

#EconTwitter and Elon Musk: Anatomy of a Failed Exodus

A study of network effects, platform disruption, and user stickiness.



Tin Cheuk Leung
Wake Forest University

Kwok Ping Tsang
Virginia Tech

The Puzzle: Why Did the Exodus Stall?

In late 2022, a major shock hit a dominant platform: Elon Musk's acquisition of Twitter. This led to controversial changes, large-scale layoffs, and widespread user dissatisfaction. A cohesive, influential community—#EconTwitter—publicly discussed migrating to alternatives like Mastodon.

Despite the negative shock and coordinated migration talk, **why did a mass exodus fail to materialize?** What does this tell us about the power of network effects?

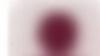
 Economist A

The new verification system is a mess. Platform is becoming unusable.

2:52 AM - Jun 10, 2024

 Economist B

With the latest changes, I'm seriously reconsidering my presence here.

 Economist C

I've set up an account on **Mastodon**. You can find me at **@[username]@[server]**. Come join!

 Economist D

The **#EconTwitter** migration seems to be gaining steam. Here's my new profile...

This Isn't Just About Twitter. It's a Test of Core Economic Principles.

The #EconTwitter episode provides a real-world laboratory for studying foundational concepts:

Network Effects & Lock-In

How strong are network effects in retaining users on an incumbent platform, even amidst dissatisfaction?

Katz & Shapiro, 1985; Farrell & Klemperer, 2007

Platform Competition & Switching Costs

What does it take to disrupt a dominant two-sided market? Our study focuses on the demand-side user perspective.

Rysman, 2009; Armstrong, 2006

Antitrust Implications

Our findings offer empirical evidence on how network effects can entrench market power, informing current debates on digital platform dominance.

