

The Advantages of Overt Censorship: How Platform Blocking Enables China to Reshape the Global Information Ecosystem

Matt DeButts*

May 22, 2026

Abstract

Why do authoritarian regimes conduct overt censorship if they have more sophisticated tools available? This paper suggests that such blatant censorship can shift digital communication onto regime-controlled platforms over time, facilitating delegated information control at low marginal cost to the regime. While censorship may produce backlash among existing media producers, these producers eventually exit the ecosystem and are replaced by more compliant entrants. I investigate this compositional explanation for censorship within the Chinese-language global podcast ecosystem. Gathering 76,212 podcasts and 7,050,824 episodes over a 12-year period along with platform user comments and engagements, I show that domestic Chinese platforms benefit from the censorship of international competitors and that the ecosystem as a whole becomes depoliticized in the wake of overt censorship, consistent with long-run compliance benefits for the regime. These results have implications for our understanding of authoritarian information controls, as well as transnational influence and digital platform governance.

Keywords: censorship, platforms, media, backlash, transnational influence

PRELIMINARY, PLEASE DO NOT CIRCULATE

*Ph.D. candidate, Department of Communication, Stanford University, mdebutts@stanford.edu

Autocratic rulers command an increasingly sophisticated array of digital tools to manipulate online information, including state-of-the-art surveillance, generative artificial intelligence, crowd-sourced propaganda, coordinated bot networks, targeted harassment, and platform-level “shadowbanning” and slowdowns (Bradshaw and Howard 2018; DiResta et al. 2019; Hassan, Mattingly and Nugent 2022; Lu et al. 2025; Roberts 2018; Stukal et al. 2017; Woolley 2020). These strategies are thought to offer advantages over blunt instruments of political control, such as nationwide platform or internet blockages, by introducing uncertainty about the extent of state intervention and thus diminishing the likelihood of public backlash (Earl, Maher and Pan 2022). The widespread adoption of these digital tools has led some scholars to herald the onset of “informational autocracies”, whereby rulers preempt the formation of public discontent by covertly shaping the information environments to which their citizens are exposed (Guriev and Treisman 2019).

Yet despite the striking advantages and increasing availability of covert digital tools, modern autocrats continue to conduct overt digital repression such as nationwide platform blockages and internet blackouts (Earl, Maher and Pan 2022; Gohdes 2015, 2020). These tactics are well known to trigger public backlash: they anger the public, increase attention to the censored information, incentivize users to circumvent state controls, and intensify suspicion of government policies and propaganda (Hassanpour 2014; Hobbs and Roberts 2018; Jansen and Martin 2015). Why, then, do states continue to use them?

Conventional accounts argue that states adopt these digital tactics when they lack capacity for more fine-grained repression or when they face acute political crises (Feldstein 2021; Gohdes 2023; Howard, Agarwal and Hussain 2011; Rydzak, Karanja and Opiyo 2020). This paper argues, in contrast, that overt repression can emerge from the opposite of those two conditions: when high-capacity states face no imminent political threat, they can endure short-term backlash to secure long-term gains in digital compliance. Specifically, I argue that overt digital repression — such as nationwide platform-based censorship — can shift media ecosystem growth onto regime-compliant digital alternatives, achieving strategic

benefits for the regime at marginal long-term cost. Producers who experienced repression exit the ecosystem and are replaced by new users who, having seen nothing else, adopt the regime-preferred norms of the surveilled alternative. Over time, the media ecosystem reaches an equilibrium that serves the interests of the ruling regime.

I illustrate these ecosystem-level effects in the case of Chinese-language podcasts, which emerged in the late 2010s as a globally distributed network of producers, hosting platforms, podcast apps, and online listeners knitting together the Chinese mainland and a constellation of global diaspora. At first, these media producers broadcast primarily on foreign-based platforms and digital apps, largely owing to podcasting’s roots in open source software architecture and Western radio journalism. In the late 2010s, however, the Chinese government targeted these foreign-based platforms with a wave of overt and highly-publicized instances of digital censorship, including a widespread blockage in June 2020. I show that these instances of overt censorship shifted growth in the medium away from foreign platforms to mainland Chinese substitutes, and that this shift coincided with a depoliticization of the podcasting medium owing to selective entry and exits by podcast producers. Political content in podcasts was associated with a greater risk of podcast cessation, and platform algorithms disproportionately recommend non-political shows. Over time, this overt censorship helped to produce a less politically potent communication medium — consistent with the long-run objectives of authoritarian information control.

This paper proceeds as follows. In Section 1, I describe why autocrats would seek to shift communication onto delegated platforms, as well as how overt censorship can shift the composition of platforms and producers. Section 2 briefly describes Chinese-language podcasting, and why it is an appropriate site to test my theory. Section 3 introduces the data, while Section 4 introduces the research design. Results are presented in Sections 5 and 6. I conclude with a discussion of the implications and limitations in Section 7.

1. Theoretical Argument

Authoritarian regimes seek to dominate online speech production while minimizing political and economic cost. One common approach to achieving this domination is to outsource speech regulation to lower-level state actors, corporate entities, and ultimately individuals themselves in the form of self-censorship (Link 2002). This delegation, with political red lines that are typically left intentionally vague, offers several strategic benefits (Fan, Pan and Zhang 2025). It is less economically costly, as lower-level actors shoulder the financial burden of algorithmically and often manually removing speech acts to comply with state diktats (Fang 2022; Huang and Wu 2023). It is politically advantageous, as the regime can blame unpopular content moderation decisions on lower-level actors and occasionally reverse them (Dimitrov 2023). It is more strategically adaptable, as non-state actors can interpret state interests in digital contexts where the state may have only limited visibility (Stern and Hassid 2012). Most importantly, the constant threat of state repression leads to widespread self-censorship — and indeed, the easiest speech to regulate is the speech that is never produced at all.¹

At the same time, a great deal of digital communication occurs outside the operational jurisdiction of these lower-level actors. Foreign-originated speech, for example, is more challenging to curtail at the source due to geographic distance and the protection of foreign sovereignty (DeButts and Pan 2026; Esberg and Siegel 2021), and often flows into regime borders via gaps in censorship (Chang et al. 2022; Lu et al. 2022). Speech on niche digital platforms may not be penetrated by the state’s repressive institutions (Corduneanu-Huci and Hamilton 2022), and speech occurring on foreign-owned platforms cannot be directly surveilled or deleted without acquiescence of the corporate parent (Pan 2017). The relative

¹Delegation to lower-level actors also has important drawbacks. Most well-established is the “authoritarian information problem”: autocratic rulers lack information about the true sentiments of the population because speech production is regulated by these lower-level actors before it reaches the center, including because it is never produced at all (Egorov, Guriev and Sonin 2009; Gehlbach et al. 2026; Wintrobe 2000). For more information on how the state seeks to address these shortcomings, see (Dimitrov 2023; Gueorguiev 2021; Lorentzen 2014; Truex 2017).

illegibility of these digital spaces is challenging for authoritarian regimes, as speech on these platforms falls between the cracks of domestic information control apparatuses erected to maximize compliance and minimize costs.

How does the state achieve political control over these hard-to-reach spaces? One option is to expand state power into new digital territory by investing in digital surveillance tools or enrolling proxy actors such as trolls, bots, or online crowds to project state influence in the absence of state institutions (King, Pan and Roberts 2017; Ong and Cabañes 2019). When speech occurs on commercially-organized digital platforms, the state can threaten legal or economic sanctions to secure political compliance (Creemers 2019), and if the company is foreign-domiciled, can also threaten diplomatic consequences to the platform’s parent country (Dukalskis et al. 2022). These approaches offer the benefit of increasing state power but come at substantial economic or political cost: the state must invest in the institutional capacity to administer new digital territories or to ensure that newly enlisted political allies are operating with state-aligned incentives. Because the spaces were previously ungoverned, the state must monitor them for ongoing compliance and faces the risk that platform operators will renege if the threat of sanction recedes.

Rather than expanding state power into new digital spaces, an alternative approach is to render those spaces inaccessible to the public via censorship. However, the literature on overt censorship is unambiguous about its unintended consequences: it results in backlash. The 2014 censorship of Instagram by the Chinese government, for example, led to not only increased adoption of censorship circumvention tools among existing users, but consumption of other politically sensitive content once circumvention had been achieved (Hobbs and Roberts 2018). Sudden disruption can catalyze revolutionary protests by providing a coordinating device for political opponents (Hassanpour 2014), and when citizens are aware of the removal of content, they often specifically seek it out, a phenomenon known as the “Streisand Effect” (Jansen and Martin 2015). Over time, public awareness of overt censorship can erode the credibility of government information sources (Chen and Yang 2019), and

increase citizens’ motivation to access alternative information precisely when the state most needs political control such as during political or economic crises (Chang 2021).

