The Vibes are Off:
Did Elon Musk Push Academics off Twitter?

August 2, 2024

Abstract

Twitter has been a prominent forum for academics communicating online, both among themselves and with policymakers or the broader public. Elon Musk's take-over of the company brought sweeping change to many aspects of the platform, including the public access to its data; Twitter's approach to censorship and mis/disinformation; and tweaks to the affordances of the platform. In this letter, we take up a narrower empirical question: what did Elon Musk's takeover of the platform mean for this academic ecosystem? Using a snowball sample of more than 15,700 academic accounts from the fields of economics, political science, sociology, and psychology, we show that academics in these fields reduced their "engagement" with the platform, measured with either the number of active accounts (i.e., those registering any behavior on a given day) or the number of tweets written (including original tweets, replies, retweets, or quote tweets). We further test whether this drop-off in engagement differed by account type, finding that verified users were significantly more likely to reduce their production of content (i.e., writing new tweets or quoting others' tweets), but not their engagement with the platform writ large (i.e., retweeting or replying to others' content).

1 Introduction

Twitter has been a prominent forum for academics communicating online, both among themselves and with policymakers or the broader public. Twitter has thus played a significant role in the practice of social scientific research in the 2010s and early 2020s.

The growth of this online academic community has weathered a number of changes in how the platform operates, but perhaps none have been as dramatic as the take-over by Elon Musk in the fall of 2022. Musk's ownership brought a range of changes to the platform, including mass layoffs at Twitter; the reinstatement of former President Donald Trump's account which had been deactivated following its role in the January 6th, 2021 riots at the capitol; the reinstatement of tens of thousands of accounts that had been suspended for violating the platform's Terms of Service; and the corresponding permission of posts containing mis and disinformation that had previously been barred, as well as a general weakening of the enforcement of these scaled-back policies (Dang, 2023; Rutenberg and Conger, 2024). Given subsequent developments in global politics, this entailed a less-moderated conversation around topics of interest to social scientists such as Russia's invasion of Ukraine, the Hamas attacks on Israel and subsequent conflict in Palestine, former president Trump's legal battles, and more. These changes each influenced the broader social network characteristics in ways that are not yet fully understood, quantitatively, but which entailed a shift in the user experience we colloquially refer to as "vibes."

In this letter we ask and answer a simple research question: did Elon Musk push academics away from the platform? Anecdotal evidence of a broader move off Twitter is abundant, perhaps best exemplified by the pervasive inclusion of profile names for alternative social media platforms in user bios after Musk's takeover. But how widespread was the desertion? When, exactly, did it happen? And perhaps most importantly, who was most likely to reduce their use of the platform?

Using a snowball sample of academics across the fields of political science, economics, sociology, and psychology, we show that there was a substantively significant reduction in engagement with the platform, especially in terms of original tweets and quote tweets.[1] Furthermore, we show that the accounts most responsive to Musk's takeover were higher profile, as proxied by the (original) blue check which indicated verified accounts. The reduction in engagement from these accounts is arguably more damaging for Twitter, and is almost certainly a greater opportunity cost for the account holders who had developed substantial online reputations using these accounts.

Conversely, the greater visibility of these accounts might also have increased the social pressure to abandon the platform and signal objections to Musk's various decisions, including reactivating thousands of alt-right accounts that had been deactivated by Twitter's previous moderators, and most notably inviting ex-president Donald Trump back to the platform which he had built his reputation on. The vibes on academic Twitter have certainly changed, and in this manuscript we attempt to provide a quantitative answer to the questions of *when, exactly?* and *for whom?*

2 The Role of Twitter in Scholarly Communication

To understand why Musk's ownership of Twitter might drive academics from the platform, it is important to note the utility that academics may have obtained in the first place. Twitter has been a prominent forum for academics communicating online, among themselves (Klar et al., 2020), with policymakers the broader public (Jester, 2022; Jünger and Fähnrich, 2020), and even their students (Sweet-Cushman, 2019).

[1]Throughout this paper, we use the term "engagement" to refer to the principle behaviors associated with Twitter: posting original content, replying to others' posts with original content, retweeting others' posts (i.e., rebroadcasting without annotation), quote-tweeting others' posts (i.e., rebroadcasting with annotation), and favoriting others' posts.

