

The Partisan Effects of Information Control Policies: Evidence from the X Ban in Brazil

Tiago Ventura¹, Christopher Barrie^{2,4}, Margaret E. Roberts³
Christopher Schwarz⁴, Joshua A. Tucker^{4,5}

¹Georgetown University,

²Department of Sociology, New York University

³University of California, San Diego

⁴Center for Social Media and Politics, New York University

⁵Politics Department, New York University

October 21, 2026

Abstract

What happens to information environments when partisan conflict over information control policies emerges in democracies? While a large literature examines information control under authoritarianism, democratic governments have increasingly intervened in major online platforms. We study a prominent case: Brazil’s 2024 national ban on the social media platform X. Using an event-study design, we estimate the causal effects of the ban and examine how partisan identity shaped responses. Drawing on a large sample of politically engaged users and ideal-point estimates of ideology, we find strong partisan asymmetries. Conservative users not aligned with the government were more likely to circumvent the ban, and right-leaning news domains became markedly more prevalent on the platform. We describe this dynamic as a “sorting ratchet”: the ban segmented the digital public sphere along partisan lines, with effects that persisted even after restrictions were lifted. When content moderation enforcement becomes a partisan issue, it may deepen polarization and reshape information environments in democratic settings.

Word Count: 6627 (without references)

1 Introduction

With online misinformation often seen as a major threat to society ([World Economic Forum, 2023](#); [Newman et al., 2024](#); [Budak et al., 2024](#)), governments increasingly face the challenge of moderating social media platforms and regulating which content can freely circulate online. Often, when governments enforce these policies, platforms push back, and conflicts escalate. Examples of such conflicts abound. In Europe, platforms have launched legal challenges against European Union fines under the Digital Services Act. When the Brazilian Congress was about to vote on a bill to moderate misinformation on online platforms, Google placed a link on its homepage urging users to lobby their representatives against it. In the context of the TikTok shutdown in the United States, the platform sent push notifications asking users to act against the proposal. These conflicts are rarely confined to government and platforms. In democracies, they often become partisan, and once they do, enforcing content moderation may become another issue along which citizens divide.

Scholars have studied partisan responses to content moderation policies in democratic countries. However, the vast majority of this scholarship deals with citizens’ preferences and how partisan alignment determines their policy views ([Appel, Pan and Roberts, 2023](#); [Kozyreva et al., 2023](#)). We know relatively little about how these policies affect users’ behavior and how polarized conflicts over content moderation policies shape the online informational ecosystem. Much of the existing scholarship on citizens’ reactions to information control policies focuses on authoritarian contexts ([Roberts, 2018](#); [Hobbs and Roberts, 2018](#); [King, Pan and Roberts, 2014](#); [Boxell and Steinert-Threlkeld, 2022](#); [Pan and Siegel, 2020](#)). In authoritarian regimes, the main cleavage runs between state and citizens; in democracies, where governments are showing increasing willingness to intervene in governance and the accessibility of platforms, the cleavage that increasingly matters the most is voters’ partisan identity ([Larreguy and Raffler, 2025](#)). As a consequence, this gap in the scholarship leaves open how content moderation interventions activate partisan responses, shape users’ online

behavior, and reshape platforms’ online ecosystems.

In this article, we consider one of the largest conflicts over content moderation policies between democratic governments and social media platforms: the country-wide ban of X/Twitter in Brazil in 2024. In Brazil, the political conflict over content moderation policies went further than in any other democracy in the world. After X (formerly Twitter) refused to comply with Supreme Court orders to remove accounts under investigation for spreading misinformation about elections and mobilizing the January 8, 2023 attacks on Brazil’s Congress, the platform closed its office in the country and withdrew its legal representative. On August 30, 2024, the Court responded with a nationwide ban on X that lasted for over a month. The Supreme Court’s ban on X had immediate and far-reaching consequences for Brazil’s digital landscape, disrupting the online routines of approximately 40 million people.¹ The ban capped a two-year political fight over content moderation, and Brazilians read it through partisan lenses. For the left, the Court was enforcing the law against a platform that shielded coup supporters. For the right, the Court was censoring the opposition.

We ask: how do partisan dynamics shape compliance² with state-imposed platform bans, and what are the consequences for the information environment in a polarized democracy? We first focus on partisan compliance with the X ban in Brazil during the period in which the ban was in place. Then, we examine the long-term effects of this ban on user behavior by analyzing changes on the platform even months after the ban ends. Some users were able to stay on the platform through the use of a VPN or comparable technology to hide their digital location; others chose to leave altogether. Using ideal-point estimates of user ideology, we identify which users remained active on the platform (i.e., did not comply), which stopped using it (i.e., complied), and their level of activity after the ban. This makes it possible to

¹Market share information estimates that roughly 20% of the Brazilian population reports using the platform at least once a month (CBS News, 2024).

²We adopt the language of experimental research to define compliance. We considered those who ignored the court ruling and stayed on the platform as the “non-compliers” with the treatment of the ban, while those who stayed off the platform were the “compliers” with the treatment of the ban.

estimate changes in the composition of news content on the platform before, during, and after the ban.

We theorize that platform bans can create a “sorting ratchet,” separating partisans into separate digital spaces. Building on the literature on information control in authoritarian regimes ([Jansen and Martin, 2015](#); [Nabi, 2014](#); [Pan and Siegel, 2020](#); [Miller, 2022](#)), we define this as a form of backlash (“Streisand effects”) to government efforts to control the information environment that is mobilized and intertwined with partisan cleavages. For example, when faced with a social media ban and a choice to exit the banned platform or circumvent it, social media users have different reputational and informational incentives that correlate with ideology. Consequently, co-partisans tend to make similar decisions either to exit or continue to use the platform, further balkanizing the digital space. We call this a ratchet because the effect does not fully reverse itself, even if or when the platform is unblocked. These effects persist, albeit at a smaller magnitude, after the ban ends. Even though social media bans reflect a growing trend in polarized democracies — Ukraine, for example, blocked Russian website VKontakte in order to reduce the spread of Russian propaganda in 2017 ([Golovchenko, 2022](#)) and India and Nepal have issued large-scale blocks of social media accounts and websites during times of political turmoil ([Digital Rights Monitor, 2025](#); [Committee to Protect Journalists, 2025](#)) — we argue this dynamic is not necessarily exclusive to bans, but it shows us broader behavioral mechanisms for partisan effects of information control policies in democracies.

We find that the ban created a partisan shock, and users with distinct partisan leanings responded differently, thereby generating shifts in the information environment and in platform engagement. During the ban, Conservatives—who were more critical of the Supreme Court’s decision—were more likely to circumvent the ban to maintain their presence on X. Liberals, on the other hand, tended to go silent on the platform during this period. This effect is large both for news circulation on the platform and for overall content creation among this

large set of politically engaged users. Furthermore, we show that these effects, although small in magnitude, persist even months after the judiciary lifted the ban in Brazil. We show that in the long run (post-ban), more conservative users increased their activity levels on X compared to the pre-ban period, and were more likely to remain active on the platform. As a consequence, the composition of political content moved further to the right, and levels of activity and engagement on the platform became more concentrated among conservative users.

Our work makes three contributions. First, we use the X ban in Brazil to illustrate how conflicts over content moderation can produce partisan responses, how partisan identity shaped users' behavior, their decisions to stay or exit the platform, and the short- and long-run effects of compositional changes on the fragmentation of digital space. In doing so, we add to prior work on the preference gap between partisans for information control policies in the online ecosystem ([Appel, Pan and Roberts, 2023](#); [Kozyreva et al., 2023](#)). Second, our findings theoretically extend research on censorship from autocracies ([Roberts, 2018](#); [Jansen and Martin, 2015](#); [Pan and Siegel, 2020](#)) to polarized democracies, and show that citizens' reaction takes a different form when the underlying cleavage is partisan: it segments the public sphere rather than uniting it against the state. Third, we add to work on partisan compliance with public policy ([Goldstein and Wiedemann, 2022](#); [Krupenkin, 2021](#)) and on the partisan motives behind platform regulation ([Esberg, 2020](#)) by documenting the informational consequences of partisan non-compliance at scale.

Bans are only the extreme end of a distribution of enforcement actions that democratic governments now take against platforms, from takedown orders to fines to threats of suspension. However, this critical event allows us to illustrate how conflict escalation around information control policies in democracies activates partisan mechanisms. We build a theoretical framework around social media bans as an illustration of politicized enforcement of information control policies, but our framework is not restricted to these escalated episodes. Our work

describes how partisan identity affects users’ decisions to stay on or exit a platform, how government interventions can reshape the platform’s information environment in ways its proponents did not intend, and how the reshaping can outlast the intervention itself.

The article proceeds as follows. First, we describe the conflict between X and the Brazilian Supreme Court and develop our theoretical framework, drawing on research on censorship, circumvention, and partisan compliance. Second, we introduce the data and research design, with particular attention to the estimation of ideal points for users and news domains in Brazil. Third, we validate our ideology measures. Fourth, we present the effects of the ban on platform activity, on partisan responses in the short and long run, and on the composition of X’s news environment. Fifth, we use original survey data to test for alternative mechanisms and show evidence of VPN usage associated with the X ban in Brazil. We conclude with implications for content moderation policy in polarized democracies.

2 Content Moderation, Social Media Bans in Democracies and the X Suspension in Brazil

The judicial and political conflict that culminated in the Brazilian Supreme Court imposing a country-wide ban on X started in the aftermath of the 2022 Brazilian election. Following a polarizing and highly competitive race between incumbent President Jair Bolsonaro and former two-time President Luiz Inácio Lula da Silva, right-leaning supporters of the incumbent launched a grassroots movement to contest the election results, which included lockouts from unionized truck drivers, public demonstrations, civilian encampments in front of Brazilian military headquarters, and a strong digital campaign on social media. On January 8, 2023, after the presidential inauguration, right-leaning activists stormed Brazil’s Congress, the Supreme Court, and presidential offices, scenes resembling the January 6 attacks in the United States ([Picheta, 2023](#)).