An enduring puzzle of authoritarian information governance is that the state nevertheless frequently blocks platforms outright despite its well-established negative consequences. In 2020, for example, Pakistan banned access to Tik Tok (Lakhani 2020), and since 2022, Russia has blocked access to Instagram, Facebook, and X (formerly Twitter) (Sauer 2022) while Iran banned access to Instagram and WhatsApp (Chan and Ortutay 2025). The Chinese government bans dozens of digital platforms each year, some of which seem to pose little immediate threat to government rule: in 2018, the government blocked a humor app with 30 million active users due to allegedly “vulgar” content (Kuo 2018). If overt censorship leads to backlash, why is it still so common?

Conventional explanations for the persistence of overt censorship focus on the role of political crises and the absence of state capacity: when states face immediate threats to their rule, they discount long-run negative consequences such as censorship circumvention or popular dissatisfaction due to the short-run benefits of disrupting opposition movements (Gohdes 2015). A complementary explanation concerns state capacity. When states lack the capacity to perform targeted censorship, which is generally believed to be more effective than overt censorship,² they turn to “rudimentary” (Frantz, Kendall-Taylor and Wright 2020) or “clumsy” (Howard, Agarwal and Hussain 2011) methods such as disconnecting digital platforms (Freyburg and Garbe 2018), similarly to how low-capacity states resort to indiscriminate physical repression when targeted repression is unavailable to them (Blaydes 2018). Overt censorship is thus often associated with autocratic weakness: either the regime’s grip on state power is slipping or the state’s capacity to control information is incomplete. In either case, this account predicts that overt censorship should be observed primarily in cases of popular contention or insecure autocratic rule.

In fact, overt censorship is observed in a variety of contexts, only a portion of which coincide with political crises or low-capacity states (Howard, Agarwal and Hussain 2011).

²For a review of the efficacy of digital repression, see Earl, Maher and Pan (2022).

Pakistan’s block of TikTok, a short-form video app, for example, or China’s frequent blockage of apps and websites, are publicly announced, occur outside political crises and, in the case of China, in the context of a high-capacity state operating what has sometimes been called the most sophisticated digital apparatus the world has ever seen (King, Pan and Roberts 2013). To address these remaining instances, we need an account of overt censorship that can explain why it occurs within high-capacity states and outside of political crises.

1.1. Compositional Explanation for Overt Censorship

Online speech production features four main actors: the speech producer (used interchangeably in this paper with “media producer”), the audience (the person or persons engaging with the speech), the digital platform (the online location where the speech occurs), and the government (who seeks to influence that speech). While the speech producer may have diverse motivations for speaking, they typically seek to be heard by an audience — and are thus incentivized to coordinate with audiences on the selection of a digital platform. All else being equal, both speech producers and audiences prefer platforms with lower levels of censorship, as censorship limits the diversity of speech that can occur on the platform.³ Most of all, however, both speech producers and audiences have a strong desire to select the same platform, lest they find themselves speaking (or listening) in a deserted digital room.

Both producer and audience choice of platforms is relatively “sticky”. Once each selects a platform, they grow familiar with locating each other in digital space, understand its user interface, develop a digital history, and in many cases invest in platform-specific online relationships or communities. Indeed, it is precisely this habit-forming tendency of digital platforms that opens the possibility for digital backlash to censorship: audiences adopt circumvention tools when their habits have been suddenly disrupted (Hobbs and Roberts 2018).

³For the purposes of this simplified characterization, content moderation (e.g. the removal of hate speech or child pornography) is treated as distinct from censorship. For more information on the distinction between content moderation and censorship, see (Gillespie 2018).

Digital platforms, meanwhile, seek to accumulate producers and audiences, as they derive profit from hosting speech acts, typically through subscriptions (from either producer or audience), app downloads, or advertising revenues. They are incentivized to comply with state laws via surveillance and censorship — legal violations can be financially ruinous — but must balance investments in compliance against investments in attracting producers and audiences. Because censorship is unpopular with users, they are incentivized to perform the minimal level of censorship needed to satisfy the government. Platforms are also motivated to increase their “stickiness” as doing so enables them to retain users.

The fourth actor, the regime, seeks to maximize compliance while minimizing its own costs. Consequently, it prefers to delegate censorship to platforms (by forcing them to comply with censorious laws and regulations) and to individual speech producers (by encouraging them to self-censor) (Stockmann and Gallagher 2011). While the regime can and does invest in institutions for regulating speech directly, these investments are expensive in part due to the cost of ongoing monitoring and administration (King, Pan and Roberts 2013, 2017). The regime therefore has a strong incentive to influence producers’ and audiences’ choice of digital platform: if speakers and audiences choose freely among digital platforms, they may select less-regulated platforms because they disprefer censorship; the regime would instead prefer that producers and audiences select heavily censored platforms as the site of their digital meet-ups because platform censorship reduces the regime’s own financial outlays. The question then becomes: how does the regime ensure that these regime-compliant platforms win market share if compliance is unpopular with both producers and audiences?

I argue that overt censorship — a one-off intervention with minimal marginal cost — is a cost-effective method of advantaging high-compliance platforms in the market for digital attention. When the state censors a digital platform, they inhibit (but do not eliminate) its ability to attract producers and audiences; the platform’s existing speech producers and audiences can and do engage in censorship circumvention, including via virtual private networks (VPNs), due to the habits they have already formed. New speech producers and

audiences, however, select among the regime-compliant platforms easily available to them, growing those platforms’ market share and shifting the composition of online speech into spaces under delegated regime control. Overt censorship thus produces long-term compliance at the price of short-term backlash, consistent with states operating from strength during periods of calm rather than weakness during periods of crisis.

A key mechanism for this compositional change is the selective entry and exit of speech producers. A long tradition in political communication suggests that mass-mediated political persuasion is difficult, especially when messages contravene prior attitudes or existing media preferences (Kalla and Broockman 2018; Lazarsfeld, Berelson and Gaudet 1968). For this reason, a growing body of literature has emphasized how political actors can reinforce existing beliefs and identities rather than seek to change them (Horz 2024; Shirikov 2024), or turn out existing supporters rather than seeking to convert new ones (Nichter 2008). A compositional explanation for overt censorship extends this logic to information control by proposing that autocrats alter the composition of active media producers, incentivizing — by making their chosen platforms harder to access — the exit of speech producers who experienced overt censorship, while facilitating the entry of new producers who did not. Exiting producers are less compliant, due to either backlash or their pre-existing experiences within a less-censored platform, while new producers are more compliant due to platform norms and their inexperience with a more open alternative. Over time, these shifts reshape the composition of media producers within the digital ecosystem.

This compositional theory has two observable implications. The first is *compositional change*. An overt censorship shock should shift the composition of media production onto platforms already under the state’s delegated control, and this shift should advance the regime’s strategic objectives for information control. The theory is falsified if the shock has no effect or initiates an increase in less-compliant platform adoption, or if the ecosystem writ large does not change or becomes less compliant with regime strategic objectives. These ecosystem-level changes should also be associated with differential media entries and exits,

and the theory is falsified if the platform’s composition remains stable, or if exit and entry are not politically selective, or if selectivity runs in the opposite direction. The second observable implication is *platform-specific equilibria*: the effects of delegated censorship should mean that media production hosted on regime-compliant platforms is systematically different than media hosted on non regime-compliant platforms. The theory is falsified if there is no difference between platforms, or if the difference inclines against the regime’s strategic interest.

This theory operates with several important scope conditions. First, the censoring regime must operate with a long enough time horizon that it is willing to endure the short-term backlash from overt censorship in exchange for long-term gain. If the regime discounts the future due to short-term uncertainty, it is unlikely to engage in compositional censorship. Second, the blocked platform must be sufficiently commoditized that it possesses effective substitutes. If the platform is irreplaceable, users will have no alternate platforms to switch to. Third, the regime must command delegated control over substitute platforms, such that the compositional shifts in user activity are captured by the regime. If users shift from one non-compliant platform to another non-compliant platform, the regime fails to capture the benefits of long-term compliance within its delegated control.

2. The Chinese-language Podcast Ecosystem

I propose to evaluate this compositional explanation in the context of the People’s Republic of China’s (PRC) censorship of the Chinese-language podcasting ecosystem. “Podcast”, a portmanteau of “iPod” and “broadcast”, refers to the medium of digital audio files, released episodically by a media creator, streamed or downloaded to an individual’s personal device such as phone or computer. The medium has exploded in popularity globally in recent years, leading to a rising tide of journalistic and academic interest in the political implications of podcast listening, including its association with misinformation consumption (Rasul et al. 2025), conservative identity politics (Warner et al. 2026), and political polarization (DeMets

and Spiro 2025).

