board approval.
Informed consent: Not applicable. This study uses publicly available online information.
Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission.
Use of Large Language Models, AI and Machine Learning Tools: The authors used Codex (OpenAI) to assist with manuscript editing and bibliographic-record review. Codex also assisted with developing and checking programs for data collection and analysis, including programs used to obtain data from X and other platforms.
Conflict of interest: The authors state no conflict of interest.
Research funding: This research received financial support from the Institute for Humane Studies.
Data availability: The replication package is available from the corresponding author upon reasonable request.

Appendix

A Exposure Construction and Sample Coverage

An indirect exposure event requires a retweet path A–B–C, no direct follow relationship from A to C, and A distinct from C. We identify the source through the referenced-post identifier before using the documented handle fallback. The event

Table A.1: Sample and indirect-exposure coverage.

Sample or exposure quantity	Count
Economists in history archive	2,321
Economists in balanced analysis	2,239
Preferred cumulative observations	29,107
Archive recipient-events	5,859
Recipients with archive exposure	1,713
Nonzero recipient-weeks	4,123

Analysis economists have observed archive activity after 2021. Indirect events exclude A = C. Archive event dates are 2022-01-08 to 2023-02-28. The main outcome panel is October 1–December 31, 2022. Nonzero recipient-weeks use the aligned Stata calendar.

data retain both identifiers and the source-resolution method. Distinct source counts use C rather than B.

The keyword rule includes `mastodon`, `mastadon`, `econtwitter.net`, `econtoot`, `fediverse`, and literal `m*stodon`. Bounded terms are `elephant app`, `elephant site`, `toot site`, `federated place`, `masto`, `trunk`, and `tusk`. Bounds exclude adjacent letters, digits, and underscores. Text review supports these terms, but euphemisms can remain ambiguous. Keyword matching does not establish that every matched post advocates migration.

B Outcome Windows and Inactivity

Tables A.2 and A.3 change only the starting week of the outcome panel. The first complete Stata weeks beginning in October, September, and August start on October 1, September 3, and August 6. The late-July comparison starts on July 30. All windows end on December 31. Exposure always accumulates from January 1 through the prior week, including exposure before the outcome window.

The replication materials also retain the focused comparison from July 30–October 28 and November 5–December 2, its extension through December, and a sample active before the episode. The focused general-posting estimates are positive and imprecise. Thus the longer-window progression does not imply a negative association under every timing comparison. The inactivity diagnostic below holds the October–December outcome window fixed and changes only accumulation after last observed posting. Additional logged posting outcomes appear in Table A.5.

Table A.2: Weekly posting probability across outcome windows.

	Oct–Dec	Sep–Dec	Aug–Dec	Late Jul–Dec
Cumulative exposure	−0.0154**	−0.0136*	−0.0098	−0.0078
	(0.0075)	(0.0075)	(0.0078)	(0.0078)
Constant	0.6783***	0.6761***	0.6659***	0.6632***
	(0.0033)	(0.0027)	(0.0024)	(0.0023)
Observations	29,107	38,063	47,019	49,258
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.528	0.516	0.513

Dependent variable: Any post. All columns use the full balanced population, economist and Stata-week fixed effects, and economist-clustered standard errors. Starting dates in column order are October 1, September 3, August 6, and July 30, 2022. Every window ends on December 31. The August and September starts are the first complete Stata weeks beginning in those months. Exposure always accumulates from January 1 through the prior week. Only the outcome-panel start changes. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Logged total posts across outcome windows.

