A Data Collection Details

First, we acquired a list of 1,261 colleges and universities. For each of these institutions, we tasked a crowdworker with identifying the website of all four of the social science departments in our study (sociology, political science, psychology and economics), for a possible total of 5,044 websites. The crowdworker either identified the faculty directory or indicated that none was available or that that institution did not have that department.

For each of these websites identified in the first round, we recruited a second set of crowdworkers to navigate there and to copy the name, email address and gender presentation of each person listed on the faculty page. The crowdworkers were asked to search this information on Google and on the Twitter platform itself to try to find that person's Twitter account.

Armed with the resulting "seed" list of academic accounts, we then expanded the sample by scraping the list of accounts each of these individuals followed on Twitter. We ordered these accounts by how many followers they had in our seed sample, and took the most popular 10,000 to expand our data. This yielded a set of roughly 42,000 accounts that we characterize as "social science Twitter", although we underscore that many of these may not be individuals or even academics at all. In addition, several accounts show up in more than one snowball sample, with particular overlap in our snowball samples of economists and political scientists. When we include only distinct accounts, our number of 42,000 total accounts drops to a little over 33,000.

While the sample was collected in pursuit of a broader research question, in the fall of 2022 we saw an opportunity to use these accounts to characterize what effect, if any, Elon Musk's takeover of Twitter would have on the academic community writ large. In the waning days of the freely available Twitter API, we scraped the full set of posts made by each of these accounts dating back to January of 2020. Doing so meant that we had to rely on the "full archive", which contains only original posts written by the owners of these accounts, precluding our ability to systematically count their retweets and replies all the way back to 2020. (Time constraints combined with quota limits meant that we only have the full set of posts from January 2020 to February 1st, 2023 for 15,761 of the accounts.)

In addition to the historic set of original posts from the full archive, we also scraped each account's 3,200 most recent tweets (including both original posts, replies, and retweets) on four separate occasions: June of 2021, June of 2022, November of 2022, and February of 2023. Furthermore, beginning in early November of 2022, we began two daily scrapes of the accounts: one that simply took a snapshot of how many followers the account had each day, and the other that gathered the ten most recent posts each day.

B Methodological Details of BCPA

This method assumes that observations are drawn i.i.d. from two underlying distributions, and searches for the break point that separates these two distributions (Aminikhanghahi and Cook, 2017).

We bootstrap 5,000 users 100 times at random from our data and estimate structural break points based on either the total number of daily active accounts among them, or on the total number of daily tweets they produced, disaggregating based on tweet type (retweets, replies, quote tweets, and originals). As illustrated, across all types of tweets (rows) and either type of measure of engagement (columns), the modal inflection point is November 18th, 2022, with the exception of the total number of retweets observed, where the inflection point is identified as November 23rd, 2022. Aggregating across behavior measures and tweet types, November 18th is chosen in 71.5% of bootstrapped samples, followed by November 23rd in 19.8% of bootstrapped samples, with the remaining inflection dates appearing in less than 3% of bootstrapped samples. Estimated on the full data, the inflection points are either November 23rd for retweets, or November 18th for all other types of engagement.

C Methodological Details of Who Left?

Formally, we run a simple difference-in-differences of the form:

$$y_{it} = \alpha_i + \beta_1 \text{Verified}_i + \beta_2 \text{Post}_t + \beta_3 \text{Verified}_i * \text{Post}_t + \lambda_1 t + \lambda_2 t^2 + \varepsilon_{it} \qquad (1)$$

where y_{it} is the logged number of tweets written by user i at time t, Verified_i is an indicator for whether user i was verified as of November 30th, 2022 (i.e., prior to the change in verification status), and Post_t is an indicator for whether the outcome is observed prior to (0) or following (1) November 19th, 2022. Importantly, α_i are account fixed effects, meaning we are identifying variation *within* users over the period of analysis, mitigating concerns of selection bias driving our results. Furthermore, we control for curvelinear time trends ($\lambda_1 t$ and $\lambda_2 t^2$). We estimate the preceding specification subsetting the data to the types of tweets being written (original tweets, retweets, reply tweets, and quote tweets) and to the period starting on October 19st, 2022, one month prior to the inflection date identified above. Our conclusions are robust to looking at the full period of data collection (starting on January 1st, 2021), and to a three way interaction where tweet type is interacted with verification status and the post indicator.