The podcasting ecosystem features the same actors as digital speech production more generally. The speech producer is the podcast creator: an individual or group of individuals who records audio into digital format such as .mp3 or .m4a. That creator then selects a podcast hosting platform, typically a corporate entity such as AnchorFM or Buzzsprout, who hosts the creator’s audio on corporate servers and compiles a Really Simple Syndication (RSS) “feed” for distribution on the open internet.⁴ The hosting platform also provides commercial services such as dynamic insertion of advertisements or listener analytics to help creators understand their audience as well as monetize their content. That platform then pushes content onto a podcast client (also known as “podcatcher” or “podcast app”), which is the user-facing app from which individual listeners access the podcast content.⁵ In podcasting, the “digital platform” is thus disaggregated into the podcast hosting platform (selected by the producer) and the podcast client (an app, selected by the audience), with the two linked by RSS feeds. The regime, meanwhile, can intervene by either blocking the hosting platform or the podcast client or both.

Crucially for my theory, podcast creators have discretion to select among multiple podcast hosting platforms for audio upload and distribution, but once they select a platform, swapping to other platforms is tedious as doing so requires them to migrate their RSS feed, re-accumulate listener analytics and monetization integrations, and gain familiarity with new platform user interfaces. This friction makes podcast hosting platforms a classic example of sticky platform effects: creators are strongly incentivized to retain access to their hosting platform rather than abandoning it in the event of blockage. This close connection between podcast creators and their hosting platform makes overt censorship a likely candidate for backlash as, after adoption, podcast platforms cannot be easily substituted for each other (Hobbs and Roberts 2018). For new creators, however, platforms are far more substitutable

⁴Podcasting hosts can also create and upload their own RSS feed directly, without the podcast hosting platform, but doing so is uncommon as podcast hosting platforms bundle audio storage, analytics, and monetization tools that most creators prefer not to manage themselves.

⁵Podcast hosting platforms sometimes offer their own podcast client (e.g. Spotify) but more often they do not.

as they have not yet uploaded content or developed familiarity with the host.

Prior to 2020, Chinese podcast creators could choose among several foreign hosting platforms — which do not typically perform wide-scale in-house self-censorship on behalf of the government — and domestic platforms, which do.⁶ As an emerging medium with low barriers to entry, podcasting tended to attract large numbers of amateur producers. In the early days of Chinese-language podcasting, many of these amateur producers selected foreign hosting platforms to establish their shows as foreign platforms offered a more robust amateur support pipeline compared to existing mainland Chinese audio platforms, such as Ximalaya and Qingting, who served more professionalized creators, were closed-ecosystem (i.e. did not use RSS), and tended to exert greater creative control over producer content. Podcast audiences, too, skewed richer and more educated than the typical Chinese citizen, with 63.4% owning an iPhone, 86.4% having earned a Bachelor’s degree, and 72.6% accessed podcasts through Western-owned podcast clients such as Apple Podcasts or Pocket Casts (PodFest China 2020). The Chinese podcast market was large and growing: prior to its blockage, Pocket Casts listed China as its seventh-biggest market worldwide (Byford 2020).

This regulatory and architectural distinctiveness attracted a producer pool whose political composition diverged from mainland Chinese digital media more generally. The medium’s roots in Western radio and its diaspora-linked listenership made it a natural site for transmission of foreign ideas — most prominently transnational feminism (Luqiu 2024; Meng and Wang 2025; Wang and Lin 2025; Yang 2023), which the CCP regards as threatening given its self-appointment as guardian of women’s interests (Zheng 2005). Exiled political movements such as Falun Gong also operated podcast channels in the global directory, though most were already individually blocked at the government’s request. The medium’s relative infancy allowed for producer experimentation, while its audio-based content may have provided some protection from predominantly text-based government censorship (Qin 2019).

⁶Domestic platforms operate with official PRC licensing and thus must comply with state censorship or risk legal liability if they fail to do so.

2.1. Overt Government Censorship: June 2020 Shock

The growing influence of audio content attracted the attention of government censors, and by 2019, the PRC government began a “Rectification Campaign” via a series of widely publicized crackdowns (China Cyberspace Network 2019). These crackdowns culminated in the Great Firewall of China’s (GFW) June 2020 simultaneous blocking of seven foreign podcast hosting platforms, jointly responsible for 13.6% of all Chinese-language podcasts available globally at the time — the largest single censorship event the ecosystem has experienced, before or since (Hoang et al. 2021). Alongside the GFW block of hosting platforms, the PRC also censored two podcast clients (Pocket Casts and Castro) which had together commanded almost a quarter of audience market share at the time, and forced Apple to remove dozens of podcasts from its global directory (Hern 2020). For podcast creators uploading content to affected platforms, the ban meant that mainland Chinese users could not access their audio without a VPN. For podcast listeners, users of affected podcast clients had to download circumvention tools or switch to regime-compliant alternatives such as Apple, who performed widespread censorship on behalf of the regime.

This shock provides a strong empirical grounding to test a compositional explanation for overt censorship. The ecosystem’s existing producers skewed somewhat international and political, both of which are points of political sensitivity to the ruling regime, and prior to the crackdown, many producers and audiences met together on digital platforms under incomplete control by the ruling regime. These platform choices were sticky, meaning that overt censorship was likely to result in the adoption of circumvention tools, but the medium was growing rapidly, providing an opportunity to shift the ecosystem’s future growth onto platforms over which the regime had more complete control. Moreover, the CCP is and was one of the world’s most sophisticated informational autocracies, meaning that targeted censorship — of individual podcast creators or audiences — remained within their capacity, making platform-level intervention more likely a strategic choice. The case of podcasting also satisfies the theory’s three scope conditions: the PRC was not in political crisis at the

time of the overt censorship,⁷ the platforms were sufficiently commoditized that creators and audiences could substitute one for the other in the case of blockage, and the regime exercised delegated control over substitute platforms to which new growth could be shifted.

3. Data

I use this June 2020 censorship shock, as well as a subsequent June 2021 platform-specific block (of a single platform, Firstory), to evaluate the effect of overt censorship on media ecosystem growth and the overall composition of media production. I gather four sources of data.⁸

The first source of data is the global universe of Chinese-language podcasts. To gather this data, I follow prior research by querying Apple’s iTunes Search and Lookup APIs, which operate as the *de facto* search directory for the global podcast ecosystem, much like a podcast-specific digital Yellow Pages (Sullivan 2019). This one-time gather was performed in January 2026. Because Apple’s API does not permit direct enumeration of all podcasts in the directory,⁹ I employed a two-phase discovery procedure. In the first phase, starting with a seed list of 391 unique Chinese-language search terms across 22 Apple Podcasts storefronts (over the four primary Chinese-speaking territories of mainland China, Hong Kong, Macau, and Taiwan alongside 18 other countries selected for geographic representativeness), I searched each [storefront, term] pair and paginated to exhaustion (see Appendix 8).¹⁰ In the second phase, I iterated through all podcasts gathered in Phase 1, segmented all collected podcast titles and artist names into bigrams and trigrams, then used this bootstrapped

⁷Although Covid-19 had broken out a few months ago and the Hong Kong security law had recently passed, neither directly threatened regime rule domestically during June 2020. More likely, instead, is that June 2020 represented a “focal point” for preemptive digital repression (Truex 2019).

⁸I received Institutional Review Board (IRB) approval for this research, including trace data collection and qualitative interviews, from the Stanford University IRB in November 2025 (Protocol number 81191).

⁹Some scholars (e.g. Litterer, Jurgens and Card (2025)) use third-party aggregators such as PodcastIndex as a substitute for the Apple directory. I opt for the Apple directory to ensure end-to-end supervision of the data collection process and because Apple’s directory is often used to define the podcast universe itself.

¹⁰In practice, nearly all Chinese-language podcasts were widely available in every other country besides mainland China, as podcasts are typically available by default in other storefronts unless either the podcast host or the country’s government requests otherwise. The average podcast availability overlap across countries, other than mainland China, was over 99%.

bigram and trigram list to re-query each podcast storefront until 100 consecutive searches yielded no new podcasts. This procedure resulted in 202,682 unique podcasts, and returned each podcast’s RSS feed (the external address where the episode information and audio content can be located, used to connect podcast platforms with podcast clients), as well as titles, author information, dates of publication and a unique Apple-assigned identifier. I deduplicated on the basis of RSS feed.¹¹

Chinese-language search terms often surface non-Chinese language podcasts. To subset to Chinese-language podcasts only, I opted for a conservative measure that checked each podcast title and description for the presence of any traditional or simplified Chinese character, while excluding all podcasts that contain Japanese or Korean characters (which share some characters with Chinese).¹² This additional criterion narrowed the pool to 76,212 distinct podcasts, of which 26,837 published any episode in 2025.¹³ Of these 76,212 podcasts, I successfully downloaded RSS feed information from 75,418 of them, producing a panel dataset containing 7,050,824 distinct podcast episodes spanning twelve years (2014-2025). Each episode contained metadata including the podcast creator’s assigned title, description, author, and publication date, as well as the URL where raw audio can be accessed. A podcast’s period of activity extends from its first published episode to its last. The average interval between a podcast’s published episodes is 43 days.