As a response to the January 8 attacks, Brazil’s Supreme Court initiated judicial inquiries to investigate groups involved in the plot. The Supreme Court requested that accounts

allegedly involved in the attacks or accused of spreading misinformation to be removed from the platforms. While other companies, such as Meta and Google, complied, X escalated the conflict, denying the judicial requests (Santos, 2024). In the second half of August, X's representative in Brazil resigned, and the platform leadership refused to name a new legal representative. The Brazilian Supreme Court considered this move an attempt to evade the legal system and ordered the platform's ban on August 30, 2024. The ban lasted for about a month, and ultimately ended on October 8 once X complied with the demands of the Brazilian Supreme Court (Associated Press, 2024).

2.1 Relevant Literature

There is relatively little work on information control policies, including full country-wide bans of social media platforms, in democracies. The literature on censorship in autocracies has shown that social media censorship creates costs for consumers of information, either by deterring engagement through fear of repercussions or by creating practical barriers to access (or “friction”, as in Roberts (2018)). In the case of Brazil, the ban imposed both types of costs on users – circumventing the ban was inconvenient, requiring users to download a VPN to access the website outside of Brazil, and the government threatened fines for accessing X. Whether users continue to circumvent these policies despite these costs depends on a variety of factors. First, users have to be sufficiently technologically savvy to continue to access a blocked website – knowing how to download, install, and work a VPN requires a certain amount of digital literacy. Second, users have to be motivated enough to circumvent it – the benefit of continuing to access the website must outweigh the additional costs that access now imposes (Chen and Yang, 2019). In general, empirical evidence has shown that overall use of blocked websites declines after a block, given that for many people, the benefits of circumvention do not exceed the cost (Hobbs and Roberts, 2018).

The type of information blocked is important for determining whether users circumvent the block. Entertainment websites seem to be more difficult for users to exit in comparison

to websites that provide purely news (Zuckerman, 2015; Hobbs and Roberts, 2018). Social media websites, therefore, could be stickier than pure news websites, which may have more ready substitutes for users who are typically not that interested in politics. Unlike news websites – or even entertainment websites – social media websites also hold the promise of a “social” connection by linking people to their preferred networks of other users, something that may also be hard to replace.

In some cases, censorship can provoke circumvention—a “Streisand effect”—as users seek out the restricted information out of protest or renewed curiosity (Jansen and Martin, 2015; Nabi, 2014; Pan and Siegel, 2020; Miller, 2022). Censorship in these cases can even create collective action, spurring unrest in opposition to government policy (Boxell and Steinert-Threlkeld, 2022; Hassanpour, 2014). Censorship circumvention can also be more likely during political crises, when use of the media tends to be heightened (Chang et al., 2022). In the context of authoritarian regimes, evidence of backlash is often seen as a conflict between state and society, with citizens reacting against information control policies in a somewhat uniform manner. In democracies, as these conflicts escalate and governments decide to moderate certain types of content or promote a full platform ban, citizens’ reactions emerge and can potentially reinforce political cleavages (Kozyreva et al., 2023; Appel, Pan and Roberts, 2023). In Brazil, the X ban was highly politicized, aligned with a political crisis, and aimed at blocking supporters of one political party/candidate, who were largely involved in misinformation sharing (Mello, 2020). Thus, unlike in authoritarian contexts, the incentives of backlash against a ban, or other information control policies, would not be uniform across the political spectrum, and any type of citizens’ reaction emerges across partisan lines.

2.2 Theoretical Framework

Drawing these threads together, we suggest that compliance with a platform ban – in either a democratic or autocratic regime – can be theorized as a choice between **exit (compliance)**

and **circumvention (non-compliance)**. Users weigh the higher costs of staying – technical friction, legal risk, and hassle – against the benefits of continued participation. Those benefits are not only about getting information: for politically engaged users, the platform is also a place for *voice*: coordination, agenda-setting, and influence. When a ban becomes politicized – i.e., associated with being supported by and/or opposed by particular partisan actors – partisanship predictably changes this cost-benefit calculation in ways that differ across partisan groups. In a highly politicized country with user choice in which platform to use, this can create a **sorting ratchet**, where the ban filters out users of a particular partisanship, while retaining users with the opposite ideology.

Partisan incentives for deciding whether to circumvent the ban can be either reputational (expressive) or informational (instrumental). Reputationally, exiting versus circumventing sends a signal: complying can communicate respect for rules and institutional legitimacy, while circumventing can communicate defiance or loyalty, depending on what one’s audience rewards. For left-leaning users in the Brazilian setting, for example, non-compliance may carry reputational and normative downsides – the risk of looking like one is aligning with the opposite partisan “camp” – so exit can be attractive even when circumvention is possible.

Instrumentally, exit can make strategic sense if the relevant audience is shrinking or if the platform becomes “owned” by the other side during the ban. Users consider the value of the information environment and the costs of switching: if co-partisans leave, the remaining content and social cues become less congenial and less useful, lowering the private payoff to staying. A credible alternative outlet (e.g., Bluesky) strengthens these mechanisms: it reduces what users give up by exiting (they can comply without going silent), makes coordinated exit easier to sustain, and can create feedback effects in which partisan exit reshapes who remains on the platform.

Both reputational and instrumental logics predict that left-leaning Brazilians ought to react differently to the ban of X than right-leaning Brazilians, albeit for different reasons. This

differential response creates a “sorting ratchet,” polarizing the banned website. Moreover, as habits form during the ban, we expect this effect to persist even after the ban is lifted.

Our theoretical framework centers around platform bans. However, the activation of partisan responses does not depend on these particular events. The sorting ratchet as an outcome emerges once information control policies are enforced and become politicized, and reputational and informational incentives to stay in a particular social media platform change. For example, court orders to remove accounts, fines, threats of suspension, and deplatforming (Mitts, 2025) all can potentially turn the platform into a political object. When platforms comply with these orders, unlike X in Brazil, remaining active on a platform that bans ideologically misaligned users (or content) signals support for these policies. At the same time, with these users (or content) banned or removed, information one considers harmful becomes scarce, and information gains increase. Compared to country-wide bans, these more modest actions should slowly produce similar partisan sorting outcomes at a smaller, but likely cumulative, magnitude. The Brazilian case gives a first indication of this. As we show below, the partisan gap in activity opens in mid-August, when the conflict between X and the Court had escalated but no block was in place, and peaks only once the ban is enforced.

3 Data and Methods

To model changes in the partisan composition of the informational environment of X during the nationwide ban, we work primarily with web links (URLs) from Brazilian news domains shared by users on the platform. News sharing represents a substantive part of how content propagates and reaches users in graph-based social media platforms (Aruguete, Calvo and Ventura, 2023; Blanchett Neheli, 2018; Hanusch and Nölleke, 2019). An emerging literature in political communication has considered different aspects of news sharing, for example, what shapes users’ decisions to share news online (Aruguete, Calvo and Ventura, 2023), which features make news more easily sharable online, and how media outlets and journalists react to news sharing online (Blanchett Neheli, 2018; Hanusch and Nölleke, 2019). Domain sharing

enables researchers to comprehend the news environment on specific social media platforms, as these links represent the primary means by which users introduce news information to the platform and share it with their peers. Critically, domain sharing represents a valuable source of behavioral data about both the users’ decisions to propagate links and how media organizations are activated in the informational environments of the platform (Eady et al., 2025; Aruguete, Calvo and Ventura, 2021). In what follows, we discuss our data sources, present our modeling strategy to estimate the ideology of news domains in Brazil, and validate our measures for the Brazilian case.

3.1 Estimating Ideological Leaning of Users and Political News in Brazil

Pre-ban News Sharing Dataset: To estimate the ideological position of users and political news in Brazil before the X ban, we collected 90 days of data, starting before August 30, via the now retired X Decahose API. This endpoint gave us access to 10% of all public tweets, selected via a real-time sampling algorithm. We filtered these tweets to only those in Portuguese, the primary language spoken in Brazil, and containing a URL. In total we collected around 14 million tweets.

News Sharing Ideology Score Estimation: Political science research has, over the years, built an extensive suite of models using behavioral data to estimate the ideology of politicians and voters, without relying on survey data. For example, similar methods have used roll-call voting (Lewis et al., 2019), party manifestos (Benoit, Laver and Garry, 2003; Rheault and Cochrane, 2020), congressional speeches (Izumi and Medeiros, 2021; Lauderdale and Herzog, 2014), and campaign contributions (Bonica, 2014) to estimate the ideology of politicians in the U.S. and abroad. Subsequent work showed how we can conduct similar ideological scaling routines using the follow networks of users and politicians (Barberá, 2015), news sharing (Aruguete, Calvo and Ventura, 2021; Eady et al., 2025), and language features of social media content (Temporão et al., 2018). Our paper uses a version of these methods

introduced by [Eady et al. \(2025\)](#), which relies on news sharing behavior as inputs. The model uses a matrix $C_{u,m}$, where u_i represents every user, m_j represents every unique news domain, and every cell in C is the sharing count of domain m_j by user u_i in the *Pre-ban News Sharing Data*. With a matrix of thousands of rows and hundreds of columns, estimating the scaling model using a Bayesian approach as in [Eady et al. \(2025\)](#) becomes computationally expensive. Therefore, we opt for a more computationally efficient approach using correspondence analysis ([Greenacre, 2010](#)) on the matrix $\mathbf{C}$. This approach has been shown to yield similar results compared to Bayesian estimation ([Macdonald et al., 2025](#); [Barberá et al., 2015](#)).