	Oct–Dec	Sep–Dec	Aug–Dec	Late Jul–Dec
Cumulative exposure	−0.0588***	−0.0422**	−0.0301	−0.0251
	(0.0203)	(0.0215)	(0.0232)	(0.0235)
Constant	1.4874***	1.4705***	1.4308***	1.4203***
	(0.0090)	(0.0078)	(0.0072)	(0.0070)
Observations	29,107	38,063	47,019	49,258
Economists	2,239	2,239	2,239	2,239
R^2	0.790	0.767	0.741	0.735

Dependent variable: Log posts. All columns use the full balanced population, economist and Stata-week fixed effects, and economist-clustered standard errors. Starting dates in column order are October 1, September 3, August 6, and July 30, 2022. Every window ends on December 31. The August and September starts are the first complete Stata weeks beginning in those months. Exposure always accumulates from January 1 through the prior week. Only the outcome-panel start changes.
** $p < 0.05$, *** $p < 0.01$.

Table A.4: Outcome-conditioned stock-freezing diagnostic.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Frozen cumulative exposure	0.0057	−0.0059	−0.0502**	−0.0022
	(0.0075)	(0.0046)	(0.0207)	(0.0045)
Constant	0.6690***	0.0579***	1.4822***	0.0480***
	(0.0031)	(0.0019)	(0.0087)	(0.0019)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Exposure is frozen after the last observed post in this outcome window. This outcome-conditioned diagnostic is not the preferred estimator. ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Additional logged posting outcomes.

	Log retweets made	Log replies made	Log quotes made
Cumulative exposure	−0.0434***	−0.0485***	−0.0228**
	(0.0165)	(0.0159)	(0.0105)
Constant	0.8900***	0.7809***	0.4061***
	(0.0073)	(0.0071)	(0.0047)
Observations	29,107	29,107	29,107
Economists	2,239	2,239	2,239
R^2	0.764	0.743	0.684

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects.
** $p < 0.05$, *** $p < 0.01$.

C Direct-Follow Comparison

Direct and indirect stocks start on January 1, 2022 and end at the prior week. Direct content may be an original, reply, quote, or retweet post from a followed account. The indirect measure requires a validated retweet path and excludes direct following of C. Their overlap at B's post means adding counts would double-count available content. The joint model is descriptive, not a mediation decomposition. The direct-only and joint specifications appear in Tables A.6 and A.7.

Table A.6: Direct-follow cumulative exposure.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Direct cumulative exposure	−0.0093**	−0.0007	−0.0327**	0.0016
	(0.0047)	(0.0027)	(0.0135)	(0.0025)
Constant	0.6821***	0.0562***	1.4987***	0.0452***
	(0.0054)	(0.0031)	(0.0154)	(0.0029)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.297	0.790	0.397

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Direct exposure is log of one plus matched posts by followed accounts, accumulated from January 1, 2022 through the prior week. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Direct exposure counts dictionary-matched posts of every reference type by followed accounts. It measures potential availability, not reading. ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Joint indirect and direct cumulative stocks.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Indirect cumulative exposure	−0.0101	−0.0073	−0.0422*	−0.0053
	(0.0086)	(0.0055)	(0.0239)	(0.0051)
Direct cumulative exposure	−0.0059	0.0017	−0.0186	0.0034
	(0.0054)	(0.0034)	(0.0159)	(0.0029)
Constant	0.6827***	0.0567***	1.5014***	0.0456***
	(0.0055)	(0.0031)	(0.0155)	(0.0029)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Both exposure counts enter as log of one plus the stock accumulated from January 1, 2022 through the prior week. Indirect paths exclude self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Direct exposure counts dictionary-matched posts of every reference type by followed accounts. It measures potential availability, not reading. * $p < 0.10$, *** $p < 0.01$.

D Count Transformations and Received Engagement

The raw and logged specifications use identical economist-week observations. Counts are right-skewed, which motivates logged main outcomes. Received counters refer to authored posts created in the week, measured when collected. No authored content contributes zero, not a missing counter. These models do not condition on realized positive posting. The raw posting and received-engagement specifications appear in Tables A.8–A.10.

Table A.8: Raw posting-count robustness.