One drawback of podcast data is that we know little about podcast listener behavior, as podcast client download rates of RSS feeds are not observable to third-party researchers. Thus while we can know what content is being produced, we can say little about which

¹¹I also pruned a platform-specific Chinese-language mass ingestion event between August 2022 and July 2023 occurring on Spreaker and Anchor alone, in which the Apple directory anomalously identified a 140% month-on-month growth in Anchor podcasts and a 9000% growth on Spreaker platforms, after which the additional podcasts disappear. I removed all new podcasts only active during that period from the analysis. I am still investigating why this ingestion event occurred.

¹²While many podcast metadata include language flags, manual inspection revealed they are sometimes unreliable as the language flag can differ from the observed text of the title and description.

¹³Apple does not generally remove RSS feeds after a podcast ceases publishing unless it is specifically requested to do so (Sullivan 2019), and podcast hosting platforms have different policies for dormant RSS feeds. If a platform removed an RSS feed prior to January 2026 it was not picked up in my data and is omitted from the analysis. Note also that if a podcast does not submit its RSS feed to Apple it will not be discovered via the directory search. The vast majority of podcast creators submit their directories to Apple as it facilitates audience discovery.

shows are popular or how listeners are engaging with them. To address this deficit, I gather a second source of data from Xiaoyuzhou, a PRC-licensed domestic podcast platform and client founded in March 2020.

Unlike other China-based audio platforms such as Ximalaya and Qingting, Xiaoyuzhou was explicitly marketed as, and designed to be integrated within, the global podcast ecosystem. Its hosting services directly targeted amateur creators, and its podcast client (i.e. app) used Apple’s public directory of podcast feeds to self-populate at the time of foundation. It also offered a series of user-facing features to promote engagement such as the ability to comment on podcasts with clickable timestamps, interact with other users, and assess the popularity of podcasts via listener and subscriber counts. Because Xiaoyuzhou operates as a podcast client as well as podcast platform, its app contains a mixture of “external” podcasts hosted on other platforms, identified via Apple’s directory and streamed through its interface (in its capacity as a podcast client), as well as “internal” podcasts hosted by Xiaoyuzhou directly (in its capacity as a podcast host). While it is not unique in serving as both host and client (Spotify, for example, also hosts and distributes podcasts) these dual roles are relatively uncommon and make it especially valuable for research.

Xiaoyuzhou does not expose podcast information through a public API, but information is gatherable via the internet if one already possesses the podcast or episode unique identifier. I gather these identifiers from three sources. The first is CommonCrawl: I queried the CommonCrawl CDX index across 47 indices spanning 2020-2025 and extracted all domain matches to Xiaoyuzhou. I identified 2,569 podcasts this way. The second is via Google, where I used a search API to perform site-specific searches for 76,212 Apple-surfaced podcast titles.¹⁴ I identified an additional 24,650 podcasts this way, of which 1,409 were already identified via CommonCrawl. The third is by using RegEx to search among already-gathered Apple podcast descriptions for all Xiaoyuzhou URLs and unique identifiers, as podcast creators often provide links to their podcasts streamed on other clients. I identified 8,483 podcasts this way. Finally, I checked each of these existing data for links to other Xiaoyuzhou URLs.

¹⁴<https://serper.dev/>

After combining these sources, I generated 34,841 unique Xiaoyuzhou podcast IDs, from which I was then able to gather play counts, favorites, and “top” 10 user comments (selected by platform algorithm) as well as all replies to those comments, for a total of 5,989,976 comments and replies. I then matched each Xiaoyuzhou podcast to its Apple RSS feed via shared RSS information, or secondarily via self-declared Xiaoyuzhou versions in the Apple descriptions, enabling a one-to-one match for 10,905 podcasts across datasets. In effect, this match allows us to detect user engagement on Xiaoyuzhou for a subset of podcasts from the overall Chinese-language podcast universe.

Third, I conduct semi-structured qualitative interviews with Chinese-language podcast creators based outside mainland China. To date, I have conducted three interviews of approximately 60-90 minutes each, as well as 10 pilot interviews (additional interviews ongoing), recruited through both publicly listed email addresses and snowball sampling. I made efforts to select interviewees from multiple content genres (technology, current affairs, culture) and degrees of professionalization (full-time creators, regular hobbyists, episodic producers). The focus on diaspora-located creators is both theoretical and practical: theoretically, diasporic producers are least likely to worry about physical repression, thus highlighting the role of platform-specific effects during the production process. More practically, interviewing PRC citizens inside China is becoming increasingly challenging due to PRC privacy laws.

These interviews help me understand the motivations and considerations of the podcast creators: why they start a podcast, how they use platforms to interact with audiences, and how platform selection changed in response to the GFW’s blocking events that anchor the empirical analysis. I also asked gently about platform speech regulations, such as whether the platforms guide producers when discussing politically sensitive topics, as well as implicit norms about what is and is not sayable on Chinese-language platforms. These interviews generate descriptive understanding of speech producers’ incentives and their interactions with platforms and audiences, while also surfacing mechanisms not directly visible in the quantitative data, including individual entry and exit decisions, the role of intermediaries in

shaping platform market share, and the cohort-level evolution of professional norms within the Chinese-language podcasting ecosystem.

Fourth, I gather data on Apple Podcast’s recommended podcasts, across six storefronts, once a day for 14 days in May 2026.¹⁵ I then match this data with data from the PodcastIndex,¹⁶ linked by Apple ID, for all podcasts gathered from six country-specific storefronts, including the creator-uploaded label for content category (e.g. “society”, “politics”, “news”, and “culture”). This produces a dataset of recommended shows, and their categories, cross-nationally.

4. Research Design

I employ two plausibly exogenous shocks to the podcast ecosystem to estimate the effect of platform censorship on podcast ecosystem growth. The claim to exogeneity is derived from the indiscriminate nature of the block: it operated upon all podcast creators who selected a blocked platform to host their content, regardless of creator or podcast content. The first censorship event is the more extensive, in June 2020, and was performed across both hosting platforms and podcast clients, and I thus also measure its impact both on the ecosystem as a whole as well as podcasts hosted on a specific platform, AnchorFM, then the world’s second most popular podcast hosting platform (Hoang et al. 2021). The second block, in June 2021, affected primarily Firstory, a Taiwan-based podcast platform.

4.1. Compositional Change

I look at four outcomes measures to detect compositional change. The first is podcast platform adoption. If overt censorship is redirecting new growth away from censored platforms, we should see a reversal in podcast creators’ adoption rates of blocked platforms. To dis-

¹⁵Data gathering is ongoing as of May 2026. This preliminary 14-day window captures the baseline against which longer-term data collection will be benchmarked. Recommendations available at [https://podcasts.apple.com/\[COUNTRYCODE\]/new](https://podcasts.apple.com/[COUNTRYCODE]/new), where [COUNTRYCODE] is the ISO country code for each storefront.

¹⁶<https://podcastindex.org/>

tinguish the effect on mainland users specifically, I create a proxy for mainland content by removing Chinese-language podcasts with traditional Chinese character descriptions, in effect excluding Hong Kong and Taiwan podcasts from this analysis.¹⁷ I measure adoption rate by detecting a podcast’s first published episode and aggregating at the monthly level. While monthly aggregation is somewhat noisy, larger time units of aggregation obscure important month-to-month variation induced by the censorship blocks.

The second outcome measure is the average political value of each podcast episode. A core strategic objective for the PRC is to prevent the circulation of grievances among the populace which can be used for organized opposition (King, Pan and Roberts 2013). Because a manual check of podcast episodes revealed strikingly few overt political podcasts, I define “politics” broadly to include content that discusses politics or general news, as the Party tightly regulates who is allowed to produce “news” via specific government licensing; unlicensed production of the news is technically illegal (Cyberspace Administration of China 2017). I operationalize this category using contrastive sentence embeddings (Rodman 2020). More specifically, I embed the concatenated titles and descriptions for all episodes of a hand-curated basket of 10 politics/news podcasts and 10 entertainment/lifestyle podcasts, matched on publication dates, aggregate at the show level and average across all shows in each basket to produce a “politics” (p^*) anchor and “entertainment” (e^*) anchor centered on each set of shows in multi-dimensional embedding space (for shows used to anchor the embeddings, see Appendix 10). Then, for each podcast episode i , I measure the net difference in cosine similarity to each anchor ($(p^* - i) - (e^* - i)$). While these scores are not interpretable in absolute terms, a political score of 0 indicates that an episode is equally similar to both politics and entertainment anchors, while positive scores mean an episode is more similar to political shows and negative scores mean it is more similar to entertainment shows.¹⁸

¹⁷Hong Kong podcast audio is not generally intelligible to mainland audiences because it is usually in Cantonese. Taiwanese podcasts are intelligible if in Mandarin but can be challenging to understand if it includes the local dialect, Taiwanese.