To estimate the ideal points, we must select a set of users u and media organizations m that allows for the identification of the *ideological* latent multidimensional space. We address this estimation challenge by restricting the counts matrix to politically focused domains and politically engaged users. We start with all unique domains shared more than 50 times, then retain only Brazilian political and news domains, identified using a zero-shot GPT-4o classification prompt. The model labels domains belonging to media organizations and to social or political organizations, including parties, social movements, civil society groups, and government entities. We then manually review these labels, removing non-political domains mistakenly included, such as gaming, sports, and celebrity news sites. We also exclude political news domains shared by fewer than 10 users. The full prompt and final domain list are provided in the Supplementary Material (Sections [A](#) and [B](#)).

On the user side of the input, we build a set of politically engaged ordinary users ([Eady et al., 2025](#); [Barberá, 2015](#)). We define these as users who shared at least five different political news domains in the *Pre-ban News Sharing Dataset*. This filtering decision is critical when working with ordinary users, and not just political actors, and allows us to restrict our analysis to users who a) provide enough information, b) were active on X over the pre-ban period, and c) engaged consistently with news content on X.

After these filtering decisions, our final count matrix consists of 9,061 users and 242

political news domains. We use this final matrix to compute ideology estimates for users and domains. To identify the model, we use the simple solution suggested in [Jackman \(2001\)](#) of multiplying the scores by -1 and standardizing them using a z-score normalization, so that the distribution is centered at zero; high values indicate right-leaning ideology, and low values, left-leaning.³

3.2 Measuring the Effects of the Ban

Politically Engaged Users Dataset: To study the effects of the X ban, we assembled a panel of politically engaged ordinary users observed before and after the ban. We began with the Pre-ban News Sharing Dataset and selected users who had shared at least five distinct political news domains in that dataset. We then scraped all available timelines for these users from June 1 to December 31, 2024, using public *Nitter* instances.⁴ We successfully collected timelines for 7,471 users, yielding approximately 6.7 million tweets and 430,000 shares of political news.⁵ Because prior research shows that a small fraction of users generates most social media content ([Grinberg et al., 2019](#); [Baribi-Bartov, Swire-Thompson and Grinberg, 2024](#)), this non-random sample likely captures much of the platform’s political information production and circulation during the period.

³The sharing space on X had a higher presence of right-leaning users and domains, and the smaller subset of left-leaning users disproportionately shared from more distinctive and less frequent news domains. As a consequence, the CA estimation projects these users far from the center of the distribution. This produces a highly asymmetric score distribution, with a long tail for the left-leaning users, that reflects substantive heterogeneity in news-sharing behavior. To improve the interpretability of this scale, we apply a groupwise recentering in the user distribution: scores are split into left and right subsets centered at zero, standardized within each group, and then rejoined at the center point. This monotone transformation compresses extreme tails, spreads out the dense middle, and yields a more balanced display without altering the underlying rank information that CA provides. We provide extensive validation of the scores in the next section, and in SM section D.

⁴We scraped timelines, rather than extending the broader collection used for the Pre-ban News Sharing Dataset, because our team lost access to the Twitter Decahose. Scraping this fixed sample of active users was therefore necessary to study the ban’s longer-run effects on the X information environment.

⁵We attempted to collect timelines for all 9,061 users identified as relevant and used in the ideology estimation, but experienced substantial attrition. Missing accounts reflect several causes: account deletion, handle changes without a numerical crosswalk ID, suspension, private settings, and, in at least one case, deletion of all tweets. Section G of the Supplementary Materials compares missing and observed users.


3.2.1 Causal Effects of the Ban: Event Study Model

To identify the causal effects of the ban, we use a Poisson event-study model (flexible difference-in-differences) to analyze the data collected from June 1st to December 31st, 2024. Using the first month of our dataset as a reference (June), we produce monthly estimates for partisan responses for both users' posting activity and overall news sharing. This allows us to capture the dynamics of the ban's effect on right-leaning user activity relative to left-leaning users in the months before escalation of the conflict that led to the ban. More formally, we use the following model:

$$\begin{aligned} y_{ij} &\sim \text{Poisson}(\lambda_{ij}) \\ \lambda_{ij} &= \exp \left(\alpha_i + \tau_j + \sum_{t=1}^6 \beta_t \cdot \text{Right-leaning}_i \cdot \text{Month}_t \right) \end{aligned} \tag{1}$$

where the outcome y_{ij} is a count variable parameterized with a Poisson distribution indicating the number of tweets posted by user i on day j , and Right-leaning_i is a binary variable indicating whether user i has a positive ideology score (ideology score > 0) compared to left-leaning and centrist users (ideology score ≤ 0). The parameters of interest, β_t , are month-specific interaction coefficients that capture the partisan effect (difference-in-differences) between right-leaning and left-leaning users in month t relative to the first month of our time series before the ban, which we set as a baseline. Lastly, we add user and day fixed effects in α_i and τ_j , respectively, which capture different rates of posting and sharing over time and by user, and errors are clustered at the user level. We estimate the same model to identify the partisan responses at the domain level, using the domain sharing per day as an outcome, and ideology scores at the domain level as the parameter of interest, and errors clustered at the domain level.

Figure 1: Top 50 Political News Domains Shared 90 days before the Ban



4 Validation

In Figure 1, we present the 50 most shared political news domains in the pre-ban period. The two most frequently shared domains are *revistaoeste.com* and *jornaldacidadeonline.com.br*. These are well-known right-leaning digital outlets in Brazil (Aruguete, Calvo and Ventura, 2021; Mont’Alverne et al., 2024). Their high prevalence in the data already indicates how the information environment on X, even before the ban, leaned more conservative. Other, more traditional media outlets are also widely shared in our data, such as *uol.com.br*, *globo.com*, *cnnbrasil.com.br*, and *estadao.com.br*. Official government portals (*gov.br*, *senado.leg.br*, and *votosdeputados.com.br*), small personal journalism blogs (*paulofigueiredoshow.com*) and local news organizations (*oliberal.com* and *opovo.com.br*) are also included as domains in the counts matrix.

We then validate the ideological estimation using well-known domains in the dataset. Figure 2 presents the ideological placement of the top 50 news domains in the Pre-ban News Sharing Dataset (see SM figure 10 for the total shares for each of these 50 domains). The ideology scores for the political news domains have strong face validity, with an overall structure between more left-leaning and right-leaning organizations consistent with what experts in the Brazilian political and media landscape might expect. For example, the most left organizations in the top 50 domains are well-defined left-leaning media organizations, such as the jornalggn.com.br, brasil247.com, revistaforum.com.br and iclnoticias.com.br. The same appears on the other side of the distribution, with well-known right-leaning political news domains all clustered together, such as atrombetanews.com.br, jornaldacidadeonline.com.br and vistapatria.com.br. In addition, more traditional mainstream media organizations, which are more widely shared due to their reputation and diverse audience, are clustered around zero, like globo.com, istoe.com.br, terra.com.br and uol.com.br.

To further validate the ideology scores, we compare the estimates with survey-based ideology estimates of news media organizations in Brazil. In the survey-based measure, media ideology estimates are calculated based on the self-reported ideological placement of each organization’s audience using a representative sample of Brazilian news consumers (Mont’Alverne et al., 2024). These estimates are available for twenty-four organizations, which we match with our behavioral sharing measure. Figure 3 presents the results. Overall, both measures align relatively well, with a high Pearson correlation of 0.85. The farthest off the regression line is *jornaldacidadeonline.com.br*, which has a more extreme position in the survey-based measure than in the sharing data. This is likely due to the fact that, by using the sharing data, our estimates incorporate cross-cutting sharing from distinct online communities, which can pull some more extreme organizations to the center, compared to measures based exclusively on the ideology of the organizations’ audience (Green et al., 2025).

Figure 2: Ideological Position for Top 50 Political News Domains Shared 90 days before the Ban




Figure 3: Validating Sharing Ideological Scores with Survey Measures



5 Results


We present results as follows. First, we show the effects of the country-wide ban on overall posting and news-sharing behavior on the platform. After establishing these general effects, we examine the role of partisan identity in users’ responses to the ban in the short and long run — months after the platform ban was lifted. Our results show that the ban affected both activity levels and user composition on the platform — right-leaning users were more likely to circumvent the ban in the short run, increasing their activity levels compared to left-leaning users who mostly left or went silent during the ban. Although smaller in magnitude than during the ban period, these effects persist for months after the ban is lifted. We conclude with additional descriptive statistics on downstream effects on X’s information space, focusing on how partisan responses affected the news environment and overall platform engagement.

5.1 The Effects of the Ban on Posting Frequency

Figure 4 tracks over time all tweets and political news shared by politically engaged users. We split the timeline for the analysis into three periods: before the ban, from June 1 to August 30; the ban period, indicated in gray in figure 4, from September 1 to October 8; and after the ban, from October 8 to December 31, 2024. Every dot in the graph indicates the daily average of tweets posted or news shared across the entire period.

Figure 4 demonstrates some important trends. First, the X ban substantially reduced overall levels of activity and news sharing on the platform, as evidenced by the strong discontinuity after August 30. Second, despite the successful enforcement, some users remained active on the platform during the country-wide ban. For example, the politically engaged users posted a daily average of 37,207 tweets and 2,587 political news shares per day across all users in the months before the ban. During the ban, our dataset registers a reduction to 11,305 and 395 daily averages, summing all users’ regular tweet activity and news sharing, respectively. These are orders of magnitude fewer counts per day. Considering our sample comes from active users, whom we would expect to have higher incentives to

Figure 4: Overall Tweet Counts and News Sharing During the Ban



stay on the platform, the average effect of the ban in the whole population is likely more substantial. Third, we also observe an upward trend before the ban, both in tweet counts and in the political news share. This spike seems to start in mid-August, which coincides with the escalation of the conflict between the Brazilian Judiciary and the X leadership.⁶ Lastly, after the ban is lifted, user activity on the platform returns to the pre-ban level. In other words, despite the country-wide ban, part of the politically engaged users in the long run slowly returned to the platform, producing as much regular content (daily tweets) and sharing similar volumes of political news as before the ban. Therefore, we do not observe anything similar to a “Streisand effect” in which the platform usage increases as a form of backlash, as described in the context of authoritarian regimes (Jansen and Martin, 2015; Nabi, 2014; Pan and Siegel, 2020; Miller, 2022).