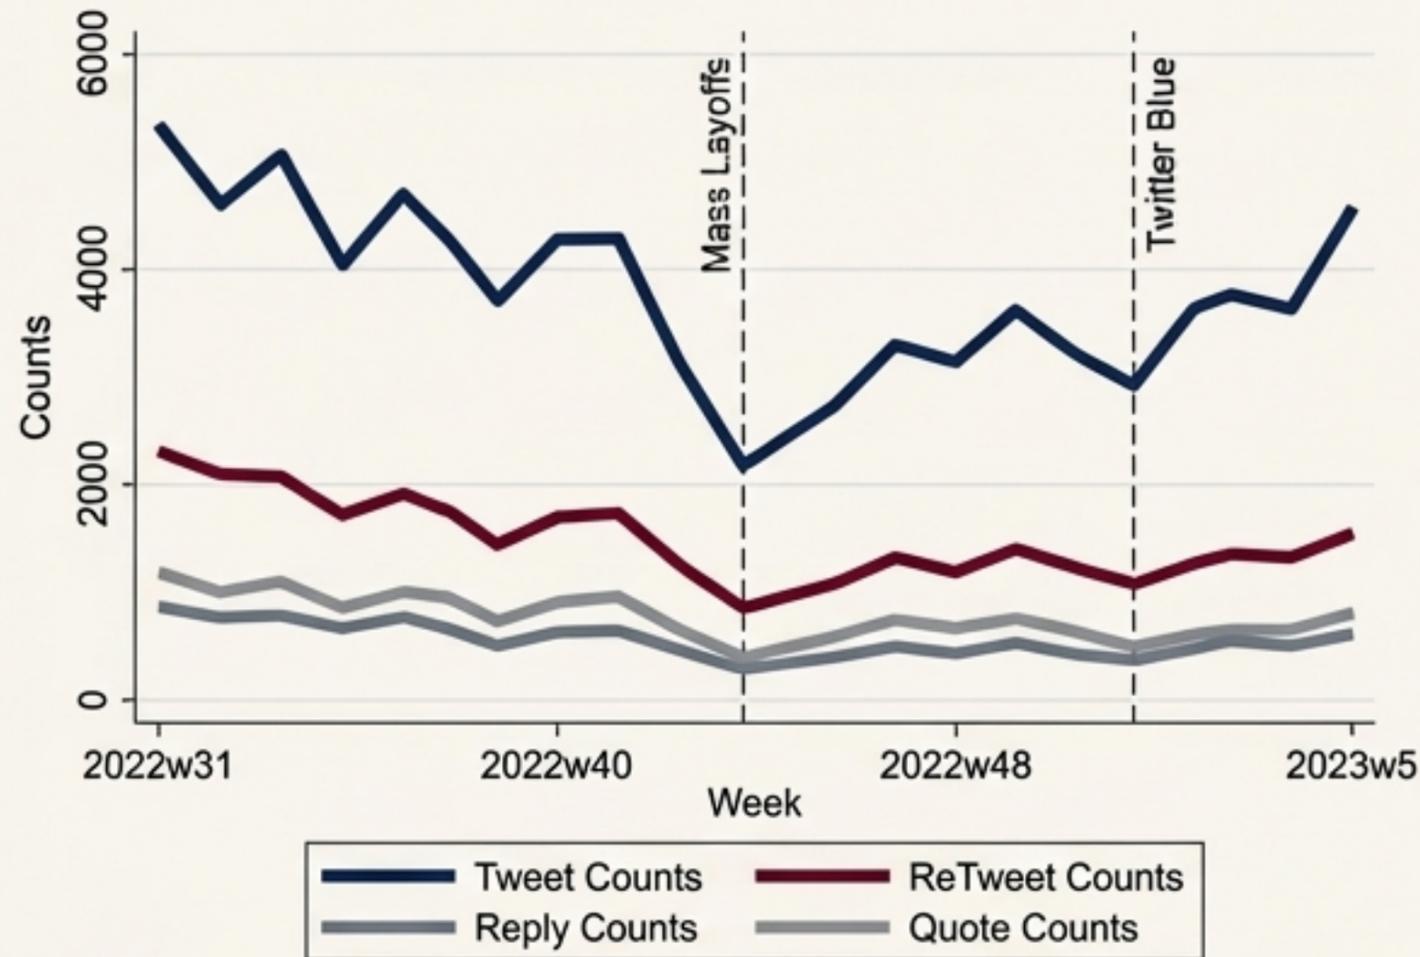
Khan, 2017; Hovenkamp, 2021

The Investigator's Ledger: A Multi-Faceted Dataset

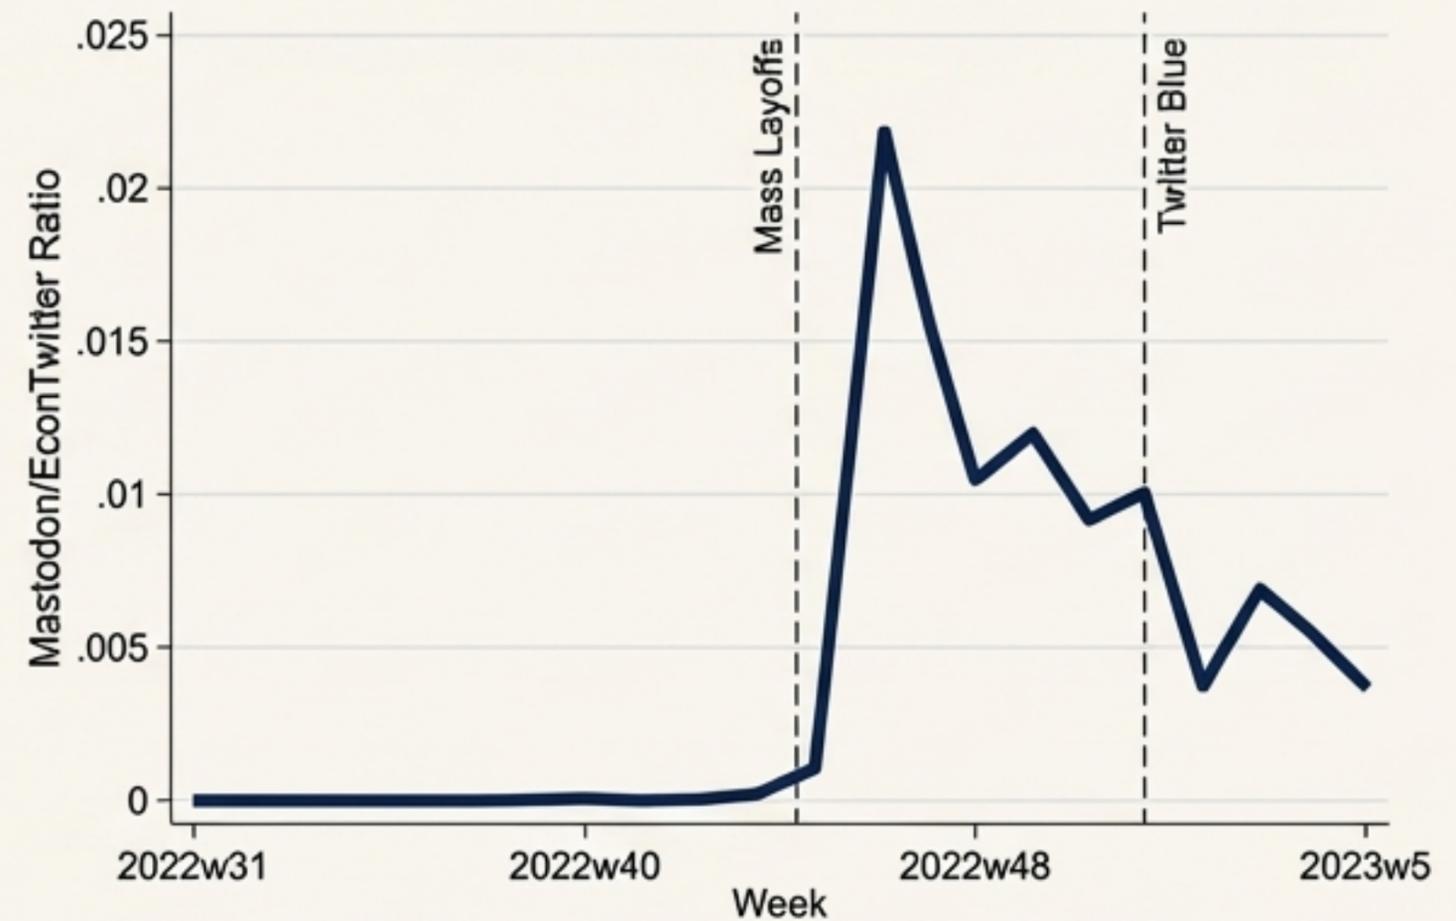
#EconTwitter Community Data	Economist Panel Data	Long-Run Follow-Up Data
<ul style="list-style-type: none">• Source: Twitter/X API (until March 2023)• Scope: 680,494 tweets containing the #EconTwitter hashtag.• Purpose: To track community-level trends and migration-related discussions.	<ul style="list-style-type: none">• Source: 2,321 economists identified from the RePEc database.• Scope: 4.3 million tweets, providing user-level weekly activity panels.• Purpose: The core dataset for our econometric analysis of individual behavior.	<ul style="list-style-type: none">• Source: Manual scraping of X, Bluesky, and Mastodon profiles.• Scope: Snapshot of activity as of May 2025.• Purpose: To assess long-term platform persistence, abandonment, and multi-homing.

First Clue: A Community in Turmoil, A Spike in Migration Talk

#EconTwitter activity dipped sharply post-layoffs but rebounded by early 2023.



Talk of Mastodon surged immediately after the shock.



The events triggered a significant, but perhaps temporary, behavioral response and a clear search for alternatives.