¹⁸Episodes with concatenated titles and descriptions with 10 characters or fewer (4.8% of corpus) are dropped. Episodes with concatenated titles and descriptions longer than 510 tokens are back-truncated to fit. I embed podcast episodes using moka-ai/m3e-base, an SBERT-style sentence embedding model. Full-text transcription of podcast audio is ongoing as of May 2026, with analysis to be conducted on raw transcripts

I call this outcome variable *political score*. For this outcome measure, each observation is an episode political value. Evidence of compositional change would consist of a decline in political content across the podcast ecosystem.

To assess changes in the first outcome measure, platform adoption, I use a regression discontinuity in time (RDit), a type of quasi-experimental design that estimates whether an outcome changes discontinuously at a known temporal cutoff. The design involves fitting each platform’s birth-rate trajectory as a smooth function of time on either side of a cutoff, in my case the Great Firewall’s blockage, and reporting change to the levels and slopes at the cutoff.¹⁹ The identifying assumption for RDit is that no other shock causes a discontinuity in the treated platforms’ birth rate at the block date, an assumption defensible in my case because the treatment is a sharp externally imposed event with a precise calendar date.

For each blocked platform I fit two RDit specifications by OLS on a ± 24 -month bandwidth (24 months before and after the block). Let y be monthly podcast births. Let t be months from the block. Let D be an indicator variable where $D = 1$ after the block and $D = 0$ before it. The linear specification is:

$$y = \alpha + \beta_1 t + \delta D + \tau(tD) + \varepsilon$$

and the quadratic specification adds curvature:

$$y = \alpha + \beta_1 t + \beta_2 t^2 + \delta D + \tau(tD) + \gamma(t^2 D) + \varepsilon$$

In the linear fit, α is the pre-block birth rate extrapolated to the cutoff, β_1 is the estimated pre-block linear growth rate, δ is the estimated level change at the cutoff, i.e. the change

upon completion. To ensure that this measure substantively captures a show’s political or news content, an independent coder also hand-coded a sample (stratified by podcast and year) of 470 podcast titles and episodes for the presence of politics (labeling each show a “1” for if it “mentions” politics and a “2” if it focuses on politics), and find that this Likert label correlates strongly with the embedding measure: Pearson’s $r = 0.574$. Additional hand-validation is ongoing, with intercoder agreement and LLM-assisted validation labeling planned for summer 2026.

¹⁹I adopt this approach in part because the natural alternative — a synthetic control built from comparable unblocked hosting platforms — is foreclosed in our setting by a SUTVA violation, as the substantive hypothesis under test is that the blocks redirected new Chinese-language podcast production toward mainland platforms, which means the obvious donor units are themselves affected by the treatment via creator-side spillover.

in births per month directly before/after the block, and τ is the post-block change in slope (post minus pre). Because we are interested in long-term ecosystem growth after the block, our coefficient of interest is τ . The quadratic specification adds $\beta_2 t^2$ and $\gamma(t^2 D)$, letting the trajectory curve on each side of the cutoff: β_2 is the pre-block curvature and γ the change in curvature at the block. I treat it as a robustness check and visual aid rather than the primary specification, since with curvature the post-block slope is not constant and τ no longer summarizes growth in a single number.

To assess changes in the second outcome measure, I average all episode political values produced across the entire global ecosystem in each month, with bootstrapped confidence intervals ($b = 1000$) constructed from repeated re-sampling within each month. I then use a placebo permutation test to construct p-values for the observed 3-month window before and after June 2020 change in average episode political values, with placebo interventions drawn from a five-year period surrounding the intervention month, restricted so that full pre/post periods can be observed (April 2018 - September 2022).²⁰ In effect, this placebo test asks the question: if we randomly select another month during this time period and measure the observed change in political values in the three months before and after this placebo month, what is the likelihood that the observed difference in means would exceed the difference from June 2020? Because only June 2020 affects all stages of the podcast production process — simultaneously blocking individual podcasts, platforms, and clients — I perform these calculations on the June 2020 shock alone.

Third, I decompose the observed changes in political values to test the mechanism: is the observed change in political values due to within-unit change — consistent with individual podcast creators changing their political-ness over time — or due to differential entries and exits, as the pool of podcast producers itself changes in the wake of a censorship event? A compositional explanation for overt censorship would suggest that observed changes are partially attributable to changes in *who* produces. Podcast producers who experienced the

²⁰To avoid treatment contamination of the placebo draws, one month on either side of June 2020 was excluded from the analysis.

overt censorship shock should die off over time, either due to censorship or self-censorship or self-selection, and new entrants should, on average, be less political owing to their lack of experience in a pre-censorship equilibrium or the practice of self-censorship at the outset. To calculate this measure, I extend the pre/post censorship to six months (three months may erroneously flag podcasts as entered or exited because they do not produce a podcast in the pre or post-window), and decompose the total change in mean political score ($M_{post} - M_{pre}$) into three components: an exit term measuring the pre-block mean if soon-to-exit outlets are removed ($C_{pre} - M_{pre}$), a within-outlet term measuring change in outlets present in both windows ($C_{post} - C_{pre}$), and an entry term measuring the post-block mean with and without entrants included ($M_{post} - C_{post}$): M denotes each window’s episode-weighted mean across active shows, and C denotes the same mean restricted to shows publishing at least one episode in each window. Bootstrapped 95% confidence intervals ($b = 1000$) come from a two-stage resample where shows are drawn with replacement (preserving status as exiters/continuing/entrants) and episodes resampled with each show’s window.

Finally, to more directly test whether political value predicts exit, I test whether the pre-shock political value of a show, controlling for the number of episodes published in the 12 months leading up to June 2020 (logged), is associated with an increased likelihood of ceasing publication in the 12 months after the censorship shock, due to either censorship or self-selected exit.²¹ I normalize shows’ political values using Z-scoring across the at-risk podcast universe, and conduct a logistic regression of the following form:

$$\Pr(E_i = 1) = \text{logit}^{-1}(\alpha + \beta z_i + \gamma \log(1 + n_i^{\text{pre}}))$$

Where the outcome variable $\Pr(E_i = 1)$ represents the probability that podcast i posts no episodes in the twelve months post-intervention, z_i represents the show’s normalized political score (relative to other shows) during the 12-month period prior to June 2020, and n_i^{pre} controls for the number of episodes published in the pre-period window. β , representing

²¹I extend the post-period to 12 months, rather than six, to diminish the stock of “false exits”. At a six-month pre/post window, 23% of shows classified as exits subsequently publish an episode after the six-month window (299/1290); at a 12-month window, that figure drops to 10% (160/1691).

the change in log-odds of exit associated with one standard deviation increase in political score, is the coefficient of interest. If β is positive, shows that are more political than mean are more likely to exit; if β is negative, political shows are more likely than entertainment shows to remain.

4.2. Platform-specific Equilibria

Part of a regime’s motivation to move content onto regime-surveilled alternative platforms is the delegation of censorship to the platform operators: content on these platforms should be systematically associated with greater regime compliance. As entrants to an ecosystem, new media producers frequently “learn” from existing media content, and are likely to replicate the speech norms already prevalent on a platform (Danescu-Niculescu-Mizil et al. 2013). The platform itself can also regulate speech production, both via self-censorship on the regime’s behalf and algorithmic promotion (Lu and Pan 2022). Both social learning and algorithmic promotion may influence the creation of new shows. As one Chinese-language podcast producer put it: “When the wheels start to spin, more and more creators jump into this field. Creators in China do not have their own ideas. A creator’s idea is — Oh, what show is popular [...] if such-and-such programs on this platform are more popular or more easily recommended, then I will make such a show.”²²

I construct two tests of platform-specific equilibria. The first is the prevalence of politically sensitive keywords, specifically of the names of the last three chairmen of the Chinese Communist Party: Xi Jinping, Hu Jintao, and Jiang Zemin. Chinese political leaders are both quite relevant to Chinese audiences and also quite sensitive: pilot interviews with podcast creators suggest that they are discouraged from mentioning leaders’ names when uploading episode descriptions to PRC-surveilled platforms. I thus conduct a simple keyword search of each episode title and description across both Apple’s global database (i.e. accessible in countries outside of China), its domestic database (post-censorship for Chinese

²²Interview with Podcast producer C, mixed English and Chinese. May 2026.


audiences), as well as on Xiaoyuzhou. I search each database for the leader’s full name in both simplified and traditional Chinese, and measure the raw number of matching episodes. It is important to reiterate that Apple is a podcast client not a platform, while Xiaoyuzhou is both a platform and a client. Concretely, this means that Apple removes podcasts *after* they have been released (client-side), while Xiaoyuzhou performs both self-censorship (platform-side) as well after release (client-side). Platform-specific equilibria would be consistent with diminished counts of leaders’ names on Xiaoyuzhou’s storefront.