Twitter has thus played a significant role in the practice of academic and particularly social scientific research in the 2010s and early 2020s. Ke, Ahn and Sugimoto (2017) find that Twitter use is actually most common among social scientists; in this paper, we restrict our analysis to academics working in the disciplines of Political Science, Economics, Sociology and Psychology.

Has this been a positive development? Perhaps it is too soon to say. In broad strokes, some observers point to Twitter as a site for revolutionary political organization by marginalized groups (Tufekci, 2017), or for direct democratic communication between elected officials and constituents (Barberá et al., 2019); others, to Twitter as precipitating the decline of traditional party organization and the bedrock liberal values of toleration and free speech (Gerbaudo, 2019).

In more concrete terms, existing research paints a mixed picture. On the one hand, Twitter has undeniably increased the volume and visibility of academic communication, enabling public conversations that otherwise would have remained within faculty lounges or not taken place at all. In Grossmann (2021)'s account of recent social science reforms, Twitter plays an important role in enabling conversations about research best practices that include the recent "credibility revolution" and conventions around pre-registration, data sharing and open access publication.

Twitter has also provided opportunities for scholars from diverse fields and different levels of seniority to interact. Searles and Krupnikov (2018) provides valuable advice on how faculty mentors should guide their graduate students on how to use the platform to promote their work, for example, illustrating how existing seniority hierarchies were made more porous on Twitter. Especially in the wake of the COVID19 pandemic, Twitter was a valuable resource to preserve existing academic networks for knowledge dissemination. Bandula-Irwin and Kitchen (2022) describe how the platform could even be used as a replacement for traditional in-person conferences. Given that the pandemic seemed to disproportionately

affect female academic's continued productivity (Kim and Patterson, 2022), Twitter may have prevented the problem from being further exacerbated.

On the other hand, many of the structural inequalities associated with legacy academia are reproduced on Twitter, with women and junior scholars being less central and less amplified than the male peers (Bisbee, Larson and Munger, 2022). There is also evidence that this is the case for the related case of journalists on Twitter (Usher, Holcomb and Littman, 2018).

But beyond these broader descriptions, Twitter is a social network where academics can feel seen by their peers, where they can learn in-group signifiers, and where they can express themselves in ways that produce positive feelings of solidarity (Jünger and Fähnrich, 2020). In this sense, the academic Twitter community operates according to the same psychological incentives as any other online community (Bail, 2021). And, for most of Twitter's existence at least, the broader platform was characterized by many of the qualities of the social science community itself: it was younger, better educated, and more progressive in its politics (Pew Research Center, 2018).

3 The Musk Effect

Why might Musk's acquisition of the corporation have driven academics off of the platform? There were a number of significant policy changes, including changes to Twitter's verification process, regulations on account names, and the removal of the free API that many scholars had relied on for empirical research across a range of areas of scientific inquiry.

Some academics also expressed distaste for Musk's vibe. Musk had long used the platform in ways that many academics disliked. His use of Twitter's "polls" feature to determine several of his early policies evinced an unscientific understanding of the principle

of random sampling. He endorsed and indeed amplified dangerous and unfounded conspiracy theories about US elected officials. More broadly, he demonstrated a puerile childishness that was often at odds with the predominantly liberal worldview and at least nominally staid professionalism of Twitter's academic community.

We argue that some combination of these features of the threat and then reality of Musk's ownership of the Twitter corporation influenced academics to either quit Twitter altogether or at least to reduce their engagement with the platform (i.e., to "disengage"). The policy changes and personality of Twitter's new owner were hard to avoid and may have made the experience of using the platform less palatable. On the other hand, these same attributes may have stimulated a type of ideological boycott, in which academics disengaged with Twitter as a political strategy, indicating their intellectual and moral opposition. We expect this pressure to be greater among more popular users, since the visibility of their continued use of the platform (and the assumed tacit endorsement of it) would be greater. Conversely, these same users are also more sensitive to a sunk cost logic in which the prospect of shutting down a "valuable" Twitter account is more costly.

Data limitations make a careful analysis of these mechanisms beyond the scope of our paper. Nevertheless, we examine heterogeneous effects in our analysis, finding that more popular accounts were more likely to disengage following Musk's takeover. These patterns are consistent with the social pressure explanation for disengagement described above, although we are careful to acknowledge the limitations of a causal interpretation of our data.