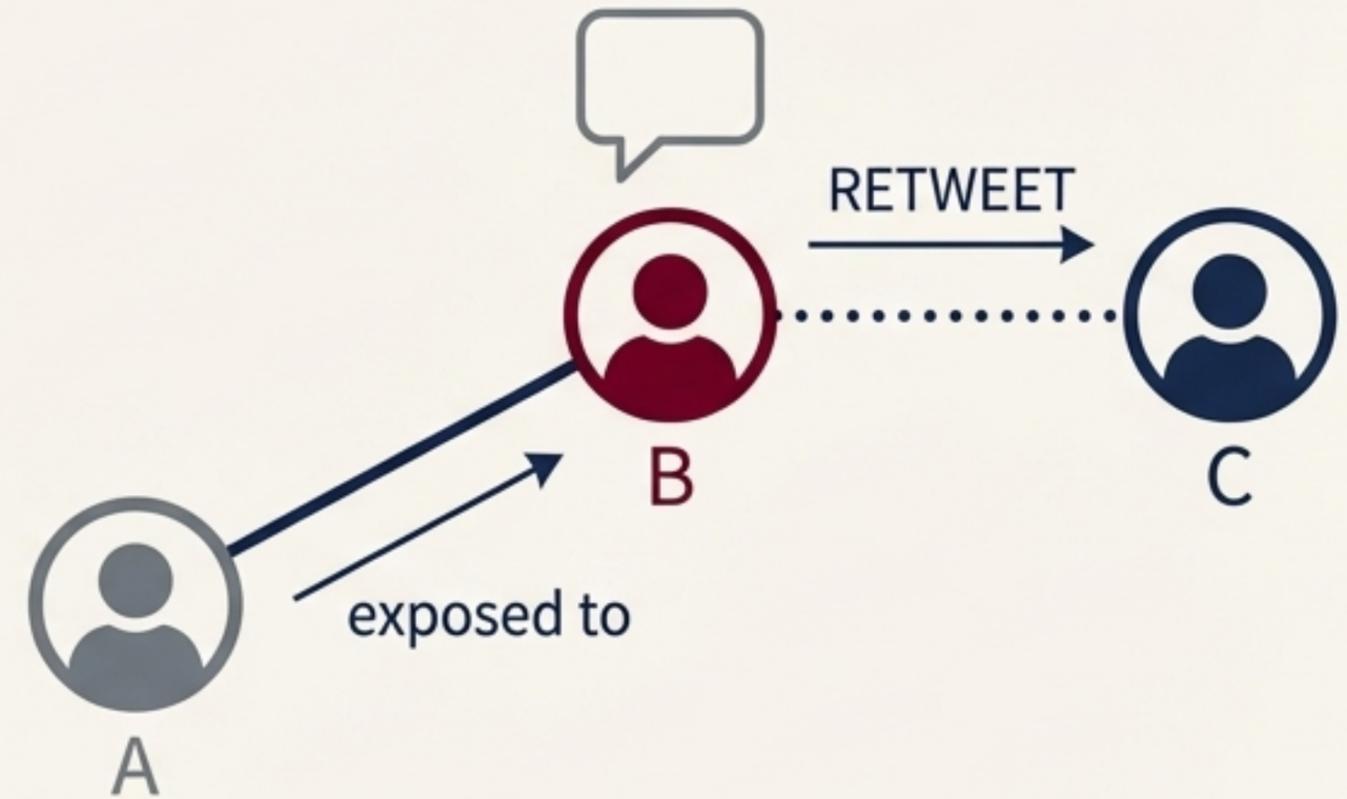
The Investigator's Tool: Isolating Peer Influence from Homophily

The Challenge

How do we measure exposure to migration talk? A simple count of followed accounts' posts is confounded by homophily (economists with similar views follow each other). This is **Manski's** (1993) reflection problem.

Our Solution: An Indirect "Peers of Peers" Exposure Measure

- **Definition:** We measure **RetweetMastodon_it** as the number of Mastodon-related retweets an economist i sees from accounts they follow (**B**), where the *original author* of the tweet (**C**) is *not* followed by i .
- **Logic:** This design, inspired by Bramoullé et al. (2009), reduces selection bias since economist i did not choose to follow the original author **C**. It captures how peers transmit information through the network.



Who Was Most Exposed to Migration Talk?

The 'High Exposure' group was more embedded and active on the platform.

Characteristic	Low Exposure	High Exposure
Avg. Following	825	1,923
Avg. Weekly Tweets	9.9	15.2
Avg. Followers	14,635	7,894

Interpretation

- Economists with higher exposure to Mastodon talk were not necessarily the most 'famous' (fewer followers on average).
- Instead, they were more active participants who followed more accounts.
- These observable differences underscore the need to control for user heterogeneity in our models.