	All posts	Retweets made	Replies made	Quotes made	#EconTwitter
Cumulative exposure	−0.7552	−0.1738	−0.4145	−0.0829	−0.0001
	(0.5623)	(0.2826)	(0.2813)	(0.0671)	(0.0096)
Constant	13.1188***	4.5892***	5.1775***	1.2666***	0.0833***
	(0.2507)	(0.1260)	(0.1254)	(0.0299)	(0.0043)
Observations	29,107	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239	2,239
R^2	0.674	0.710	0.594	0.605	0.469

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are raw levels. Associations are not causal effects. *** $p < 0.01$.

Table A.9: Logged counters received by authored posts.

	Likes received	Retweets received	Replies received	Quotes received
Cumulative exposure	−0.0657**	−0.0340	−0.0438**	−0.0194
	(0.0303)	(0.0216)	(0.0183)	(0.0132)
Constant	1.8415***	0.9341***	0.8914***	0.3960***
	(0.0135)	(0.0096)	(0.0081)	(0.0059)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.753	0.725	0.762	0.681

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Metrics are collection-snapshot counters on original, reply and quote posts created in that week. Plain retweets are excluded. These are not reactions arriving during the week or engagement per post. No authored posts implies zero. ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Raw counters received by authored posts.

	Likes received	Retweets received	Replies received	Quotes received
Cumulative exposure	−60.2287	−10.2734	−4.5862*	−0.5553
	(40.9352)	(7.4311)	(2.4240)	(0.6590)
Constant	278.9877***	41.4139***	18.4728***	4.4255***
	(18.2508)	(3.3131)	(1.0807)	(0.2938)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.609	0.502	0.538	0.510

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are raw levels. Associations are not causal effects. Metrics are collection-snapshot counters on original, reply and quote posts created in that week. Plain retweets are excluded. These are not reactions arriving during the week or engagement per post. No authored posts implies zero. * $p < 0.10$, *** $p < 0.01$.

E Network-Position Alternatives

In-degree checks whether the PageRank result depends on recursive weighting. March follow-network position is measured after the episode and remains descriptive. Supplier weighting keeps unmatched sources distinct. The interaction-network sample excludes unmapped users but retains isolates. The network-position and follower-size specifications appear in Tables A.11–A.14.

Table A.11: Recipient pre-episode in-degree.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.0119	−0.0011	−0.0259	0.0016
	(0.0079)	(0.0042)	(0.0206)	(0.0041)
Exposure x pre-episode in-degree	−0.0097*	−0.0164***	−0.1141***	−0.0138***
	(0.0051)	(0.0035)	(0.0162)	(0.0037)
Constant	0.6794***	0.0579***	1.4889***	0.0480***
	(0.0034)	(0.0020)	(0.0090)	(0.0019)
Observations	29,029	29,029	29,029	29,029
Economists	2,233	2,233	2,233	2,233
R^2	0.543	0.299	0.791	0.398

October 1–December 31, 2022. Balanced mapped sample where position is observed, with sample sizes shown. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Position is standardized log PageRank on the January 1–October 26 reply/retweet graph unless labeled otherwise. Follower interactions use uncentered log(1+followers), so the exposure main term in those models is evaluated at zero followers. March follower/follow-graph information is post-episode. See text for coverage and supplier weighting. * $p < 0.10$, *** $p < 0.01$.