5.2 Partisan Responses to the Ban

We now discuss how decisions to circumvent the ban were motivated by partisan considerations. We explore how these partisan responses, in both the short and long term, shaped the informational environment on X in two critical aspects: the content produced by users and the decision to stay on the platform. Figure 5 presents marginal effects from the Poisson

⁶On August 17, for example, the X leadership announced the closure of the platform’s office in Brazil and the removal of all the Brazilian staff from the country over allegations of threats from the Supreme Court. See <https://www.politico.com/news/2024/08/17/x-brazil-censorship-supreme-court-00174500>.


event-study model that identifies the changes in the daily posting and domain sharing between right-leaning and left-leaning users and domains. On the x-axis, we separate the data monthly, using the first month of our data as the baseline. The top plot illustrates differences in overall tweet activity, while the bottom plot displays differences in news sharing.

Figure 5 depicts clear causal evidence of partisan responses to the ban on the levels of activity on the platform. The ban increased by 5.8 times ($\beta = 1.76$, $\exp(\beta) = 5.8$, $z\text{-score}=12.6$) the posting rate of right-leaning users in comparison to the activity of left-leaning users. While, in absolute terms, the levels of activity on X decreased as a result of the ban, compared to left-leaning users, right-leaning users increased their activity at a substantially higher rate, and were by far the main partisan group to defy the ban. Similar differences are observed in news-sharing rates between right-leaning and left-leaning media organizations. However, results are less precise due to the smaller sample size of sharing rates for news domains.

Notably, the partisan response that spiked in the ban period starts before its formal enforcement on August 30. This is consistent with our theoretical framework: the sorting ratchet is driven by the platform’s politicization, which is strongest when the ban is enforced. It shows that while the ban produces a more significant change in the composition, the conflict around information control of accounts connected to spreading misinformation on X, and the platform reaction to the Court orders, slowly increased partisan sorting in the platform.

In the SM Section H, we show that these results are robust to narrower bandwidths (around the ban). First, Figure 22 shows that using 10-day and 30-day bandwidth before and after the ban, instead of the first month of our data, as a reference, depicts similar partisan responses. Second, we present results that hold across different model specifications, using an OLS model rather than a Poisson. Furthermore, we show similar differences across other engagement metrics, such as likes, quotes, retweets, and replies received by right-leaning

Figure 5: Event Study Estimates: Poisson Regression on Daily Rates of Posts and Shares on X



Note: Full regression tables on SM Tables 4 and 5

users compared to left-leaning users. Lastly, we show that the results are robust to distinct cutoffs for the partisan leaning of users.

5.3 Changes in the User Composition


In addition to changes in activity levels, we show that decisions to remain active on the platform during the ban and to return to it in the long run are conditional on the partisan leaning of users. The upper plot in Figure 6 presents the number of active users in the dataset of politically engaged users on the y-axis, and the users’ corresponding ideological estimates on the x-axis, divided into three time periods: before, during, and after the ban. We also present the median and mean ideological scores in the graphs, which we calculated using a weighted measure of the users’ ideological scores. The bottom part shows the total number of tweets produced by these users in the same time period.⁷

Across all periods, the most ideologically extreme users on both ends of the distribution produced a disproportionate share of content on the platform, consistent with prior work showing that ideological motivation drives social media engagement (Aruguete, Calvo and Ventura, 2023; Eady et al., 2025; Wojcieszak et al., 2022). At the same time, the distribution is right-skewed in every period, and this skew becomes more pronounced after the ban. In all three periods, the modal bin lies well to the right, and both the mean and median are positive, indicating that right-leaning users were already overrepresented on X before the ban. We also observe modest shifts in the aggregate distribution over time: the median ideology score is 3.15 before the ban, 3.29 during the ban, and 3.22 after the ban. These differences are even sharper in tweet volume. As Figure 6B shows, centrist and left-leaning users largely go silent during the ban, while right-leaning users continue to produce most of the platform’s content.


Figure 7 unpacks these results further, focusing on users who were active in the pre-ban

⁷In SM Figure 11, we plot the most shared domain across bins of the users’ scores, which provides the reader an intuitive understanding of partisan leaning across the users’ ideology bins.

Figure 6: The Effects of the Ban on Platform's User Composition



A: Counts of Active Users (Who Posted at Least One Tweet) Before, During, and After the Ban



B: Total Number of Tweets Posted Before, During, and After the Ban

period but left X during or after the ban. In the pre-ban period, we identified 7,291 users. During the ban, only 4,945 tweeted at least once, either via VPN or via other mechanisms circumventing the ban. And after the ban was lifted, we identified only 6,635 users who remained active on the platform. In other words, 2,346 users fully complied with the ban, and 656 never returned to the platform. Figure 7 presents a cumulative distribution function of the ideological distribution of the pre-ban user (N=7,291), the dropout sample during the ban (N=2,346), and the dropout sample after the ban (N=656) ⁸.

The ban produced a shock on the platform, affecting users across the entire ideological spectrum. However, these effects are stronger among left-leaning users who proportionally left X during the ban at a considerably higher rate. As depicted in Figure 7, in the pre-ban distribution, around 70% of the users have a positive ideological score, indicating a right-leaning alignment. In both dropout samples, the cumulative share of active users who left the platform drops to 50% of the distributions. This indicates that, although users across the entire distribution are affected by the platform ban, in the short run and the long run, users who are more supportive of the ban do not return to the platform.


5.4 Downstream Effects of the Ban on X’s Informational Environment

Having presented causally identified effects of the ban on partisan responses from users and changes in the user base composition, this section now introduces downstream effects in the aggregate informational environment makeup of the platform. Here, we discuss two aspects: changes in the news information environment on the platform and a higher concentration of multiple engagement metrics on right-leaning users on X as a consequence of the ban.

We start with Figure 8, which shows the ideological leaning of political news on the platform over time. Each dot represents the daily ideological leaning on X, calculated as the weighted average of the ideology scores of domains shared on the platform that day (weighted

⁸In the long run, this is not a negligible decline in the base of users. It represents close to 9% of the data we collected for active users, who arguably had higher costs to leave the platform. As we mentioned before, these effects are potentially larger on the overall population of X users

Figure 7: Users Dropout from X After the Ban



Note: Higher values in the CDF curve indicate a higher share of users leaving the platform in each time period


by the number of shares received daily by the domain). As in the validation results, higher values indicate more conservative leaning. Figure 8 depicts three main trends. First, even before the ban, the political news environment on X was slightly more right-leaning. Second, the ban produced a substantial shock in the news environment; the small volume of content produced while the ban was in action, likely due to right-leaning users circumventing the ban using VPNs, was politically charged and came almost exclusively from right-leaning outlets. Lastly, the partisan shock of the ban produced a small effect in the long run, as the news environment on the platform was more right-leaning even three months after the ban was lifted. These changes in the partisan leaning of the political content on the platform take place in parallel with an increase in the overall usage depicted in the post-ban period in Figure 4.

When considering all news shared by these authors, the median ideology score across the three time periods was 0.27 before the ban, 1.26 during the ban, and 0.30 in the months after the ban had been lifted. These differences indicate that the political news environment on X became approximately 1.05 standard deviations more right-leaning ($p\text{-value} < 0.01$) during the ban and 0.04 standard deviations ($p\text{-value} < 0.01$) more right-leaning in the three months following the ban⁹. To provide more substantive context for these scores particularly for the ban period, we compared them to scores from well-known news domains in our data. Before the ban, the political news environment was equivalent to <https://www.metropoles.com> (score = 0.3). During the ban, the news environment would be equivalent to content coming from right-leaning extremist portals, such as jornaldacidadeonline.com.br (score = 1.29). In the supplementary materials, we replicate this analysis using ideology estimates for politically engaged users (SM Figure 13).

We discuss now how the production of content on X became more concentrated among right-leaning users as a consequence of the ban. Figure 9 presents the empirical cumulative

⁹P-values calculated via pairwise t-tests on every news share, comparing the three time periods.


Figure 8: Ideological Response to X Ban



Note: Every dot indicates the ideological score of the news domains weighted by the number of times these domains were shared for each day between June 1st and December 31st. Higher values indicate more right-leaning ideology from users and domains

distribution function of levels of engagement (likes, quotes, replies, and retweets) by politically engaged users on the y-axis, and their respective ideological scores on the x-axis. Across all measures, we see that the volume of engagement produced on the platform became more concentrated among right-leaning users during the ban, as expected, but remained so months after in the long run. For example, when we look at the share of likes received before the ban, tweets produced by right-leaning users received 73% of all likes in our dataset. This concentration jumps to approximately 90% of all likes during the ban, and remains at 80% during the period after the ban is lifted. Similar differences appear in other engagement metrics, such as retweets, quotes, and replies.

Figure 9: Concentration of engagement with content across users' ideological leaning



In the SM, Figure 15, we show similar effects when considering the engagement of news domains as an outcome, but with more modest differences for the news sharing engagement

concentration in the long-run. Additionally, SM Figure 14 shows that this concentration also holds when examining activity levels (tweets posted). These results are in line with the findings from the event-study model presented earlier.