4 Data and Methods

In early 2021, we used Amazon's Mechanical Turk platform to begin the process of collecting what we believe to be the most comprehensive dataset of US-based social scientists on Twitter. This procedure had several steps, detailed in the Appendix.

The collection and analysis of large quantities of social media trace data is now *de rigeur* across the quantitative social sciences and indeed many other realms of social, political and economic life. There are thousands of peer-reviewed academic studies using the same methods that we use here. As is common with these studies, our respective institutional ethics review boards deemed this exempt from full review. Still, there are legitimate concerns about the ethics of potentially causing public scrutiny of social media users who are not accustomed to (and who are not obviously seeking) attention beyond a local network. For this reason, we do not identify any accounts by name in this manuscript, and our replication materials anonymize the accounts.

Our analysis consists of a series of descriptive results, relying on a before/after comparison of several different measures of engagement with the platform in order to reach a causal conclusion. These measures include the number of active accounts on Twitter, the frequency with which these accounts tweeted, and the levels of engagement with their content, measured as either likes, retweets, replies, or quote tweets.

For the sake of simplicity, we focus on Musk's official takeover on October 28th, 2022 as the cut point around which to compare before/after changes in academic use of the platform. However, we also implement a Bayesian change-point analysis which locates the structural break in Twitter usage by academics around November 19th, which happens to be the day when Musk announced he would reinstate Trump's Twitter account (see Figure 4 below).

Causally identifying Elon Musk's impact on social scientists on Twitter is complicated by two alternative stories. The first is that Musk's ownership started just prior to the standard academic winter break, during which social scientists disengage from the platform naturally, rendering any empirical association we document spurious. Second is the potential for reverse causality. While it is hard to imagine that social scientists themselves would have any influence over either Musk's interest in purchasing the platform or the previous owners' interest in selling, there was a steady decline in engagement among our sample that mirrors

a broader disengagement that paved the way for Musk's takeover.

We implement a variety of methods to address these concerns, although we do not claim that our findings are perfectly causally identified. First, we calculate a per-account change in activity between the weekly data we measure in 2022 relative to that observed during the same week in 2021. While this approach is coarse, it does account for seasonal variation in engagement that might otherwise produce a dramatic but spurious association between Musk's takeover and engagement among users in our sample.

Second, we let the data speak by implementing Bayesian Change Point Analysis (BCPA). This method detects structural breaks in time series data under the assumption that observations are drawn from one of two independent and identically distributed random processes. We bootstrap sample 5,000 accounts at a time to demonstrate that the method selects November 18th, 2022 in 71.5% of these bootstrapped samples as the break point separating one era of social scientist engagement with the platform from another. Importantly, this is one month after Musk's official takeover, suggesting that it is unlikely to reflect a reverse causality story in which declining academic engagement prompted Musk's interest. Even more tellingly, this breakpoint coincides with Musk's promise to reinstate former President Donald Trump's Twitter account, which had previously been suspended by Twitter's previous ownership on January 8th, 2021, in response to Trump's role in the Capitol Riots.

Third, we implement a simple difference-in-differences specification with account fixed effects, subsetting to one month prior to and following the cut point identified with the BCPA. While our second difference is of theoretical interest – specifically examining whether verified accounts were more or less likely to disengage – the first difference, estimated with account fixed effects, cubic polynomial time trends, and subsetting to a period just prior to and just following Musk's takeover, provides further evidence against the alternative explanations for the association between Musk's ownership and social science Twitter's disengagement.

5 Results

Did Musk's takeover of Twitter co-occur with a decline in the academic community of scholars who had been using the platform to share and discuss research? A simple overtime plot of the number of users posting on a given day suggests that the answer is a qualified "yes." As illustrated in the left-panel of Figure 1, while the average number of daily active accounts had been declining since 2021, it experienced a dramatic drop in early November of 2022, just as Musk took over. More strikingly, the level of engagement (operationalized by the total number of tweets written by academics, right panel of Figure 1) with the platform exhibits an even more precipitous decline slightly later in November.