Table A.12: Recipient March following-network PageRank.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.0198**	−0.0020	−0.0568***	0.0004
	(0.0078)	(0.0044)	(0.0201)	(0.0042)
Exposure x March PageRank	0.0117**	−0.0097**	−0.0053	−0.0070*
	(0.0059)	(0.0039)	(0.0204)	(0.0042)
Constant	0.6790***	0.0575***	1.4871***	0.0477***
	(0.0034)	(0.0019)	(0.0089)	(0.0019)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced mapped sample where position is observed, with sample sizes shown. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Position is standardized log PageRank on the January 1–October 26 reply/retweet graph unless labeled otherwise. Follower interactions use uncentered log(1+followers), so the exposure main term in those models is evaluated at zero followers. March follower/follow-graph information is post-episode. See text for coverage and supplier weighting. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Follower-size interaction.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.1166***	0.0076	−0.0940	−0.0010
	(0.0279)	(0.0159)	(0.0906)	(0.0164)
Exposure x log followers	0.0131***	−0.0017	0.0045	−0.0002
	(0.0033)	(0.0020)	(0.0120)	(0.0022)
Constant	0.6790***	0.0579***	1.4877***	0.0481***
	(0.0034)	(0.0020)	(0.0089)	(0.0019)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced mapped sample where position is observed, with sample sizes shown. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Position is standardized log PageRank on the January 1–October 26 reply/retweet graph unless labeled otherwise. Follower interactions use uncentered log(1+followers), so the exposure main term in those models is evaluated at zero followers. March follower/follow-graph information is post-episode. See text for coverage and supplier weighting. *** $p < 0.01$.

Table A.14: Intermediary-position-weighted exposure.

	Any post	Any #EconTwitter	Log posts	Log #EconTwitter
Cumulative exposure	−0.0328**	−0.0029	−0.0809**	−0.0024
	(0.0155)	(0.0096)	(0.0396)	(0.0085)
Unmatched-B exposure	0.0124	−0.0050	0.0163	0.0005
	(0.0162)	(0.0109)	(0.0428)	(0.0103)
B-position-weighted exposure	0.0098	−0.0005	0.0122	−4.6e-05
	(0.0065)	(0.0043)	(0.0180)	(0.0038)
Constant	0.6784***	0.0581***	1.4875***	0.0481***
	(0.0034)	(0.0020)	(0.0092)	(0.0020)
Observations	29,107	29,107	29,107	29,107
Economists	2,239	2,239	2,239	2,239
R^2	0.545	0.298	0.790	0.397

October 1–December 31, 2022. Balanced full sample. Economist and Stata-week fixed effects, economist-clustered standard errors in parentheses. Exposure is log of one plus indirect events accumulated from January 1, 2022 through the prior week, excluding self-source events. Binary outcomes are levels, counts are log(1+count). Associations are not causal effects. Position is standardized log PageRank on the January 1–October 26 reply/retweet graph unless labeled otherwise. Follower interactions use uncentered log(1+followers), so the exposure main term in those models is evaluated at zero followers. March follower/follow-graph information is post-episode. See text for coverage and supplier weighting. ** $p < 0.05$, *** $p < 0.01$.

Table A.15: Descriptive exposure-increase contrasts for later overlap.

Model	Change (pp)	SE (pp)	p-value	N
Baseline	0.556	0.074	0.000	2,239
Recipient position	0.230	0.079	0.003	2,233
Intermediary position	0.577	0.076	0.000	2,239

Average fitted probability change for Twitter/X plus alternative positive counters when each observed event count is multiplied by 1.10. Controls and observed position held fixed. HC1 delta-method standard errors. Change and standard error (SE) are in percentage points (pp). N is the number of economists. This is a descriptive fitted contrast, not an intervention on tweets supplied or a causal effect.

For supplier B, let $K_{i,t-1}$ be cumulative events with mapped supplier position, $U_{i,t-1}$ cumulative unmatched events, and $\bar{Z}^{B}_{i,t-1}$ the event-weighted mean standardized supplier position. The supplier model adds $\log(1 + U_{i,t-1})$ and $\log(1 + K_{i,t-1})\bar{Z}^{B}_{i,t-1}$ to the cumulative-exposure regressor. The latter contribution is zero when no mapped event exists. That convention does not assign zero centrality to an unmapped account. Source-C alternatives use the same rule and have weaker coverage. Long-run fitted contrasts apply the corresponding full-archive counts.

F Academic-Data Coverage and Source Quality

We link Twitter accounts to RePEc author profiles using direct handle evidence and other high-confidence identity matches. Ambiguous cases are not treated as established links. Table A.16 distinguishes identifiers from accounts because one RePEc author is linked to two Twitter accounts.