5.5 Evidence of VPN Usage During the Ban: Testing for Alternative Paths

Our results so far show that users did circumvent the ban, and that this decision was mostly driven by partisan reasons, with right-leaning users staying active on X while left-leaning users went quiet. To circumvent the ban, users in Brazil needed to sign up for VPN services to alter their origin location. However, an alternative story to our argument is that VPN access correlates with other variables, particularly income, which is also associated with right-wing preferences in Brazil (Zanotti, Rama and Tanscheit, 2023; Layton et al., 2021; Rennó, Moreira and Tamaki, 2026). In this section, we provide evidence of VPN use in Brazil, its use in the context of the ban, and test this alternative hypothesis related to our core findings. To do so, we use an original survey of Brazilian internet users fielded in April 2026 ($n = 1,694$). The survey asks about VPN use: whether participants use a VPN, have used one in the past, have never used a VPN, or do not even know what VPN stands for. We model VPN usage using three VPN outcomes: uses a VPN now, has ever used a VPN (now or in the past), and used a VPN in the past but not now.

First, descriptively, out of 1,694 respondents, only 17.9% report currently using a VPN, and a further 24.0% report having used one in the past but not currently. Of the remaining participants, 35.4% have never used a VPN and 22.7% report not knowing what the tool is. In short, VPN usage in Brazil is relatively rare, with more than half of respondents reporting they have never used a VPN. Critically, past users outnumber active users of VPN in Brazil.

Next, we model predictors of usage with logistic models regressing the three VPN outcomes on a set of sociodemographic variables, including ideology and income, digital literacy, social media platform usage, and macro-region fixed effects.¹⁰ Table 1 presents the

¹⁰The models are estimated on the 1,427 respondents with complete covariates, and on 1,172 for the *used*

results. Some interesting descriptive patterns emerge. First, VPN use in Brazil is strongly associated with digital literacy across all three VPN outcomes (current users, ever users, and past-but-not-current users). Second, neither income nor ideological leaning (right vs center and left vs center) statistically correlates with any of the three VPN outcomes. Income, which is arguably higher among right-wing voters in Brazil, does not predict VPN use, and partisan affiliation on either side does not show a distinct baseline of VPN access. These null effects align with our argument that the ban pushed right-wing voters to circumvent the decision, that this group had no differential prior VPN access, and that income, a predictor of partisan leaning in Brazil, cannot explain these partisan effects. Lastly, we find suggestive evidence that temporary usage of a VPN (used in the past, not now) is associated with currently using X. This finding also illustrates the dynamics of circumvention: X users are more likely than users of other platforms to report having used a VPN in the past but not now.

6 Conclusion

This paper examines how a major social media ban in a democracy reshaped the information environment and how partisan dynamics conditioned those effects. We show that Brazil’s ban on X substantially reduced use of the platform, but compliance was sharply partisan, producing a social media “sorting ratchet.” During the ban, right-leaning users were far more likely to remain active, and they accounted for most of the content posted and news shared on the platform. Although users from across the ideological spectrum gradually returned after the ban was lifted, the effects did not fully reverse. Three months later, X remained more conservative than before the ban: right-leaning users were more active, posted more frequently, and captured a larger share of engagement.

More broadly, our findings show that conflicts over platform regulation in polarized democracies can generate durable partisan shifts in participation. A large literature has examined the formation of echo chambers within major platforms ([Sunstein, 2018](#); [Bakshy,](#)

a VPN in the past but not now outcome, which excludes current VPN users.

Table 1: Regression Table: Predictors of VPN Use in Brazil

	Uses VPN now	Uses VPN now + used in the past	Used VPN in past, not now
Left (vs. center)	−0.102 (0.180)	0.030 (0.149)	0.115 (0.181)
Right (vs. center)	−0.066 (0.175)	−0.001 (0.145)	0.118 (0.178)
Income Brackets	0.020 (0.087)	−0.025 (0.074)	−0.063 (0.091)
Education	0.077 (0.108)	0.069 (0.091)	0.022 (0.109)
Digital literacy	0.253*** (0.081)	0.569*** (0.069)	0.659*** (0.083)
Age	0.005 (0.006)	−0.011** (0.005)	−0.017*** (0.006)
Female	0.066 (0.150)	−0.387*** (0.121)	−0.550*** (0.150)
White	−0.208 (0.148)	−0.130 (0.122)	−0.011 (0.148)
X (Twitter) user	−0.009 (0.156)	0.199 (0.125)	0.338** (0.151)
Instagram user	−0.027 (0.384)	0.050 (0.315)	0.189 (0.389)
Facebook user	0.258 (0.224)	0.129 (0.170)	0.032 (0.198)
TikTok user	0.299* (0.160)	0.029 (0.126)	−0.096 (0.149)
YouTube user	−0.146 (0.338)	0.114 (0.276)	0.312 (0.356)
Reddit user	0.420** (0.188)	0.408** (0.169)	0.262 (0.193)
Num.Obs.	1427	1427	1172
Region FE	Yes	Yes	Yes

* p < 0.1, ** p < 0.05, *** p < 0.01

All models include IBGE macro-region fixed effects.

[Messing and Adamic, 2015](#); [Del Vicario et al., 2016](#); [Eady et al., 2019](#)). Our results suggest that platform bans and other forms of platform governance can accelerate a related process at the platform level, deepening fragmentation by shifting which partisans remain active in shared online spaces.

From a policy perspective, the Brazilian case highlights a central trade-off in democratic digital governance. Measures intended to curb misinformation may also reallocate control over the information environment if partisan groups respond differently to enforcement. Platform bans may therefore be difficult to sustain politically and may produce unintended downstream effects even when they are temporary.

Our study also has important limitations. First, our inferences are based on a sample of politically engaged users in Brazil’s X ecosystem—users who frequently shared news and likely account for a disproportionate share of political content on the platform ([Grinberg et al., 2019](#); [Baribi-Bartov, Swire-Thompson and Grinberg, 2024](#)). Because post-ban restrictions on API access prevented broader data collection, we cannot fully assess how well this group represents the wider user population. Second, our analysis is limited to behavior on X and cannot capture spillovers to other platforms. Future research should examine these multi-platform dynamics, including whether bans shift partisan users and political discussion across platforms ([Mitts, 2025](#); [Thomas and Wahedi, 2023](#)).

References

- Appel, Ruth E, Jennifer Pan and Margaret E Roberts. 2023. “Partisan conflict over content moderation is more than disagreement about facts.” Science Advances 9(44):eadg6799.
- Aruguete, Natalia, Ernesto Calvo and Tiago Ventura. 2021. “News sharing, gatekeeping, and polarization: a study of the# Bolsonaro election.” Digital journalism 9(1):1–23.
- Aruguete, Natalia, Ernesto Calvo and Tiago Ventura. 2023. “News by popular demand: Ideological congruence, issue salience, and media reputation in news sharing.” The International Journal of Press/Politics 28(3):558–579.
- Associated Press. 2024. “Musk’s X to Be Reinstated in Brazil after Complying with Supreme Court Demands.” NPR .
URL: <https://www.npr.org/2024/10/08/nx-s1-5146510/brazil-x-twitter-court-reinstated-elon-musk>
- Bakshy, Eytan, Solomon Messing and Lada A Adamic. 2015. “Exposure to ideologically diverse news and opinion on Facebook.” Science 348(6239):1130–1132.
- Barberá, Pablo. 2015. “Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data.” Political analysis 23(1):76–91.
- Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker and Richard Bonneau. 2015. “Tweeting from left to right: Is online political communication more than an echo chamber?” Psychological science 26(10):1531–1542.
- Baribi-Bartov, Sahar, Briony Swire-Thompson and Nir Grinberg. 2024. “Supersharers of fake news on Twitter.” Science 384(6699):979–982.
- Benoit, Kenneth, Michael Laver and John Garry. 2003. “Estimating policy positions from political texts: A comparison of sentence-by-sentence versus whole-text approaches.” American Journal of Political Science 47(2):319–331.

- Blanchett Neheli, Nicole. 2018. “News by numbers: The evolution of analytics in journalism.” Digital Journalism 6(8):1041–1051.
- Bonica, Adam. 2014. “Mapping the ideological marketplace.” American Journal of Political Science 58(2):367–386.
- Boxell, Levi and Zachary Steinert-Threlkeld. 2022. “Taxing dissent: The impact of a social media tax in Uganda.” World Development 158:105950.
- Budak, Ceren, Brendan Nyhan, David M Rothschild, Emily Thorson and Duncan J Watts. 2024. “Misunderstanding the harms of online misinformation.” Nature 630(8015):45–53.
- CBS News. 2024. “Brazil’s Supreme Court orders X platform suspended amid feud with Elon Musk.”. Accessed: 2025-08-22.
URL: https://www.cbsnews.com/news/brazil-x-platform-suspended-elon-musk-feud-judge-alexandre-de-moraes/?utm_source=chatgpt.com
- Chang, Keng-Chi, William R Hobbs, Margaret E Roberts and Zachary C Steinert-Threlkeld. 2022. “COVID-19 increased censorship circumvention and access to sensitive topics in China.” Proceedings of the National Academy of Sciences 119(4):e2102818119.
- Chen, Yuyu and David Y Yang. 2019. “The impact of media censorship: 1984 or brave new world?” American economic review 109(6):2294–2332.
- Committee to Protect Journalists. 2025. “Nepal orders ban on major social media platforms.”. Accessed: 2026-03-04.
URL: <https://cpj.org/2025/09/nepal-orders-ban-on-major-social-media-platforms/>
- Del Vicario, Michela, Gianna Vivaldo, Alessandro Bessi, Fabiana Zollo, Antonio Scala, Guido Caldarelli and Walter Quattrociocchi. 2016. “Echo chambers: Emotional contagion and group polarization on facebook.” Scientific reports 6(1):37825.