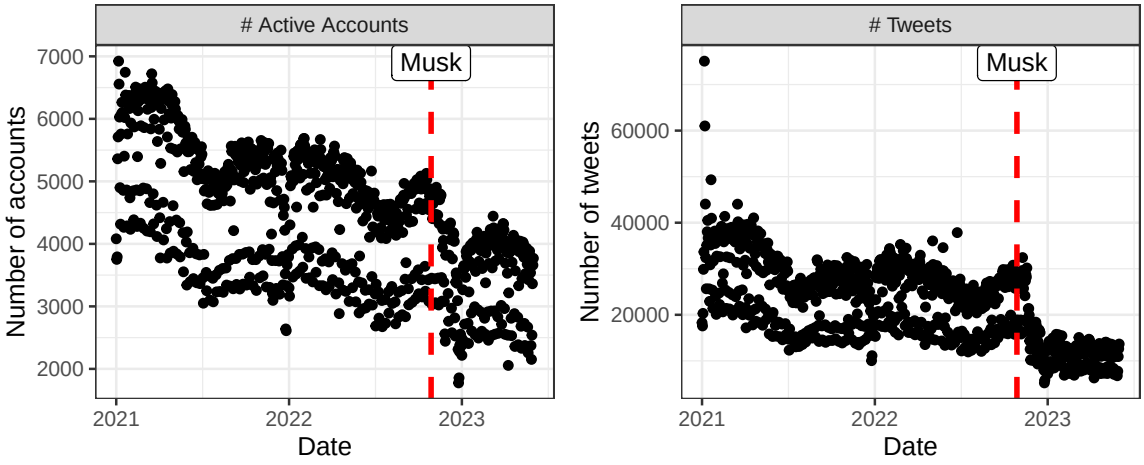

Figure 1: Descriptive time series plots of academic use of Twitter. Left panel depicts the total number of users who posted each day. Right panel depicts the total number of tweets posted each day. Vertical dashed lines indicate the date Musk officially took over as Twitter CEO on October 28th, 2022.

Note, however, that Musk's takeover occurs just prior to the typical academic winter break. As previous years illustrate, this is generally a period of reduced Twitter activity in our Social Science Twitter sample, with especially pronounced reductions during the week between Christmas and New Years. The 2022/2023 winter break period, though, exhibits a stronger decline than that observed in the 2021/2022 break, one that starts much earlier

than the preceding year (around the time of Musk's take-over), and exhibits no evidence of a sharp rebound to the pre-break levels.

Nevertheless, it may be that the decline in both active users and average tweets per day that is associated with Musk's takeover is actually just a spurious correlation driven by a seasonal trend. To investigate this claim, we calculate a per-account annual change by week. Specifically, for each of our 15k accounts, we first calculate their average activity per week, and then calculate the annual change measuring how many tweets were written in a given week in 2022 versus the amount written by the same user in the same week of 2021. There are of course many other factors at play here such as changes in an individual's position in the academy. Nevertheless, with fifteen thousand accounts to draw on, we argue that the these individual-level factors should produce more or less random variations and therefore amount to attenuation bias that our large sample is sufficient to overcome. We present these annual comparisons in Figure 2, coloring the bars green if 2022's weekly average engagement was higher than the same week in 2021, and red if it was lower. Again, despite other fluctuations over the course of the year, the drop off following Musk's takeover is stark and persistent. In particular, the post-Musk takeover sees the greatest decline in average weekly tweets for all of 2022 relative to 2021, with the exception of the first week in January, whose 2022 engagement is being compared to the Capitol Riots on January 6th, 2021 and the attendant flurry of tweets from social scientists of all stripes. Even after controlling for seasonal trends by calculating the annual change in activity, we still note a striking decline in academic Twitter use that co-occurs with Elon Musk.