Table A.16: RePEc linkage and historical-variable coverage.

Measure	IDs	Accounts	Account %
Accepted RePEc linkage	2,172	2,173	97.1
Profile pages retrieved	2,168	2,169	96.9
Profile research references retrieved	2,153	2,154	96.2
Structured person record retrieved	2,168	2,169	96.9
Usable journal count	1,593	1,594	71.2
Usable unique-project count	979	979	43.7
Usable earliest-work year	1,041	1,041	46.5
Usable earliest-journal year	1,638	1,639	73.2

Denominator is 2,239 economist accounts. IDs count distinct accepted RePEc authors. One RePEc author is linked to two Twitter accounts. Usable means the measure-specific date, ownership and completeness checks pass.

Table A.17: Bibliography source-review evidence.

Check	Sampled	Assessed	Disagreed	%
Work ownership	47	45	2	4.4
ReDIF type mapping	47	45	0	0.0
Inspected duplicate examples	47	45	0	0.0
Observed earliest year	47	45	2	4.4
Sampled groups of publication versions	25	25	0	0.0

The review of bibliographic records combines 42 stratified draws with five additional cases selected for targeted review. The person-level checks use the 45 assessable histories and cover earliest and recent references and selected duplicate examples, not every work. A separate random sample covers 25 groups of records identified as possible versions of the same research work. Records were reviewed by an AI assistant, without independent human verification. The disagreement rates are not estimates of population error rates.

Table A.18: Completeness by observed history size.

Observed article records	IDs	Journal %	Project %	First work %
Zero observed	144	74.3	64.6	64.6
1-9 observed	816	90.3	65.0	66.2
10-24 observed	631	75.1	39.6	44.7
25-49 observed	367	56.4	24.0	28.3
50 or more observed	214	31.8	8.4	10.3

Rows group linked IDs by observed journal records through 2021, not true lifetime productivity. Percentages show measure-specific usability. Missing histories are not zero publications.

Table A.19: Descriptive statistics for available academic histories.

Measure	N	Mean	Median	p90	Missing %
Journal articles	1,593	14.21	9.00	33.00	26.66
Unique projects	979	19.75	13.00	45.00	54.93
Working-paper versions	979	16.79	9.00	42.00	54.93
First work year	1,041	2,008.05	2,010.00	2,017.00	52.07
First journal year	1,638	2,007.54	2,009.00	2,018.00	24.59
Years since first work	1,041	12.95	11.00	24.00	52.07
Years since first journal	1,638	13.46	12.00	26.00	24.59

Unit is distinct accepted RePEc ID. Counts cover records through 2021 that meet the authorship, date, and version-reconciliation criteria in Appendix F, with complete RePEc enumeration for the relevant measure. Missing percentages use all 2,172 accepted IDs. Years since first work is 2021 minus the earliest indexed year meeting the date and authorship criteria, not PhD age. Journal or serial status does not certify prestige or peer review.

Linking an account to RePEc does not ensure a complete publication history. Bibliography records can have unresolved ownership, missing years, and multiple versions of the same work. We classify dated journal records separately from broader unique projects and use time since the earliest dated work we observe as a proxy for career length. Counts require source-backed authorship, dates through 2021, reconciled versions, and complete RePEc enumeration for the measure. Earliest-year measures require that no unresolved record could precede the observed first year. Histories held for unresolved ownership are excluded. Missing information on institutional prestige, PhD pedigree, and historical citations remains missing.

We use the academic-history data descriptively. Source review identifies ownership and date errors, and usable coverage falls with larger observed publication histories. Incomplete and selectively available academic histories limit comparisons of platform use by professional standing. Source-review results, completeness by observed history size, and descriptive statistics appear in Tables A.17–A.19.

G Endpoint Counts and Profile Timing

The exact endpoint counts supplement Figures 3–4. Higher and lower groups use full-archive exposure. Profile-creation dates are observed for located alternative

Table A.20: Exact platform-state counts.