- Digital Rights Monitor. 2025. "INDIA: More Pakistani YouTube channels, Instagram accounts blocked." . Accessed: 2026-03-04.
- URL:** <https://digitalrightsmonitor.pk/india-more-pakistani-youtube-channels-instagram-accounts-blocked/>
- Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky and Joshua A Tucker. 2019. "How many people live in political bubbles on social media? Evidence from linked survey and Twitter data." Sage Open 9(1):2158244019832705.
- Eady, Gregory, Richard Bonneau, Joshua A Tucker and Jonathan Nagler. 2025. "News sharing on social media: mapping the ideology of news media, politicians, and the mass public." Political Analysis 33(2):73–90.
- Esberg, Jane. 2020. "Censorship as reward: Evidence from pop culture censorship in Chile." American Political Science Review 114(3):821–836.
- Goldstein, Daniel A.N. and Johannes Wiedemann. 2022. "Who Do You Trust? The Consequences of Partisanship and Trust for Public Responsiveness to COVID-19 Orders." Perspectives on Politics 20(2):412–438.
- Golovchenko, Yevgeniy. 2022. "Fighting propaganda with censorship: A study of the Ukrainian ban on Russian social media." The Journal of Politics 84(2):639–654.
- Green, Jon, Stefan McCabe, Sarah Shugars, Hanyu Chwe, Luke Horgan, Shuyang Cao and David Lazer. 2025. "Curation bubbles." American Political Science Review pp. 1–19.
- Greenacre, Michael J. 2010. "Correspondence analysis." Wiley Interdisciplinary Reviews: Computational Statistics 2(5):613–619.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson and David Lazer. 2019. "Fake news on Twitter during the 2016 US presidential election." Science 363(6425):374–378.

- Hanusch, Folker and Daniel Nölleke. 2019. “Journalistic homophily on social media: Exploring journalists’ interactions with each other on Twitter.” Digital Journalism 7(1):22–44.
- Hassanpour, Navid. 2014. “Media Disruption and Revolutionary Unrest: Evidence from Mubarak’s Quasi-Experiment.” Political Communication 31(1):1–24.
URL: <https://www.tandfonline.com/doi/abs/10.1080/10584609.2012.737439>
- Hobbs, William R and Margaret E Roberts. 2018. “How sudden censorship can increase access to information.” American Political Science Review 112(3):621–636.
- Izumi, Mauricio Y and Danilo B Medeiros. 2021. “Government and opposition in legislative speechmaking: using text-as-data to estimate Brazilian political parties’ policy positions.” Latin American Politics and Society 63(1):145–164.
- Jackman, Simon. 2001. “Multidimensional analysis of roll call data via Bayesian simulation: Identification, estimation, inference, and model checking.” Political Analysis 9(3):227–241.
- Jansen, Sue Curry and Brian Martin. 2015. “The Streisand Effect and Censorship Backfire.” International Journal of Communication 9:656–671.
URL: <https://ijoc.org/index.php/ijoc/article/view/2498>
- King, Gary, Jennifer Pan and Margaret E Roberts. 2014. “Reverse-engineering censorship in China: Randomized experimentation and participant observation.” Science 345(6199):1251722.
- Kozyreva, Anastasia, Stefan M Herzog, Stephan Lewandowsky, Ralph Hertwig, Philipp Lorenz-Spreen, Mark Leiser and Jason Reifler. 2023. “Resolving content moderation dilemmas between free speech and harmful misinformation.” Proceedings of the National Academy of Sciences 120(7):e2210666120.
- Krupenkin, Masha. 2021. “Does partisanship affect compliance with government recommendations?” Political behavior 43(1):451–472.

- Larreguy, Horacio and Pia J Raffler. 2025. "Accountability in developing democracies: The impact of the internet, social media, and polarization." Annual Review of Political Science 28(1):413–434.
- Lauderdale, Benjamin E. and Alexander Herzog. 2014. "Measuring political positions from legislative speech." Political Analysis 22(3):277–291.
- Layton, Matthew L., Amy Erica Smith, Mason W. Moseley and Mollie J. Cohen. 2021. "Demographic Polarization and the Rise of the Far Right: Brazil's 2018 Presidential Election." Research & Politics 8(1):1–7.
- Lewis, Jeffrey B, Keith Poole, Howard Rosenthal, Adam Boche, Aaron Rudkin and Luke Sonnet. 2019. "Voteview: Congressional roll-call votes database." See <https://voteview.com/> (accessed 27 July 2018) .
- Macdonald, Maggie, Megan A Brown, Joshua A Tucker and Jonathan Nagler. 2025. "To moderate, or not to moderate: Strategic domain sharing by congressional campaigns." Electoral Studies 95:102907.
- Mello, Patrícia Campos. 2020. A máquina do ódio: notas de uma repórter sobre fake news e violência digital. Companhia das Letras.
- Miller, Andrew Cesare. 2022. "# DictatorErdogan: How social media bans trigger backlash." Political Communication 39(6):801–825.
- Mitts, Tamar. 2025. Safe havens for hate: The challenge of moderating online extremism. Princeton University Press.
- Mont'Alverne, Camila, Amy Ross Arguedas, Sayan Banerjee, Benjamin Toff, Richard Fletcher and Rasmus Kleis Nielsen. 2024. "Exposure to Partisan News and Its Impact on Social Polarization and Vote Choice: Evidence From the 2022 Brazilian Elections." The International Journal of Press/Politics p. 19401612241292700.

- Nabi, Zubair. 2014. “Resistance censorship is futile.” First Monday 19(11).
URL: <https://firstmonday.org/ojs/index.php/fm/article/view/5525>
- Newman, Nic, Richard Fletcher, Craig T Robertson, A Ross Arguedas and Rasmus Kleis Nielsen. 2024. Reuters Institute Digital News Report 2024. Technical report.
URL: <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2024>
- Pan, Jennifer and Alexandra A Siegel. 2020. “How Saudi crackdowns fail to silence online dissent.” American Political Science Review 114(1):109–125.
- Picheta, Rob. 2023. “The violent attack on Brazil’s government was months in the making. Here’s what you need to know.” CNN, January 9:1–2.
- Rennó, Lucio, Thiago Moreira and Eduardo Ryo Tamaki. 2026. “Beyond Electoral Fortunes: The Consolidation of a Far-Right Alignment in Brazil.” Journal of Politics in Latin America 18(2):235–260.
- Rheault, Ludovic and Christopher Cochrane. 2020. “Word embeddings for the analysis of ideological placement in parliamentary corpora.” Political analysis 28(1):112–133.
- Roberts, Margaret. 2018. Censored: distraction and diversion inside China’s Great Firewall. Princeton University Press.
- Santos, SF. 2024. “Musk’s X banned in Brazil after disinformation row.” BBC News 31.
- Sunstein, Cass R. 2018. # Republic. In # Republic. Princeton university press.
- Temporão, Mickael, Corentin Vande Kerckhove, Clifton van Der Linden, Yannick Dufresne and Julien M Hendrickx. 2018. “Ideological scaling of social media users: a dynamic lexicon approach.” Political Analysis 26(4):457–473.
- Thomas, Daniel Robert and Laila A Wahedi. 2023. “Disrupting hate: The effect of deplat-

forming hate organizations on their online audience.” Proceedings of the National Academy of Sciences 120(24):e2214080120.

Wojcieszak, Magdalena, Andreu Casas, Xudong Yu, Jonathan Nagler and Joshua A Tucker. 2022. “Most users do not follow political elites on Twitter; those who do show overwhelming preferences for ideological congruity.” Science advances 8(39):eabn9418.

World Economic Forum. 2023. The Global Risks Report 2023. Technical Report 18 World Economic Forum.

URL: https://www3.weforum.org/docs/WEF_Global_Risks_Report_2023.pdf

Zanotti, Lisa, José Rama and Talita Tanscheit. 2023. “Assessing the Fourth Wave of the Populist Radical Right: Jair Bolsonaro’s Voters in Comparative Perspective.” Opinião Pública 29(1):1–23.

Zuckerman, Ethan. 2015. Cute cats to the rescue? Participatory media and political expression. In From Voice to Influence: Understanding Citizenship in a Digital Age, ed. Danielle Allen and Jennifer S. Light. Chicago: The University of Chicago Press pp. 131–154.

Acknowledgments: We gratefully acknowledge that the Center for Social Media and Politics at New York University is supported by funding from the John S. and James L. Knight Foundation, the Hewlett Foundation, Craig Newmark Philanthropies, the William and Flora Hewlett Foundation, and the Siegel Family Endowment. This research was also supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

The Partisan Effects of Social Media Bans

Supporting Materials (SM)

Contents

A Prompt structure for Classifying Political News Domains in Brazil	1
B List of Domains	1
C Validation: Domains Ideology Estimation	3
D Validation: Users' Groupwise Recentering	3
E Additional Results: Partisan Responses to the Ban	7
F Additional Results: The Effects of the Ban on Social Media Engagement	9
G Analysis of Uncollected Users	10
H Robustness Checks: Event Study Models	15
I Regression Tables	20

A Prompt structure for Classifying Political News Domains in Brazil

Listing 1: Prompt structure and parameters

```
# Prompt template
prompt_text <- paste0(
  "Please examine this particular domain and determine if it is a news website, ",
  "a government agency website, or a political/social movement website in Brazil.\n",
  "Here is the domain: '", domain, "'\n",
  "Start your responses with 'binary' using 1 if the domain matches one of these
  categories, ",
  "and 0 otherwise, followed by an 'explanation' describing your decision in one sentence.
  "
)

# Function parameter specification
function_info <- list(
  name = "get_domain",
  description = "Analyzes domains and decides if it is a news website, a government
  agency website, or a political/social movement website in Brazil",
  parameters = list(
    type = "object",
    properties = list(
      binary = list(type = "integer", description = "1 if the domain matches the target
      criteria, otherwise 0"),
      description = list(type = "string", description = "Explanation of the logic in one
      sentence")
    ),
    required = c("binary", "description")
  )
)
```