5.1 Bayesian Change Point Analysis

Thus far, we have relied on descriptive plots that suggest academics reduced their use of Twitter in response to Musk's takeover. While the patterns are striking, this approach nevertheless prevents us from making inferential statements about these patterns absent our

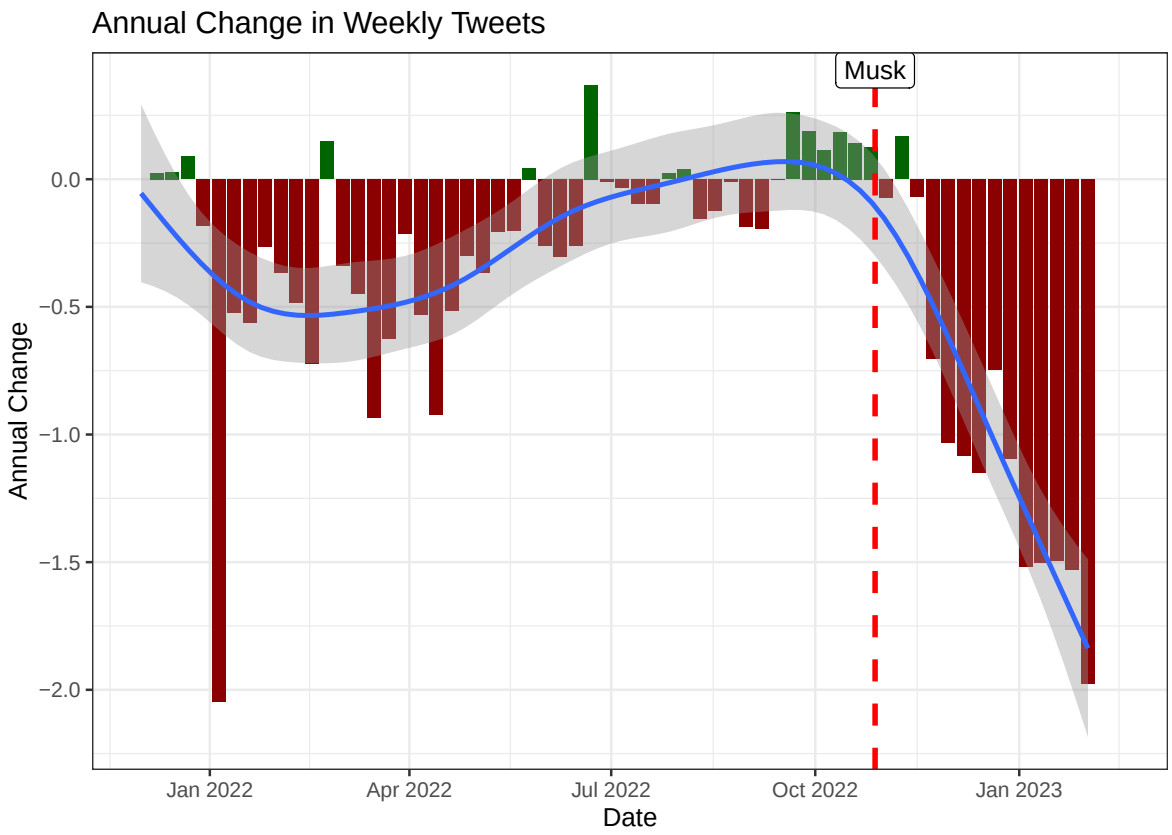

Figure 2: Account-Level Change in Tweets Posted Per Calendar Week

assumption that October 28th is the pivotal cause. To overcome this limitation, we turn to Bayesian change-point analysis. For details, see the Appendix.

Figure 3 visualizes the results where the y-axes indicate the number of bootstrapped simulations in which each date (x-axes) was chosen as the breakpoint.

Taken together, there is clear and consistent evidence pointing to between November 18th and the 23rd as when the academics comprising our social science sample disengaged from the platform.

This period is proximate to Musk's take-over on October 28th, 2022, but clearly with a bit of a lag. A review of the events of the period provides a post hoc explanation: November 19th was when Musk reinstated Donald Trump's twitter account and announced it with the tweet depicted in Figure 4. Although Musk's actions over the latter half of 2022 were