The second test of platform equilibria is whether regime-compliant platforms promote non-political content to podcast audiences. Because this test is audience side, I turn to the Apple podcast client, and compare its storefront recommendations across six distinct storefronts worldwide over a 14-day period in 2026. For cross-national comparability, I use the Apple store’s self-assigned “politics” and “news” labels, and assess whether Chinese-language political shows are recommended in the Chinese storefront at a lower rate than political shows in other countries’ storefronts. While purely descriptive, lower political recommendations on the PRC storefront would be consistent with regime objectives to diminish the salience of political content. Client-side promotion of non-political podcasts should be detected in diminished promotion of political podcasts on the Chinese storefront only, and would be consistent with a platform-specific equilibrium.

5. Overt Censorship and Compositional Shift

Foreign podcast platform adoption rate in the wake of the two censorship events slowed and reversed after both censorship shocks. In Figure 1, Panel A1 depicts the change in platform adoption rate for simplified Chinese podcasts on AnchorFM before and after its block in June 2020, and Panel A2 depicts the change for Firstory before and after its block in June 2021. In the months leading up to June 2020, AnchorFM had been among the fastest growing host platforms selected by Chinese-language creators, with 25 new podcasts in April and 55 new podcasts in May; after the block its adoption rate reversed and began to decrease ($\tau = -2.4$,

$p < 0.001$, see Panel B2). During the June 2020 block Firstory, based out of Taiwan, which at the time was not yet blocked by the Great Firewall, continued to grow; its growth reversed in June 2021 after its own blockage ($\tau = -0.70$, $p < 0.001$; Panel B2). Level changes, which tend to be noisier as podcast creators may begin adopting platforms in advance of their first published episode, show no significant change in three of the four specifications (See Panel B1).



(A) NEWLY PUBLISHED PODCASTS BY MONTH. A1=ANCHORFM; A2=FIRSTORY. GREY AND BLACK LINES SHOW LINEAR AND QUADRATIC FITS. BLOCK AT DASHED LINE.

(B) ESTIMATED RDIT COEFFICIENTS WITH 95% CONFIDENCE INTERVALS. B1=LEVEL CHANGE; B2=SLOPE CHANGE.

FIGURE 1. EFFECT OF BLOCK ON PLATFORM GROWTH

These reversals stand in contrast to the growth rates of the two largest Chinese-language platforms during this period, Xiaoyuzhou (founded in Mar 2020) and Ximalaya. Figure 2 shows the raw volume of monthly new podcasts started on Xiaoyuzhou (Panel C1) and Ximalaya (Panel C2), with the ecosystem-wide and AnchorFM censorship event (June 2020) and Firstory blockage (June 2021) events depicted at the dashed lines. While both platforms' adoption rates declined briefly during June 2020, growth quickly rebounded in both cases and, in Xiaoyuzhou's case, grew rapidly through the first half of 2022. These results are consistent with new platform growth being captured by domestic alternatives in the wake of overt censorship.

Blockage was also associated with a pronounced decline in the observed political content on the ecosystem, averaged across all episodes. Figure 3, Panel A shows the month-to-month decline in political content in the podcast ecosystem, with bootstrapped monthly confidence





FIGURE 2. DOMESTIC PRC PLATFORM GROWTH (JUN 2019–JUN 2022). FOREIGN PLATFORM BLOCKS AT DOTTED LINES.

intervals, while Panel B shows the likelihood that another month during this period would have shown a greater 3-month pre/post change in political content. The results of 51 placebo draws show that the decline in political content (0.00666) is unlikely to be observed due to chance ($p = 0.04$). As a robustness check, I also check whether the self-assigned political labels for active podcasts decline during this period and find that they do.



(A) DECLINE IN MEAN SIMILARITY TO BUNDLE OF POLITICAL PODCASTS. CENSORSHIP AT DOTTED LINE.

(B) SIMULATED 3-MONTH PRE/POST DECLINES WITH PLACEBO INTERVENTION MONTHS. OBSERVED DECLINE IN RED.

FIGURE 3. POLITICAL-CONTENT DECLINE AND PLACEBO TEST.

Figure 4 reveals that the majority of the observed change in pre/post means are attributable to exit and entry, as outlets that experience the censorship event gradually exit the producer pool and new entrants, who never experienced the censorship event, enter. On average across 1000 simulations, “exiters” contribute 0.0018 political points to the overall pool of content that, after their departure, is no longer available; entrants, meanwhile, are -0.0055 less political than the pool of shows to which they are entering. Combined, these

entries and exits reshape the composition of podcasts in the ecosystem.




FIGURE 4. CONTRIBUTION TO OBSERVED DECLINE IN POLITICAL CONTENT. CIs CONSTRUCTED BY RESAMPLING FROM WITHIN EACH POOL.

When post-shock exit is regressed on the demeaned and scaled political score (controlling for the number of previously published episodes), a show’s “political-ness” emerges as a significant predictor of podcast cessation, with a one standard deviation in political score resulting in a 44.4% increase in odds of exiting ($p < 0.001$) (See Figure 5). Over time, as new entrants are less political than exiting podcasts, the ecosystem becomes increasingly dominated by entertainment and lifestyle shows. While we do observe some within-podcast change over time, it is a smaller contribution to the observed decline than ecosystem-level compositional change.

6. Platform Equilibria are Different

PRC-operated platforms contain shows that are systematically different than shows on non-PRC operated platforms. Figure 6 shows the raw count of Chinese-language podcasts or episodes that mention the last three leaders of China appearing on Apple’s global storefront (a podcast client), its mainland-specific storefront, and on Xiaoyuzhou alone. Despite both Apple and Xiaoyuzhou performing censorship on behalf of the Party, Xiaoyuzhou’s podcasts appear to be much more conservative in their mentions of PRC leaders: only nine episodes




FIGURE 5. PROBABILITY OF EXIT POST-CENSORSHIP BY PRE-CENSORSHIP POLITICAL SCORE (Z-SCORED). BINNED BY DECILE.




FIGURE 6. MENTION OF CHINESE LEADERS, BY PODCAST CLIENT.

or podcasts in the entire dataset mention either Jiang Zemin or Hu Jintao, and only 133 mention President Xi (of which several are propaganda broadcasts). Apple’s podcast store — despite being subject to the same PRC censorship — is much more likely to display shows or episodes mentioning state leaders. These numbers are not driven by the size of the dataset: while Apple’s Chinese-language global storefront is roughly six times bigger than Xiaoyouzhou, its mainland storefront is only 2.5x as big, but contains over 10 times as many episodes mentioning state leaders. These results are consistent with platform-specific equilibria, although whether this equilibria is driven by platform moderation or podcast creator self-selection will be explored in a follow-up article to this one.

Table 1 shows that comparing Apple’s Chinese mainland storefront to other countries’ reveals that Apple’s China store is less likely to recommend political podcasts (determined by self-assigned category), even relative to the Chinese-speaking regions of Hong Kong and Taiwan. As an additional check, I also evaluate whether political or news-related podcasts episodes (as determined by political score) are likely to get fewer listeners on Xiaoyuzhou — even when controlling for the producer of the podcast itself — and find that they do (See Table 2). To be clear, these associations could be driven by audiences, i.e. political-leaning episodes could be less interesting to viewers, or by platforms, i.e. platforms systematically push users away from political content, thus rendering them less popular and worthy of promotion. Rather, these results merely suggest that political content on PRC-compliant platforms appears less promoted (relative to other countries) and less consumed (relative to less political shows produced by the same creator).