Definition/state code	Lower exposure	Higher exposure
Profile 0	50	31
Profile 1	3	0
Profile 2	64	57
Profile 3	24	19
Profile 4	583	280
Profile 5	7	5
Profile 6	504	417
Profile 7	96	99
Positive 0	60	38
Positive 1	3	0
Positive 2	64	56
Positive 3	26	17
Positive 4	667	325
Positive 5	7	7
Positive 6	424	374
Positive 7	80	91

State code = 4(Twitter/X) + 2(Bluesky) + 1(Mastodon). High is above the full-sample median of two archive events, hence at least three. Lower $N = 1{,}331$, higher $N = 908$. These are endpoint compositions, not transitions. Positive denotes a positive recorded counter.

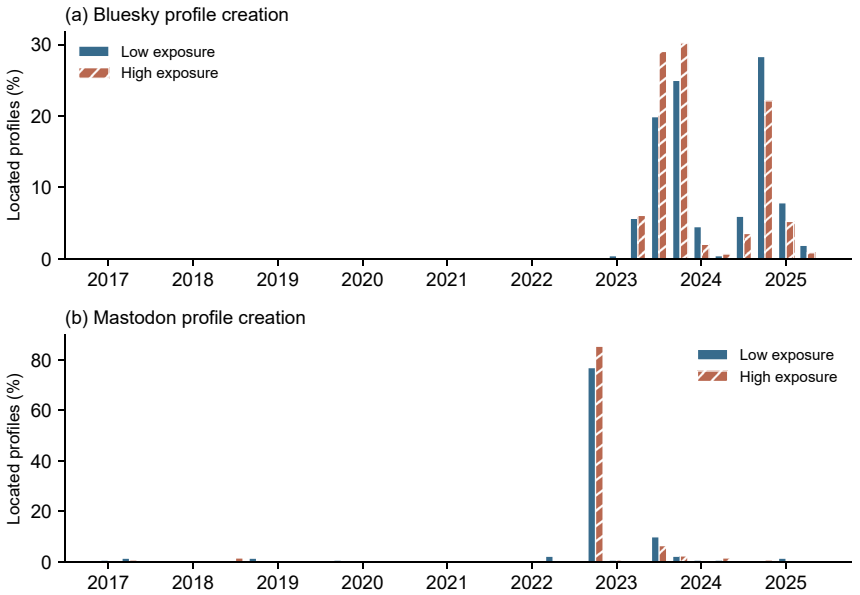

Figure A.1: Alternative-platform profile-creation timing. Notes: Dates in the May 2025 workbook for located profiles. Percentages use located profiles within each exposure group. Creation is not first posting, primary-platform choice, or migration timing. The two platforms have different launch histories.

Table A.21: Mastodon last-post dates, including missing dates.

Last-post period	Overall	Lower	Higher
Before 2022 Q4	1	1	0
2022 Q4	53	23	30
2023 Q1	57	28	29
2023 Q2-Q4	51	27	24
2024	27	12	15
2025	40	25	15
Missing date	24	14	10

Located Mastodon profiles in May 2025. Exposure groups use the archive median split. Missing dates are retained. Last observed posting is not exit or common-period cross-platform activity.

profiles (Figure A.1). Some Mastodon profiles have missing last-post dates, while one dated record precedes the episode (Table A.21). Zero-counter accounts can have recorded dates, and positive-counter accounts can lack them. We retain these inconsistencies and do not infer a common activity window.

H Variable Definitions and Reproducibility

Table A.22 collects the variable definitions used throughout the paper.

Table A.22: Variable definitions.