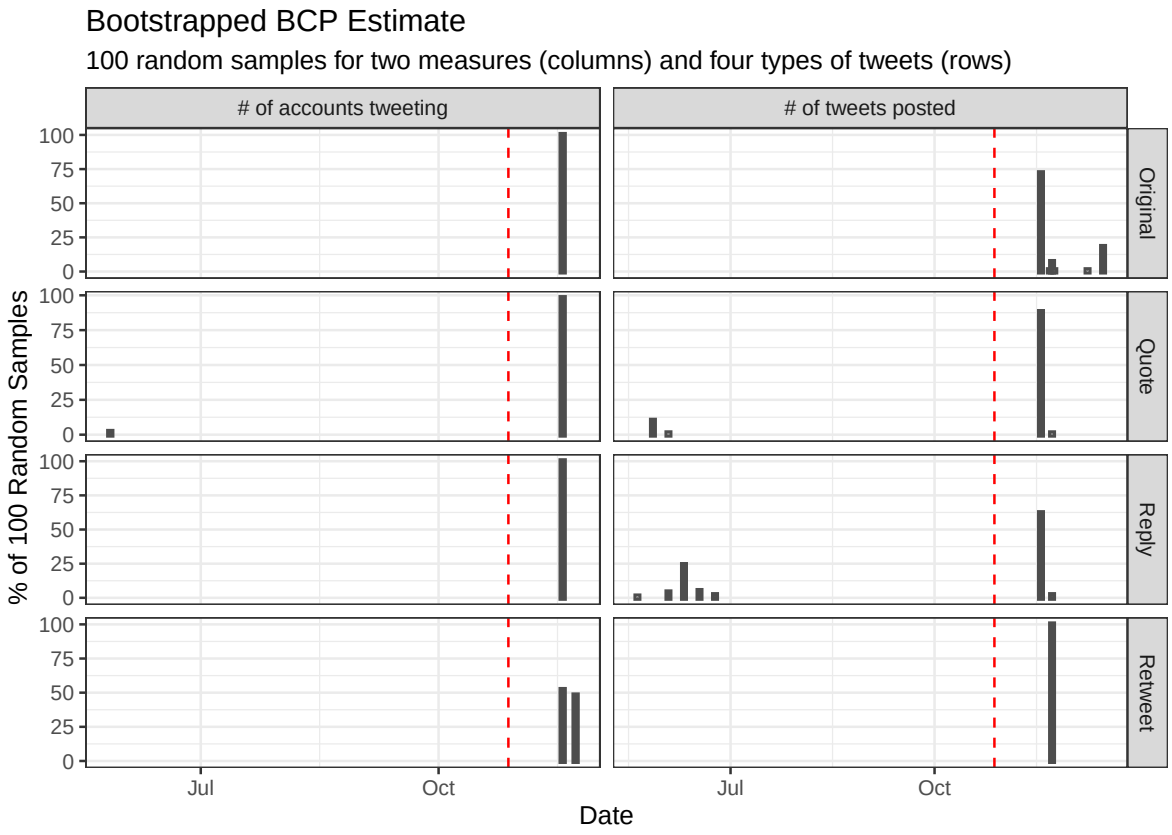

Figure 3: Agnostic Estimation of Structural Break in Tweeting Activity

regularly norm-breaking, for many social scientists, the idea of allowing Donald Trump back on to the platform which had been the focal point of his election fraud conspiracies may have been too much to take (to say nothing of both the hubris and the statistical ignorance of Musk for thinking that his "poll" actually measured the voice of the people). In the subsequent analyses, we use November 19th, 2022 as the inflection point to evaluate heterogeneous effects by account type.

5.2 Who Left?

As discussed above, we posit that the decision to disengage with the platform is a function of the account's popularity, although this intuition can cut in one of two ways.

Figure 4: Musk's tweet announcing the decision to re-activate Donald Trump's twitter account.

To test, we separate our data by whether the account was verified prior to November 30th, 2022 (i.e., when a blue check still meant something other than a willingness to spend $8 per month).

Table 1 presents the results for each tweet type in columns, highlighting that verified users are significantly more likely to reduce their engagement on the platform, but only in terms of original tweets (column 1) or quote tweets (column 4). Across all specifications, the Post_t indicator is negative and significant, indicating that – within users – engagement dropped off significantly after November 19th, 2022, relative to October 1st, 2022.