Extension: Global Implications

One implication of overt censorship — and the corresponding ascendancy of domestic platforms — is that an increasing percentage of global Chinese-language communication is occurring on domestic PRC media platforms, including communication from media producers and audiences physically located outside Chinese territory. Unlike mainland audiences, diasporic

TABLE 1. POLITICAL CONTENT IN APPLE PODCAST RECOMMENDATIONS BY STOREFRONT (MAY 2026)

Storefront	Total Recommended	Political %
United States	182	19.3
Hong Kong	189	11.6
Taiwan	189	10.9
South Korea	71	8.5
Japan	82	5.3
China	133	3.0

TABLE 2. PODCAST POLITICALNESS AND EPISODE PLAY COUNT

Dependent Variable:	Play count (logged)
Model:	(1)
<i>Variables</i>	
Political score	-0.9878*** (0.2854)
Podcast age (logged days)	-0.2143*** (0.0154)
Previous podcast episodes (logged)	0.3236*** (0.0276)
<i>Fixed effects</i>	
Podcast ID	Yes
<i>Fit statistics</i>	
Observations	178,091
R ²	0.88283
Within R ²	0.01676
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>	

Chinese-language audiences can access uncensored global platforms, such as Apple’s global Chinese-language storefront. My data reveals that increasing volumes of overseas Chinese-language users are joining Xiaoyuzhou. Since mid-2022, Xiaoyuzhou has required all users to report their IP address on the platform. Recording new user sign-ups by their first recorded comment on a podcast episode, Figure 7 shows the cumulative volume of unique overseas Xiaoyuzhou users in my dataset between 2022 and 2026. Starting at just 1,310 users in November 2022, by December 2025 the figure had grown to 37,928. While not all of these users entered the platform at the time of their first comment — my data is only a portion of all Xiaoyuzhou episodes and its comments — these results suggest that domestic platforms, which are subject to the PRC’s distributed censorship regimen, are gaining market share among the global population of Chinese-language speakers — with implications for censorship beyond PRC borders.




FIGURE 7. CUMULATIVE DIASPORIC USER BASE, 2022-2025.

7. Discussion

This paper has argued that authoritarian regimes can exchange short-term backlash for long-term compliance by shifting ecosystem growth onto digital platforms under more complete

delegated control by the regime. One important mechanism for this shift is selective entry and exit: media producers who operate against the regime’s interests — in this case, by creating political content — exit the media ecosystem, to be replaced by more compliant producers who never experienced a freer alternative. Platform-specific governance, including the possibility of self-censorship on behalf of the regime, may help to ensure that media production on regime-controlled platforms operate in accordance with regime objectives. Over time, the regime achieves content-level information control at low marginal cost.

I offer this compositional account to help explain the stubborn persistence of overt censorship among high-capacity regimes who, theory would suggest, can achieve similar strategic outcomes with more sophisticated technological alternatives, such as shadow-banning, throttling, and flooding (Roberts 2018). My contention is that overt censorship can and does operate alongside these other strategies, but may be particularly useful when a regime operates with a longer time horizon and is willing to suffer short-term backlash.

This analysis bears important limitations. For one, isolating the causal effect of overt censorship on ecosystem composition is difficult in a setting with a single major shock and a small number of affected platforms; while I can partially address these concerns with data triangulation and placebo tests, the results from a regression discontinuity-in-time still rely on an assumption of local exogeneity: that no other shock affected the treated platforms’ birth rates at the cutoff. An additional concern is that ecosystem growth onto domestic Chinese platforms may partly reflect product quality. The PRC’s consumer software market is generally recognized as being quite competitive, and Xiaoyuzhou in particular is, by many accounts, a thoughtfully designed app specifically tailored for the Chinese market. While my argument may nevertheless hold — Xiaoyuzhou and domestic apps can still benefit from the blocking of overseas competitors even as they offer impressive products on their own terms — the specific attraction of native-Chinese platforms is likely an important contributor to their observed performance over time.

Finally, this paper restricts its empirical attention to the Chinese-language ecosystem,

but the compositional mechanism almost certainly extends beyond regime borders. One implication of authoritarian “reshoring” of digital communication is that an increasing percentage of global speakers use regime-controlled platforms to speak to markets held captive behind regime information controls. As a consequence, digital audiences outside the regime’s territorial borders often channel their speech through regime-controlled platforms, rendering regime information control effectively transnational. This possibility may suggest that foreign borders offer less protection for digital communication than is commonly assumed, and transnational repression and cooptation — the regime’s reach outward into foreign jurisdictions — may be occurring alongside the reverse migration, of foreign audiences into regime control via digital platforms. This digital reshoring is an important subject for future research on information controls within a global communication environment.

Appendix

Table of Contents

8	Podcast Data Gathering	34
9	Robustness Check: Self-labels	39
10	Podcast Embedding Anchors: Politics and Entertainment	39

8. Podcast Data Gathering

Search terms span eleven semantic categories, including generic podcast terminology, simplified and traditional content categories, format descriptors, named popular shows from mainland, Taiwanese, and Hong Kong markets, the 100 highest-frequency Chinese characters, Cantonese-specific terms, and politically sensitive topics.

Search terms:

播客 (podcast); podcast (podcast); 博客 (blog); 音频 (audio); 音 (audio); 广播 (broadcasting; radio); 播 (broadcasting; radio); 电台 (radio station); 台 (radio station); FM (FM radio); 节目 (program; show); 目 (program; show); 频道 (channel); 道 (channel); 声音 (sound; voice); 音 (sound; voice); 新闻 (news); 科技 (technology); 商业 (business); 财经 (finance and economics); 教育 (education); 历史 (history); 文化 (culture); 艺术 (arts); 音乐 (music); 电影 (film; movies); 游戏 (games); 体育 (sports); 健康 (health); 生活 (life; lifestyle); 旅游 (travel); 美食 (food; cuisine); 育儿 (parenting; child-rearing); 心理 (psychology; mental health); 情感 (emotions; relationships); 故事 (stories); 小说 (fiction; novels); 悬疑 (mystery; suspense); 喜剧 (comedy); 访谈 (interviews); 脱口秀 (talk show); 创业 (entrepreneurship; startups); 投资 (investment); 职场 (workplace; careers); 学习 (learning; study); 读书 (reading; books); 社会 (society); 经济 (economy; economics); 政治 (politics); 法律 (law); 医学 (medicine); 哲学 (philosophy); 宗教 (religion); 灵修 (spirituality); 冥想 (meditation); 瑜伽 (yoga); 时尚 (fashion); 美妆 (beauty; makeup); 汽车 (cars; automotive); 房产 (real estate); 军事 (military); 航空 (aviation); 太空 (space); 环保 (environmental protection); 新 (news); 科技 (technology); 商 (business); (finance and economics); 教育 (education); 史 (history); 文化 (culture); (arts); 音 (music); 影 (film; movies); (games); 育 (sports); 健康 (health); 生活 (life; lifestyle); 旅 (travel); 美食 (food; cuisine); 育 (parenting; child-rearing); 心理 (psychology; mental health); 情感 (emotions; relationships); 故事 (stories); 小 (fiction; novels); 疑 (mystery; suspense); 喜 (comedy); (interviews); 口秀 (talk show); (entrepreneurship; startups); 投 (investment); (workplace; careers); (learning; study); (reading; books); 社 (society); (economy; economics); 政治 (politics); 法律

(law); (medicine); 哲 (philosophy); 宗教 (religion); 修 (spirituality); 冥想 (meditation); 瑜伽 (yoga); 尚 (fashion); 美 (beauty; makeup); 汽 (cars; automotive); 房 (real estate); 事 (military); 航空 (aviation); 太空 (space); 保 (environmental protection); 聊天 (chat); 对话 (dialogue; conversation); (dialogue; conversation); 谈话 (conversation; talk); (conversation; talk); 讲座 (lecture; talk); 座 (lecture; talk); 课程 (course); 程 (course); 分享 (sharing); 日更 (daily updates); 周更 (weekly updates); 更 (weekly updates); 独白 (monologue); 白 (monologue); 纪录 (documentary; record); (documentary; record); 专访 (special interview); (special interview); 深度 (in-depth); 闲聊 (casual chat); 聊 (casual chat); 夜谈 (nighttime talk); 夜 (nighttime talk); 晨间 (morning); 晨 (morning); 日谈公园 (Ritan Park; Daily Talk Park); 故事FM (Stories FM); 大内密谈 (Palace Insider Chats; Inner-Court Secret Talks); 忽左忽右 (Left and Right; Swaying Left and Right); 声东击西 (Feint East, Strike West); 得意忘形 (So pleased as to forget oneself); 剩余价值 (Surplus Value); 随机波动 (Random Fluctuations); 文化有限 (Limited Culture); 不合时宜 (Untimely; Out of Step); 海马星球 (Seahorse Planet); 小声喧哗 (Quiet Clamor); 凹凸电波 (Bumpy Radio Waves); 展开讲讲 (Unfold and Talk It Through); 无聊斋 (Boredom Studio); 跳岛FM (Island-Hopping FM); 杯弓蛇影 (Bow in a Cup Mistaken for a Snake); 问题青年 (Problem Youth); 东亚观察局 (East Asia Observation Bureau); 百果 (Bailingguo; Bilingual News); 股癌 (Stock Cancer); 台通 (Taiwan Connect); 力 (Mario); 好味小姐 (Miss Good Taste); 敏迪 (Mindi's Selected Readings); 角 (Corner International); 法客台 (Law Guest Radio; Faka Radio); 少年兄 (Formosa Brothers); 前哨站 (Reading Outpost); 天下 (Listen to the World); 者 (The Reporter); 鬼之音 (Ghost Island Media); 五夜 (Five Nights Forum); 自由 (Free Wind); 早晨 (Good Morning; Early Morning); 的 (possessive or nominal particle); 一 (one); 是 (to be); 不 (not); 了 (completed-action particle); 在 (at; in; currently); 人 (person; people); 有 (to have; there is); 我 (I; me); 他 (he; him); 这 (this); (classifier; individual); 们 (plural suffix); 中 (middle; China); 来 (come); 上 (up; on); 大 (big; large); 为 (for; as; to do); 和 (and; with); 国 (country; nation); 地 (ground; adverbial particle); 到 (arrive; to); 以 (by; with; in order to); 说 (speak; say); 时 (time; when); 要 (want; need; will); 就 (then;