Variable	Definition and treatment
Any post	One if the archived count of all posts in the economist-week is positive. Zero-post weeks remain.
Any #EconTwitter post	One if at least one archived post contains either `#EconTwitter` or `#econtwitter`. Other capitalization variants are not included.
Total posts	`count_tweet`, includes originals, retweets, replies, and quotes.
Retweets, replies, quotes made	`count_retweet`, `count_reply`, `count_quote`, mutually exclusive recorded reference types.
#EconTwitter count	`count_EconTwitter`, count using the two hashtag capitalization variants above.
Logged counts	Natural log of one plus the count. No relevant archived posts implies zero. A missing source engagement counter makes that metric's weekly sum and log missing.
Indirect cumulative exposure	A–B–C recipient-events from January 1, 2022 through the prior Stata week, excluding A = C. Main regressor is `log_preferred_stock_prior`.
Direct cumulative exposure	Dictionary-matched posts of any reference type by followed accounts, same dates and prior-week accumulation.
Received likes, retweets, replies, quotes	Snapshot public counters on original, reply, and quote posts created in the week. Plain retweets excluded. No authored content implies zero. Not reaction-week flows.
Recipient position	Standardized log PageRank in the binary directed pre-episode reply/retweet network. In-degree is an alternative. Missing linkage is not zero centrality.
Supplier position	Position-weighted exposure from mapped B or C, with unmatched exposure controlled separately.
Profile presence	Profile located in the May 2025 workbook using the platform identifier. Unlocated does not mean verified exit.
Positive recorded counter	Located profile with a positive recorded account-history post count. Not common-period recent activity.
Profile overlap	Located Twitter/X profile and at least one located alternative. The positive-counter analogue requires positive counters on both sides.
Exposure groups for endpoints	Above versus at or below the full-sample median archive event count.
Observed career proxy	Years through 2021 since earliest verified dated work, not years since PhD or entry into the profession.

The replication package reconstructs the panels, estimates, tables, and figures from archived API exports, the platform workbook, reconstruction rules, and reviewed academic extracts. It does not recover unavailable API histories or feed impressions.

References

Athey, Susan, Emilio Calvano, and Joshua S. Gans. 2018. "The Impact of Consumer Multi-Homing on Advertising Markets and Media Competition." *Management Science* 64 (4): 1574—90.

Biglaiser, Gary, Jacques Crémer, and André Veiga. 2022. "Should I Stay or Should I Go? Migrating Away from an Incumbent Platform." *The RAND Journal of Economics* 53 (3): 453—83.

Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, 3, edited by Mark Armstrong, and Robert H. Porter, 1967—2072. Elsevier.

He, Jiahui, Haris Bin Zia, Ignacio Castro, Aravindh Raman, Nishanth Sastry, and Gareth Tyson. 2023. "Flocking to Mastodon: Tracking the Great Twitter Migration." In *Proceedings of the 2023 ACM on Internet Measurement Conference*, 111—23. Association for Computing Machinery.

Jeong, Ujun, Ayushi Nirmal, Kritshekhar Jha, Susan Xu Tang, H. Russell Bernard, and Huan Liu. 2024. "User Migration Across Multiple Social Media Platforms." In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*, edited by Shashi Shekhar, Vagelis Papalexakis, Jing Gao, Zhe Jiang and Matteo Riondato, 436—44. Society for Industrial and Applied Mathematics.

La Cava, Lucio, Luca Maria Aiello, and Andrea Tagarelli. 2023. "Drivers of Social Influence in the Twitter Migration to Mastodon." *Scientific Reports* 13 (1): 21626.

Quelle, Dorian, Frederic Denker, Prashant Garg, and Alexandre Bovet. 2026. *Simple Contagion Drives Population-Scale Platform Migration.* arXiv preprint 2505.24801, Version 2. First Version (2025) Titled "Why Academics are Leaving Twitter for Bluesky". https://doi.org/10.48550/arXiv.2505.24801.

Wang, Xinyu, Sai Koneru, and Sarah Rajtmajer. 2026. *The Failed Migration of Academic Twitter: A Case Study of Precocious Adopters.* arXiv preprint 2406.04005, Version 4. https://doi.org/10.48550/arXiv.2406.04005.