B List of Domains

Table 2: Political News Domains in the Estimation Data

a24.com	abcm ais.com	abril.com.br	acaoamazonia.com.br	acidigital.com	acritica.com	agazeta.com.br
agoranoticiasbrasil.com.br	agorarn.com.br	agorars.com	ainvestigacao.com	alegretetudo.com.br	aliadosbrasiloficial.com.br	almapreta.com.br
aloalobahia.com	antropofagista.com.br	aosfatos.org	apublica.org	areavip.com.br	aredacao.com.br	arnewsnoticias.com
atarde.com.br	atrombetanews.com.br	auditoriacidada.org.br	bahia.ba	bahianoticias.com.br	bcb.gov.br	bemparana.com.br
blogbbm.com	blogdobarroto.com.br	bloomberglinea.com.br	bmcnews.com.br	bn.gov.br	bnews.com.br	boatos.org
bonifacio.net.br	brasil247.com	brasildefato.com.br	brasildefatomg.com.br	brasilparalelo.com.br	brasilsemmedo.com	braziljournal.com
brazilurgente.com.br	caixa.gov.br	camara.leg.br	campograndenews.com.br	canalrural.com.br	cartacapital.com.br	causaoperaria.org.br
central3.com.br	chumbogrossomanaus.com.br	claudiodantas.com.br	clickpetroleoegas.com.br	clicrbs.com.br	cnj.jus.br	cnnbrasil.com.br
coisasdapolitica.com	conexao1.com	conexaopolitica.com.br	congressoemfoco.com.br	conjur.com.br	consumidor.gov.br	contrafatos.com.br
correio24horas.com.br	correiobrasiliense.com.br	correiopovo.com.br	cpacbr.com.br	crusoe.com.br	deolhonosruralistas.com.br	dgabc.com.br
diariocarioca.com	diariodepernambuco.com.br	diariodobrasilnoticias.com.br	diariodocentrodomundo.com.br	diariodolitoral.com.br	diariodopoder.com.br	diariodorio.com
diariodotransporte.com.br	direitadegoiania.com.br	direitaonline.com.br	ebc.com.br	em.com.br	enfoco.com.br	epochtimes.com.br
esmaelmorais.com.br	estadao.com.br	estado.sc.gov.br	estudosnacionais.com	exame.com	f5online.com.br	fapesp.br
fepal.com.br	fiocruz.br	folha.com	folhadestra.com	folhape.com.br	folhavoria.com.br	forasalles.com.br
forbes.com.br	fpabramo.org.br	gazetabrasil.com.br	gazetadopovo.com.br	gazetapotiguar.com.br	gcm ais.com.br	geledes.org.br
globo.com	gov.br	guaiba.com.br	hojebahia.com.br	horabrasilia.com.br	i24.com.br	ibahia.com
iclnoticias.com.br	ig.com.br	infomoney.com.br	informeorwell.com	inteligenciafinanceira.com.br	intercept.com.br	istoe.com.br
istoedinheiro.com.br	itatiaia.com.br	jacobin.com.br	jornalbrasilonline.com.br	jornaldacidadeonline.com.br	jornaldebrasilia.com.br	jornaldocomercio.com.br
jornalggn.com.br	jornalopcao.com.br	jornalrazao.com	jota.info	jovempan.com.br	lance.com.br	leiaja.com
luizmuller.com	massanews.com	matinaljornalismo.com.br	medicospelavidacovid19.com.br	meioemensagem.com.br	meionews.com	metro1.com.br
metropoles.com	metroworldnews.com.br	midianews.com.br	midianinja.org	migalhas.com.br	moneytimes.com.br	movimentosempicanh
mpf.mp.br	mst.org.br	mulheresconservadoras.org	naoeimprensa.com	ndmais.com.br	newsatual.com	nexojornal.com.br
noataque.com.br	no pontodofato.com	noticiabrasil.net.br	noticiasaominuto.com	noticiasaominuto.com.br	novonoticias.com.br	nsctotal.com.br
nucleo.jor.br	oantagonista.com.br	obastidor.com.br	ocafezinho.com	ocp.news	oeco.org.br	oliberal.com
opopular.com.br	opovo.com.br	orepublicano.com.br	otempo.com.br	outraspalavras.net	paraibaonline.com.br	paulofigueiredoshov
paulogala.com.br	pensandodireita.com	piaui.co	pinoquionews.com.br	planalto.gov.br	plantaobrasil.net	pleno.news
plural.jor.br	poa24horas.com.br	poder360.com.br	politicalivre.com.br	politicos.org.br	ponte.org	portalamazonia.com
portalnapolis.com	portaldacidade.com	portalemo ff.com.br	portalnovonorte.com.br	portalpopline.com.br	portalvoce.com	pragmatismopolitico
pt.org.br	ptnacamara.org.br	ptsapaulo.org.br	publica.org.br	r7.com	radio93.com.br	radiopeaobrasil.com
redebrasilatual.com.br	reporterdiario.com.br	republicanoticias.com.br	revistacenarium.com.br	revistaesmeril.com.br	revistaforum.com.br	revistamensch.com.br
revistaoeste.com	revistaonline.news	ric.com.br	rothbardbrasil.com	rtbrasil.com	saibamais.jor.br	samiabomfim.com.br

Table 2: List of Domains (*continued*)

Domain1	Domain2	Domain3	Domain4	Domain5	Domain6	Domain7
sbt.com.br	seliganerd.com	senado.leg.br	sociedademilitar.com.br	stf.jus.br	sul21.com.br	terra.com.br
terrabrasilnoticias.com	theinvestor.com.br	tribunadonorte.com.br	tribunaonline.com.br	tribunapr.com.br	tse.jus.br	tupi.fm
uai.com.br	uhnplus.com	uol.com.br	upiara.net	urbsmagna.com	verdesmares.com.br	viomundo.com.br
visaodadireita.com	vistapatria.com.br	votosdeputados.com.br	vozascomunidades.com.br			

C Validation: Domains Ideology Estimation

Figure 10 presents the 50 most shared political news domains in the pre-ban period. The two most frequently shared domains are *revistaoeste.com* and *jornalcidadeonline.com.br*. These are well-known right-leaning digital outlets in Brazil (Aruguete, Calvo and Ventura, 2021; Mont’Alverne et al., 2024). Their high prevalence in the data already indicates how the information environment on X, even before the ban, leaned more conservative. Other more traditional media outlets also appear as highly shared in our data, such as *uol.com.br*, *globo.com*, *cnnbrasil.com.br*, and *estadao.com.br*. Official government portals (*gov.br*, *senado.leg.br*, and *votosdeputados*), small journalistic-based personal blogs (*paulofigureidoshow.com*), and local news organizations (*oliberal.com* and *opovo.com.br*) are also included as domains in the counts matrix.

D Validation: Users’ Groupwise Recentering

In this section, we validate our users’ ideology estimation with two additional pieces of information. First, we provide the most shared domain across the users’ ideological bins. This measure allows us to validate the ordering of users’ scores using substantive knowledge of the domains and their mapping to the users’ distribution (Figure 11). Second, we compare results using the original scale and our groupwise recentering, as described in the methods section in the paper (Figure 12). This figure illustrates the issue of asymmetric distribution and how a narrow group of left-leaning users skews the left tail of the distribution. Figure 12 shows how our recentering does not affect the ordering of the distribution substantively, given how the main domains used in the ideology estimation remain mostly the same across the bins.

Figure 10: Top 50 Political News Domains Shared 90 days before the Ban




Figure 11: Most Shared political news domain across users' ideological bins





Figure 12: Most Shared political news domain across users' original ideological bins



E Additional Results: Partisan Responses to the Ban

Figure 13: Average Users' ideological scores weighted by overall tweet activity






Figure 14: Concentration of Tweet Activity and News Sharing



F Additional Results: The Effects of the Ban on Social Media Engagement

Figure 15: Concentration of engagement with content across domains



G Analysis of Uncollected Users

To derive our measure of ideology, we utilized an ongoing Twitter Decahose collection to identify politically engaged ordinary users. To fill in their timelines in the target period, we attempted to scrape the profiles of the 9061 users identified and succeeded for 7471 users. The reasons for which we were unable to collect accounts were (a) the user deleting their account before our collection, (b) the user changing their Twitter handle and us not having an appropriate numerical crosswalk ID, (c) the account being suspended for rules violations, (d) the account being set to private, and (e) in at least one instance an account appearing to have deleted all of their Tweets.

Since these reasons for missingness are hardly ‘at random,’ we examined the observable differences of these collected (henceforth called ‘Matched’) and uncollected (‘Non-matched’) users based on what we knew about them from the Decahose collection: their estimated ideology, how many times they have shared news, the number of distinct news domains they shared, and the total posts previously observed. We judge that, although some of these differences are statistically significant, they are not substantively significant. Moreover, their substantive interpretation largely supports our broader theoretical expectations.

Table 3: Balance Table: Matched vs. Non-Matched Users

Variable	Matched		Non-Matched		Difference	SMD	p-value
	Mean (Matched)	SD (Matched)	Mean (Non-matched)	SD (Non-matched)			
Ideology (Dim 1)	0.025	0.990	-0.117	1.037	0.142	0.140	6.75e-07
Total News Shares	63.500	78.247	59.358	72.753	4.141	0.055	0.042
N Distinct Domains	9.781	4.331	9.755	4.024	0.026	0.006	0.818
Total Posts (Decahose)	91.919	115.767	95.302	118.545	-3.383	-0.029	0.300

Note: N matched = 7471; N non-matched = 1590. SMD = standardized mean difference. P-values from Welch’s t-test.

From our balance table, we see that non-matched users are (a) more liberal, (b) share less news, and (c) posted more than the users whose timelines we were able to collect (no difference being observed in the number of distinct domains shared). This is largely supportive of our partisan compliance story. One of the central findings of our main analysis is that conservative users were significantly more likely to circumvent restrictions, with right-leaning

news domains becoming considerably more prevalent. A disproportionate ‘disappearance’ of more left-leaning accounts is the other side of this coin: left-leaning accounts were somewhat more likely to leave the platform permanently. Given that these users produce somewhat more news, this disproportionate exodus may also help to explain why right-leaning news domains became considerably more prevalent on the platform.