The Initial Evidence: Higher Exposure Correlated with a Short-Run Dip

Method

Difference-in-Differences (DiD) estimation comparing behavior before and after the November 2022 layoffs, based on pre-existing exposure levels.

$$Y_{it} = \beta * (Layoff_t \times MastodonExposure_{it}) + \dots$$

Results

Outcome (Post-Layoff)	Effect of Mastodon Exposure
Tweet (Yes/No)	-0.0121*
#EconTwitter (Yes/No)	-0.0109*
Tweet Count	-0.539*

Interpretation

- In the weeks following the layoffs, economists with greater cumulative exposure to Mastodon discussions were significantly less likely to tweet on X.

In Plain English

A one standard deviation increase in exposure was associated with a ~6 percentage point lower probability of tweeting in a given week. This supports the “attempted exodus” narrative.

A Deeper Look: Do These Effects Hold Up Within Economists Over Time?

Limitation of DiD

The previous result relies on comparing different types of economists. Are we just picking up that less-attached economists were more exposed?

A More Rigorous Test: Economist Fixed-Effects

We now use a model that controls for all stable, individual characteristics and common weekly trends.

$$Y_{it} = \alpha * RetweetMastodon_{i,t-1} + \delta_i + \phi_t + \varepsilon_{it}$$

- δ_i : Economist Fixed Effects (controls for any permanent differences between users)
- ϕ_t : Week Fixed Effects (controls for platform-wide events affecting everyone)

This isolates the relationship between a change in an economist's exposure in one week and their activity in the *next* week, within a narrow 10-week window around the layoffs.

The Plot Thickens: Effects Shrink Substantially with Fixed Effects

When we control for individual heterogeneity, the relationship between Mastodon exposure and X activity becomes much weaker.

Table: Simplified Summary of Table 5

Outcome (Week t)	Effect of Lagged Mastodon Exposure (Week t-1)
Tweet (Yes/No)	0.00186 (Not sig.)
#EconTwitter (Yes/No)	-0.00803*
Tweet Count	-0.372*