These results are consistent with a reputational mechanism by which the most visible

Table 1: Logged tweets per day predicted by period interacted with verified accounts

Model:	Original Tweets (1)	Retweets (2)	Replies (3)	Quotes (4)
Variables				
Post 11/19	-0.025***	-0.148***	-0.036***	-0.033***
	(0.004)	(0.005)	(0.006)	(0.004)
Verified × Post 11/19	-0.032***	-0.031	-0.012	-0.027**
	(0.009)	(0.020)	(0.019)	(0.012)
Quadratic Time Trends	Yes	Yes	Yes	Yes
Fixed-effects				
Account	Yes	Yes	Yes	Yes
Fit statistics				
Observations	294,817	426,457	343,793	237,588
R^2	0.413	0.462	0.322	0.311
Within R^2	0.005	0.247	0.002	0.002

Clustered (Account) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

members of the academic community on Twitter (i.e., those with verified accounts) reduced their engagement with the platform more than the less visible accounts. However, these conclusions hold only for the more reputational behaviors of authoring original content (column 1) or quoting someone else's tweet (column 4). Conversely, we observe no significant difference between verified and unverified accounts when it comes to retweeting or replying to tweets, although both of these are in the same direction. If we were to imagine a range of behaviors on Twitter from most reputational to least, it would be arrayed from original posts, quote tweets, replies, retweets, and finally favorites (which we cannot measure in our data). As such, the patterns we document here seem to indicate that, while all academics reduced their engagement with the platform in our sample, this reduction was significantly more pronounced among those whose visibility carried larger implications for their reputation, and among behaviors that are more reputational in nature.

6 Discussion

We document suggestive evidence of academics across four social science disciplines disengaging with Twitter just as Elon Musk took over as CEO of the platform. We find that both the number of daily active accounts, and the volume of content that was posted, declined significantly follow Musk's takeover. Using Bayesian change point analysis, we further conclude that the real drop did not occur until around November 19th, 2022, three weeks after Musk's official tenure began and corresponding to his poll-driven decision to reinstate former president Donald Trump's account. Finally, we show that the decline in engagement was especially strong among "verified" academic accounts, and especially in the behaviors that create original content: writing original tweets or quote-tweeting existing tweets.

We argue that these patterns are consistent with a simple story in which Musk's ownership brought about changes in the social network that reduced the utility obtained by academics using it. In addition, the significantly greater declines in engagement among more popular accounts is consistent with a reputational driver wherein scholars with larger profiles may have been more worried about being seen as tacitly endorsing the platform. We caution that we have not been able to experimentally manipulate anything in our study, including this potential reputational mechanism, and call on future research to more carefully evaluate our explanations for why academics reduced their use of Twitter after Musk's takeover. Given the lack of a dedicated API to collect data (one of many of the onerous changes brought about by Elon Musk, which we return to below), future work may need to rely on smaller-n approaches via survey experiments or qualitative studies.

An obvious question is "who cares?"[2] We argue that those who value Twitter specifically and public-facing academic research more generally should care about these results. Although, as summarized above, the net effect of Twitter on political and intellectual life will

[2]Indeed, this is a question that has been posed sarcastically by a resurgent anti-academic Twitter community who celebrated Musk's takeover and the "blue check evacuation."

continue to be debated, there are many parties who will suffer from the decline of academic Twitter.

Those who believe that Twitter generated fruitful discussion about working papers should care about the evidence we present in this letter. Those who became aware of junior scholars' research that they would otherwise not have should care, as should those who were introduced to academics from underrepresented backgrounds in the academy. Furthermore, the non-academics who were introduced to, and followed the scholars in our dataset should also care, as should anyone who values the principle of independent research reaching a wider, non-academic audience.

Nevertheless, we acknowledge that there are many scholars who never use Twitter in the first place, and that the institutions that govern professional advancement remain mostly insulated from new channels of networking and influence such as those embodied by Twitter. Furthermore, it is possible that the decline we document will be temporary, and that those who indeed were pushed off the platform by Musk's ownership may one day return. As with any research focused on online social networks, we expect that the conclusions contained in this manuscript are fleeting, although we argue that the dynamics governing the patterns we document are durable.

This is the part of the paper where we would normally punt the outstanding questions to "future research." Except in the case of Twitter at the time of writing, our capacity to conduct the type of analyses contained in this manuscript has been destroyed. Free access—academic or otherwise—to Twitter's APIs is no longer possible.

Even the skeptic who might not worry too much about the decline of social scientists on Twitter should be concerned about the decline in transparency that has accompanied Musk's leadership caused by the removal of API access. This decline is mirrored across a number of prominent online information environments, hamstringing the ability non-corporate actors

like journalists and social scientists to study how these ecosystems contribute to society writ large.