While this interpretation is appealing, we stress that results based on these differences are not substantively significant. To demonstrate this, we fit a logistic regression predicting missingness as a function of observables, the results of which reflect the findings of our balance tests.

Table 4: Logistic Regression: Predicting Timeline Match Status

Term	Estimate	SE	Statistic	p-value	CI Low	CI High	AME	AME SE	AME p-value
Intercept	1.6299	0.0769	21.1859	0.0000	1.4787	1.7804	NA	NA	NA
Ideology (Dim 1)	0.1282	0.0265	4.8358	0.0000	0.0760	0.1799	0.0185	0.0038	0.0000
Total News Shares	0.0027	0.0007	3.8596	0.0001	0.0014	0.0042	0.0004	0.0001	0.0001
N Distinct Domains	-0.0146	0.0097	-1.5019	0.1331	-0.0335	0.0045	-0.0021	0.0014	0.1331
Total Posts (Decahose)	-0.0011	0.0003	-3.3301	0.0009	-0.0018	-0.0005	-0.0002	0.0000	0.0009

Note: N = 9,061. Dependent variable: matched (1) vs. non-matched (0). AME = average marginal effect. McFadden pseudo- R^2 = 0.0052; AUC = 0.551.

When we examine the standardized coefficients of this exercise, however, the effects are all clearly either ‘small’ or ‘very small’ based on common rule of thumb heuristics (see <https://github.com/easystats/effectsize/blob/main/vignettes/interpret.Rmd>).

These similarities can be seen quite clearly in simple distributional comparisons which put the mean-differences in perspective.

In Figure 18, we can see quite clearly the lack of meaningful distributional differences (despite the small mean differences) in all our measures of activity. Figures 17 and 19 add somewhat more nuance to the story given above. As previously discussed, we see that those users who left the platform were disproportionately to the left of those who stayed (the most likely to stay being on the far right of our estimated ideology). When comparing deciles, we see this was driven primarily by those on the far-left of our ideological spectrum. While this pattern should not be over-interpreted, we note that (a) it is the most striking difference we

Figure 16: Standardized Logistic Regression Coefficients




Figure 17: Distribution of Ideology Across Matched and Non-Matched




Figure 18: Distribution of Activity Across Matched and Non-Matched





Figure 19: Comparison of Ideology Deciles




are able to observe between accounts that we were able to collect and those that we were not, and (b) the pattern of attrition lines up with our broader theoretical expectations and findings.

H Robustness Checks: Event Study Models


In this section, we present robustness checks and extensions for the event study model presented in the main paper. First, we show that the results discussed in the paper hold across distinct outcomes. Figure 20 shows the same event study model discussed in the main paper using engagement outcomes, such as likes, replies, quotes and retweets received at the month level. For comparison, we also present the results for the total number of tweets posted as in Figure 5. In addition, we show that the results hold across different cutoffs for defining right-leaning and left-leaning users. Figure 21 shows the baseline result reported in the paper – with zero in the user ideology scale as a cutoff – compared to three other options for cutoff. Overall, the effect size of the ban increases with a more restricted definition of left and right-leaning users.

Figure 20: Event Study Estimates: Poisson Regression on Daily Rates Across Engagement Metrics



Full regression results on SM Table 4

Figure 21: Event Study Estimates: Robustness Across Cutoff for the User Partisan Leaning




Full regression results on SM Table 7

Second, we show our results are robust across distinct bandwidths around the ban. Instead of using a month-level event study as reported in the main paper, Figure 22 uses a canonical difference-in-differences model with two time periods (before and after the ban), varying the bandwidth from 30 to 10 days around the event. Results are also robust to this model specification.


Lastly, Figure 23 verifies the robustness of our findings to a distinct model specification. In the main paper, we present results from an event study using a Poisson model for the total number of posts per day and fixed effects at the day and user levels. Figure 23 relaxes the Poisson distribution assumption and models the total number of tweets with the same fixed effects at the user and month level using an OLS. Results are the same for Tweets Posted and other engagement metrics, although less precise due to the smaller sample size.

Figure 22: Event Study Estimates: Difference in Difference on Daily Rates of Posts and Shares on X



Full regression results on SM Table 8

Figure 23: Event Study Estimates: OLS Results



Full regression results on SM Table 6

I Regression Tables

Table 5: Regression Table: Poisson Event Study on Daily Rates of Posts and Engagement on X

	Tweets Posted	Likes	Replies	Quotes	Retweets
Jul x Right-Leaning	0.065 (0.050)	0.596*** (0.190)	-0.076 (0.129)	0.042 (0.299)	0.637*** (0.229)
Aug x Right-Leaning	0.236*** (0.055)	0.463 (0.351)	0.207 (0.291)	0.817* (0.451)	0.595 (0.374)
Sep x Right-Leaning	1.766*** (0.142)	1.650** (0.657)	1.503*** (0.402)	1.866** (0.787)	2.285*** (0.673)
Oct x Right-Leaning	0.406*** (0.061)	0.733* (0.402)	0.481 (0.327)	0.911 (0.714)	0.824* (0.459)
Nov x Right-Leaning	0.222*** (0.057)	1.036*** (0.366)	0.740** (0.298)	1.025 (0.639)	1.249*** (0.413)
Dec x Right-Leaning	0.229*** (0.056)	0.479 (0.324)	0.034 (0.299)	0.617 (0.692)	0.745** (0.359)
Num.Obs.	1 598 794	1 524 108	1 357 830	457 318	1 313 960
R2	0.489	0.792	0.735	0.756	0.806
R2 Within	0.009	0.012	0.010	0.015	0.014

* p < 0.1, ** p < 0.05, *** p < 0.01
Right-Leaning = ideology score ≥ 0 .

Table 6: Regression Table: Poisson
Event Study on News Domain
Sharing Rates on X

	News Shared
Jul x Right-Leaning	0.054 (0.082)
Aug x Right-Leaning	0.108 (0.072)
Sep x Right-Leaning	1.635*** (0.354)
Oct x Right-Leaning	0.190 (0.129)
Nov x Right-Leaning	0.053 (0.123)
Dec x Right-Leaning	0.009 (0.156)
Num.Obs.	46 224
R2	0.892
R2 Within	0.021

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Right-Leaning = domain ideology
score ≥ 0 .

Table 7: Regression Table: OLS Event Study on Daily Rates of Posts and Engagement on X

	Tweets Posted	Likes	Replies	Quotes	Retweets
Jul x Right-Leaning	0.980*** (0.293)	61.139*** (18.282)	0.416 (0.441)	0.035 (0.119)	11.173*** (3.714)
Aug x Right-Leaning	2.378*** (0.429)	69.221*** (20.533)	2.985** (1.340)	0.366** (0.179)	13.119*** (4.446)
Sep x Right-Leaning	8.042*** (1.082)	101.675* (55.520)	3.686 (3.035)	0.431 (0.378)	12.490 (10.044)
Oct x Right-Leaning	2.437*** (0.462)	41.106*** (15.638)	2.190*** (0.658)	0.166 (0.107)	4.786 (3.116)
Nov x Right-Leaning	1.659*** (0.435)	74.131*** (20.267)	4.487*** (1.436)	0.295** (0.139)	14.829*** (3.879)
Dec x Right-Leaning	1.823*** (0.414)	30.725* (16.228)	0.374 (0.965)	0.071 (0.196)	8.355** (3.410)
Num.Obs.	585 453	585 453	585 453	585 453	585 453
R2	0.516	0.493	0.217	0.341	0.528
R2 Within	0.002	0.000	0.000	0.000	0.000

* p < 0.1, ** p < 0.05, *** p < 0.01

Reference month: June 2024. Right-Leaning = ideology score ≥ 0 .

Table 8: Regression Table: Poisson Event Study Across Ideology Cutoffs

	Cutoff: 0	Cutoff: +/-0.5	Cutoff: +/-1	Cutoff: +/-2
Jul x Right-Leaning	0.065 (0.050)	0.062 (0.053)	0.048 (0.055)	0.037 (0.064)
Aug x Right-Leaning	0.236*** (0.055)	0.246*** (0.058)	0.228*** (0.059)	0.234*** (0.067)
Sep x Right-Leaning	1.766*** (0.142)	2.083*** (0.168)	2.343*** (0.198)	2.516*** (0.230)
Oct x Right-Leaning	0.406*** (0.061)	0.412*** (0.064)	0.465*** (0.068)	0.503*** (0.077)
Nov x Right-Leaning	0.222*** (0.057)	0.219*** (0.060)	0.228*** (0.064)	0.259*** (0.074)
Dec x Right-Leaning	0.229*** (0.056)	0.238*** (0.059)	0.264*** (0.062)	0.270*** (0.069)
Num.Obs.	1 598 794	1 535 664	1 474 460	1 303 474
R2	0.489	0.490	0.490	0.489
R2 Within	0.009	0.010	0.010	0.009

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Each column uses a different cutoff for defining left vs. right-leaning users.

Table 9: Regression Table: Difference-in-Differences Poisson Models on X Ban

	Users (30-day)	Users (10-day)	Domains (30-day)	Domains (10-day)
Post (During Ban)	-1.273*** (0.006)	-1.311*** (0.011)	-1.227*** (0.031)	-1.002*** (0.049)
Right-Leaning	-0.164*** (0.002)	-0.173*** (0.003)	0.704*** (0.007)	0.720*** (0.011)
Post x Right-Leaning	1.535*** (0.007)	1.578*** (0.012)	1.523*** (0.032)	1.266*** (0.052)
Num.Obs.	440 789	141 949	12 744	4104

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Bandwidth refers to the number of days before and after the ban.