Interpretation

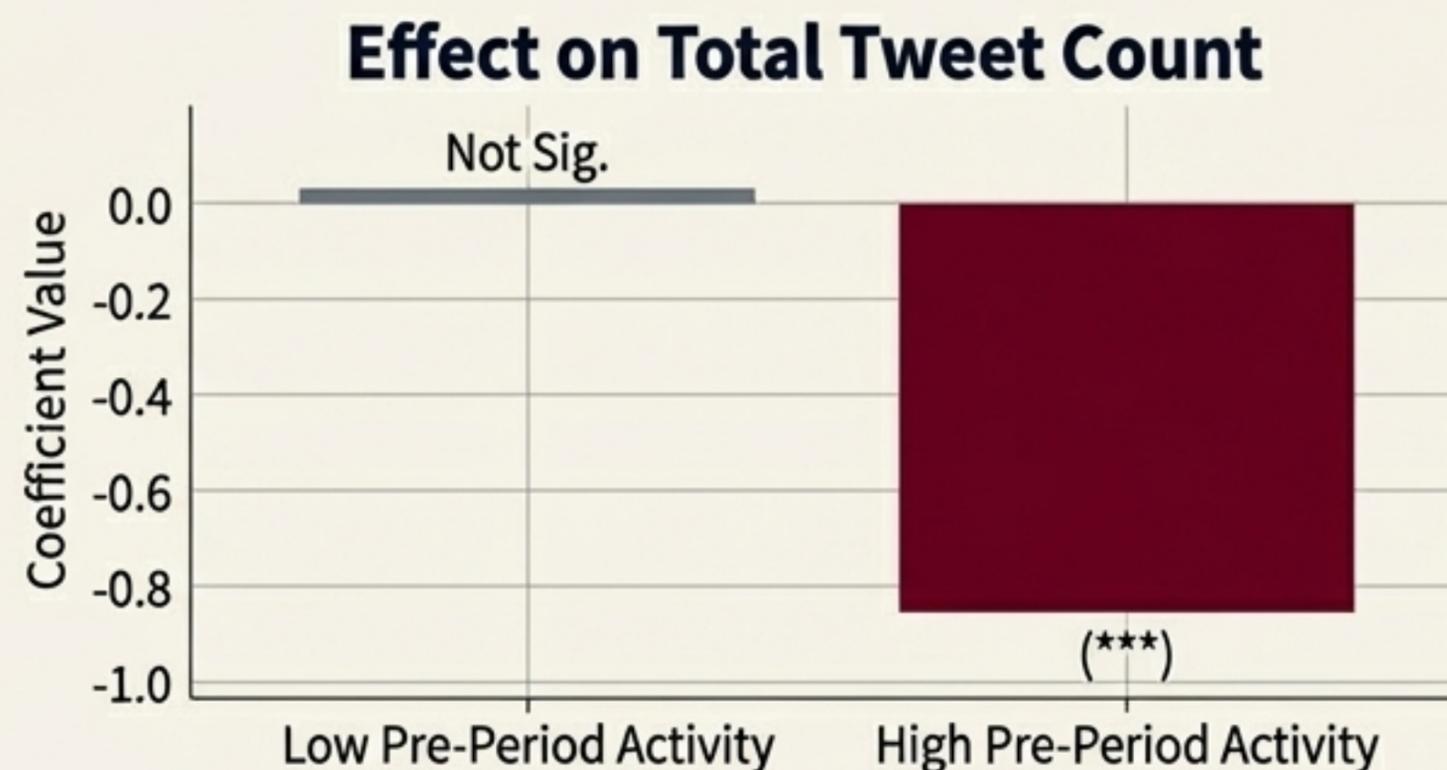
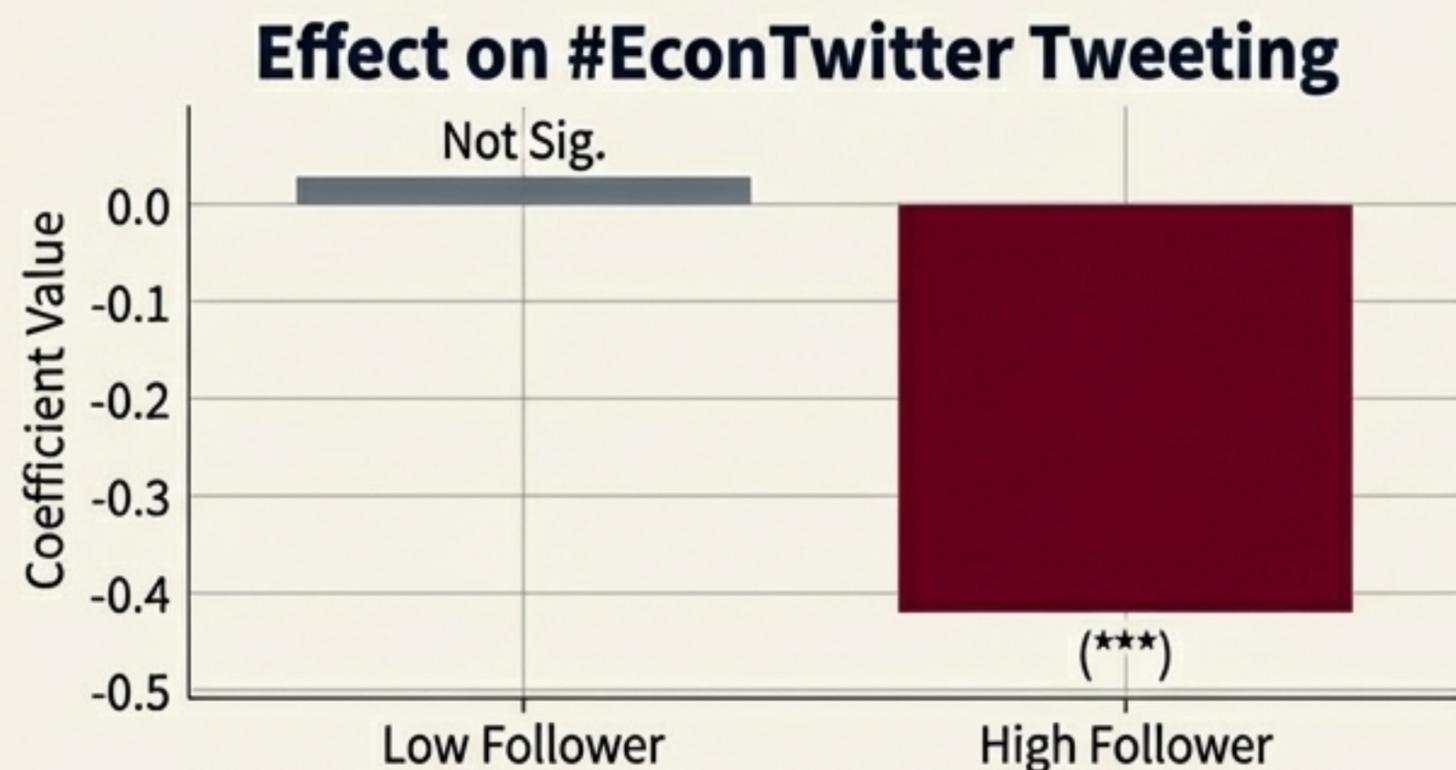
- Short-run exposure is not associated with stopping tweeting altogether.
- There is a statistically significant, but very small, reduction in the *volume* of tweets. A one standard deviation increase in weekly exposure is linked to a reduction of only ~0.5 tweets.