Note that this API access removal was not a total shock, as Freelon (2018) predicted the coming "Post-API Era" more than five years ago. As Munger (2023) argues, this kind of "temporal validity" problem is a first-order concern for research on fast-moving systems like social media. In light of this, social scientists cannot take stability for granted, but should rather embrace more epistemically humble research methods like quantitative description (Munger, Guess and Hargittai, 2021).

References

Aminikhanghahi, Samaneh and Diane J Cook. 2017. "A survey of methods for time series change point detection." *Knowledge and information systems* 51(2):339–367.

Bail, Chris. 2021. *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*. Princeton University Press.

Bandula-Irwin, Tanya and Veronica Kitchen. 2022. "The Twitter Conference as a New Medium of Scholarly Communication (and How to Host One)." *PS: Political Science & Politics* 55(4):778–782.

Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J Egan, Richard Bonneau, John T Jost and Joshua A Tucker. 2019. "Who leads? Who follows? Measuring issue attention and agenda setting by legislators and the mass public using social media data." *American Political Science Review* 113(4):883–901.

Bisbee, James, Jennifer Larson and Kevin Munger. 2022. "# polisci Twitter: A Descriptive Analysis of How Political Scientists Use Twitter in 2019." *Perspectives on Politics* 20(3):879–900.

Dang, Sheila. 2023. "Exclusive: Elon Musk's X restructuring curtails disinformation research, spurs legal fears." *Reuters.https://www.reuters.com/technology/elon-musks-x-restructuring-curtails-disinformation-research-spurs-legal-fears-2023-11-06/* .

Freelon, Deen. 2018. "Computational research in the post-API age." *Political Communication* 35(4):665–668.

Gerbaudo, Paolo. 2019. *The digital party: political organisation and online democracy*. Pluto Books.

Grossmann, Matt. 2021. *How social science got better: Overcoming bias with more evidence, diversity, and self-reflection*. Oxford University Press.

Jester, Natalie. 2022. "Twitter as higher-education community of practice: A political science perspective." *PS: Political Science & Politics* 55(4):815–819.

Jünger, Jakob and Birte Fähnrich. 2020. "Does really no one care? Analyzing the public engagement of communication scientists on Twitter." *New Media & Society* 22(3):387–408.

Ke, Qing, Yong-Yeol Ahn and Cassidy R Sugimoto. 2017. "A systematic identification and analysis of scientists on Twitter." *PLoS one* 12(4):e0175368.

Kim, Eunji and Shawn Patterson. 2022. "The pandemic and gender inequality in academia." *PS: Political Science & Politics* 55(1):109–116.

Klar, Samara, Yanna Krupnikov, John Barry Ryan, Kathleen Searles and Yotam Shmargad. 2020. "Using social media to promote academic research: Identifying the benefits of twitter for sharing academic work." *PloS one* 15(4):e0229446.

Munger, Kevin. 2023. "Temporal validity as meta-science." *Research & Politics* 10(3):20531680231187271.

Munger, Kevin, Andrew M Guess and Eszter Hargittai. 2021. "Quantitative description of digital media: A modest proposal to disrupt academic publishing." *Journal of Quantitative Description* (1):1–13.

Pew Research Center. 2018. *Internet/Broadband Fact Sheet*. Pew.

Rutenberg, Jim and Kate Conger. 2024. "Elon Musk Is Spreading Election Misinformation, but X's Fact Checkers Are Long Gone." *The New York Times.https://www.nytimes.com/2024/01/25/us/politics/elon-musk-election-misinformation-x-twitter* .

Searles, Kathleen and Yanna Krupnikov. 2018. "How Not to Get Ratioed and Other Advice for the Savvy Graduate Mentor." *Political Communication* 35(4):669–673.

Sweet-Cushman, Jennie. 2019. "Social media learning as a pedagogical tool: Twitter and engagement in civic dialogue and public policy." *PS: Political Science & Politics* 52(4):763–770.

Tufekci, Zeynep. 2017. *Twitter and tear gas: The power and fragility of networked protest*. Yale University Press.

Usher, Nikki, Jesse Holcomb and Justin Littman. 2018. "Twitter makes it worse: Political journalists, gendered echo chambers, and the amplification of gender bias." *The international journal of press/politics* 23(3):324–344.

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:

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

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.