**Conclusion: Behavior on X is highly “sticky.”
The exodus was more of a whisper than a roar.**

Unmasking Heterogeneity: The Most Connected Users Reacted Most

Question: Did the modest average effect mask stronger responses in certain groups?

Answer: Yes. The negative relationship between Mastodon exposure and X activity was concentrated among the most active and central economists.



Takeaway: This is consistent with a 'leader-follower' dynamic, where the most embedded network members were the ones reducing activity, but their actions were not enough to trigger a cascade.

The Verdict, Two Years Later: Multi-Homing, Not Migration

We tracked the 2,321 economists' presence and activity on X, Bluesky, and Mastodon as of May 2025.
We estimate a cross-sectional regression: $LongRunOutcome_i = \beta * MastodonExposure_i + \dots$

Results - Presence (Extensive Margin)

Higher 2022 Mastodon exposure is associated with:

- A **3 ppt lower** probability of being active on X.
- A **9 ppt higher** probability of being active on Bluesky.
- A **6 ppt higher** probability of being active on Mastodon.

Results - Activity (Intensive Margin)

Higher 2022 exposure is associated with:

- **~300 fewer** posts on X...
- ...and **~64 more** posts on Bluesky between 2023-2025.

The primary long-run outcome was platform diversification ('multi-homing'), not a wholesale abandonment of Twitter/X.

Solving the Puzzle: The Anatomy of a Failed Exodus

- 1 The Shock:** Musk's takeover created genuine user discontent and a spike in migration talk within #EconTwitter.
- 2 The Initial Reaction:** This talk was correlated with a brief, statistically significant dip in X activity. The 'exodus' seemed plausible.
- 3 The Attenuation:** However, more rigorous fixed-effects models show this effect was very small in magnitude, revealing immense user stickiness.
- 4 The Long-Run Outcome:** Two years later, there was no mass departure. Instead, users with high exposure were more likely to adopt *additional* platforms while largely remaining on X.

The immense value of the incumbent network proved more powerful than widespread dissatisfaction with platform governance. For a community built on connection, the cost of leaving a network where everyone else remains is simply too high.

Contribution: What This Case Teaches Us About Digital Economies

1. A Powerful Demonstration of Network Effects

We provide rare empirical evidence of network resilience. Even with a major negative shock, a coordinated “exit” campaign, and a viable alternative, the incumbent platform’s network effects were **overwhelmingly strong**. **The resilience was overwhelming.**

2. Insights for Antitrust and Platform Regulation

The findings highlight the **formidable barriers to entry** in markets dominated by network effects. User dissatisfaction alone is insufficient to foster competition. This has direct relevance for policy debates around interoperability and platform power. **The barriers are formidable.**

3. Understanding Academic Digital Communities

Specialized professional networks like #EconTwitter derive **immense value from concentration**. The desire to remain connected to peers is a powerful force for platform stability. **The value from concentration is immense.**