just); 出 (go out; produce); 会 (can; will; meeting); 可 (can; may); 也 (also); 你 (you); 对
 (correct; toward); 生 (life; to be born); 能 (can; able); 而 (and; yet); 子 (child; suffix); 那
 (that); 得 (obtain; complement particle); 于 (at; in; to); 着 (ongoing-action particle); 下
 (down; below); 自 (self; from); 之 (of; it); 年 (year); 过 (pass; experienced-action particle);
 发 (send; issue; develop); 后 (after; behind); 作 (do; make; work); 里 (inside); 用 (use); 道
 (way; path; say); 行 (go; okay; line); 所 (place; that which); 然 (so; thus); 家 (home; fam-
 ily); 种 (kind; type); 事 (matter; thing); 成 (become; accomplish); 方 (direction; method);
 多 (many; much); 经 (pass through; classics); 么 (interrogative suffix); 去 (go); 法 (law;
 method); 学 (learn; study); 如 (as; like); 都 (all; both); 同 (same; with); 现 (present; now);
 当 (when; serve as); 没 (not have; no); 动 (move); 面 (face; surface); 起 (rise; start); 看
 (see; watch); 定 (decide; fixed); 天 (sky; day); 分 (divide; minute; part); 还 (still; also;
 return); 进 (enter; advance); 好 (good); 小 (small); 部 (part; section); 其 (its; their); 些
 (some); 主 (main; host; master); 样 (kind; appearance); 理 (reason; principle); 心 (heart;
 mind); 她 (she; her); 本 (root; book; this); 前 (before; front); 开 (open; start); 但 (but);
 因 (cause; because); 只 (only; measure word); 从 (from; follow); 想 (think; want); 实 (real;
 solid); (this); (classifier; individual); (plural suffix); (come); (for; as; to do); (country;
 nation); (speak; say); (time; when); (can; will; meeting); (correct; toward); 於 (at; in;
 to); 著 (ongoing-action particle; to wear); (pass; experienced-action particle); (send; issue;
 develop); 後 (after; behind); (inside); (kind; type); (pass through; classics); (interrog-
 ative suffix); (learn; study); (present; now); (when; serve as); (not have; no); (move);
 (still; also; return); (enter; advance); (kind; appearance); (open; start); (from; follow);
 (real; solid); 听 (listen); (listen); 讲 (speak; lecture); (speak; lecture); 谈 (talk); (talk); 读
 (read); (read); 写 (write); (write); 问 (ask); (ask); 话 (speech; words); (speech; words);
 语 (language; speech); (language; speech); 声 (sound; voice); (sound; voice); 音 (sound;
 music); 影 (shadow; film); (Cantonese; Guangdong speech); (Cantonese; Yue language);
 古 (storytelling); (chat; talk); 偈 (chat; talk); 港式 (Hong Kong style); 香港人 (Hong Kong
 people); 本土 (local; native); (Yue; Cantonese); 港 (Hong Kong-produced); 政治 (politics);

民主 (democracy); 自由 (freedom); 人权 (human rights); 人 (human rights); 维权 (rights defense); (rights defense); 异见 (dissent); (dissent); 六四 (June Fourth); 西藏 (Tibet); 新疆 (Xinjiang); 香港 (Hong Kong); 台湾 (Taiwan); 台 (Taiwan independence); LGBT (LGBT); 同志 (LGBTQ; gay/queer; comrade); 性别 (gender); 性 (gender); 女权 (feminism; women's rights); 女 (feminism; women's rights); 基督教 (Christianity); 佛教 (Buddhism); 伊斯兰 (Islam); 伊斯 (Islam); 法轮功 (Falun Gong); 真实犯罪 (true crime); 真犯罪 (true crime); 悬疑推理 (mystery and deduction); 疑推理 (mystery and deduction); 恐怖 (horror); 灵异 (supernatural; paranormal); (supernatural; paranormal); 都市传说 (urban legends); 都市 (urban legends); 二次元 (anime/manga subculture; 2D fandom); 动漫 (animation and comics); 漫 (animation and comics); 虚拟偶像 (virtual idol); 偶像 (virtual idol); 加密货币 (cryptocurrency); 加密 (cryptocurrency); 区块链 (blockchain); (blockchain); 人工智能 (artificial intelligence); 机器学习 (machine learning); 器 (machine learning); 元宇宙 (metaverse); 电竞 (esports); (esports); 手游 (mobile games); 手 (mobile games); 喜马拉雅 (Ximalaya); 喜拉雅 (Ximalaya); 小宇宙 (Xiaoyuzhou); 荔枝 (Lizhi); 蜻蜓 (Qingting); 网易云 (NetEase Cloud); 易 (NetEase Cloud); 酷我 (Kuwo); 人 (Lanren Tingshu; Lazy Audio Books); 得到 (Dedao); 中文播客 (Chinese-language podcasts); 播客 (Sinophone Chinese-language podcasts); 国语播客 (Mandarin-language podcasts); 播客 (Mandarin-language podcasts); 普通话 (Mandarin; Putonghua); 普通 (Mandarin; Putonghua); 闽南语 (Southern Min; Hokkien); 南 (Southern Min; Hokkien); 客家话 (Hakka); 客家 (Hakka); 海外华人 (overseas Chinese); 海外人 (overseas Chinese); 留学生 (international students); 留生 (international students)

TABLE 3. APPLE PODCASTS STOREFRONTS SELECTED FOR CHINESE-LANGUAGE
PODCAST COLLECTION

Continent	Country or region	Apple storefront code
Asia	Taiwan	TW
Asia	Hong Kong	HK
Asia	China	CN
Asia	Macau	MO
Asia	Singapore	SG
Asia	Malaysia	MY
Asia	Japan	JP
Asia	South Korea	KR
Asia	India	IN
Asia	Saudi Arabia	SA
North America	United States	US
North America	Canada	CA
North America	Mexico	MX
South America	Venezuela	VE
Oceania	Australia	AU
Oceania	New Zealand	NZ
Europe	United Kingdom	GB
Europe	Germany	DE
Europe	France	FR
Europe	Netherlands	NL
Europe	Russia	RU
Africa	South Africa	ZA

9. Robustness Check: Self-labels

I also measure the percent of podcasts with self-assigned “politics” and “news” labels, according to the Apple metadata. The data reflect a similar decline in politics and news podcasts in June 2020.



FIGURE 8. SELF-REPORTED SHARE OF POLITICS AND NEWS SHOWS IN THE CHINESE-LANGUAGE ECOSYSTEM.

10. Podcast Embedding Anchors: Politics and Entertainment

See table on next page

TABLE 4. MATCHED CONTRASTIVE ANCHOR SHOWS

Source show	Translation	Source window	Contrastive anchor	Translation	Anchor window	Category fit
不明白播客	Not Under-standing Podcast	2022-05-24 to 2026-01-25	to 北咖啡在台	Nordic Coffee in Taiwan	2022-05-22 to 2026-01-23	to Leisure, coffee
404档案馆	404 Archive	2021-06-02 to 2026-01-23	to 中年少女坦白局	Middle-Aged Girls' Confession Session	2021-05-28 to 2026-01-23	to Leisure, relationships, fashion/beauty
亚太报道	Asia-Pacific Report	2022-05-02 to 2025-03-21	to 舒服的梳人生茶室	Comfortable Comb: Life Outlook Teahouse	2022-05-05 to 2025-03-25	to Lifestyle, personal growth
端 — 端媒新播客	Initium News — Initium Media News Podcast	2024-01-01 to 2026-01-28	to 客播 I 山背山海台九	Hakka Radio: Mountain, Sea, and Provincial Highway 9	2024-01-05 to 2026-01-28	to Travel, places
时事大家谈 - 美国之音	Current Affairs Forum – Voice of America	2024-11-26 to 2025-03-14	to 04x06 生活指南	04x06 Life Guide	2024-12-05 to 2025-03-06	to Parenting, lifestyle, mental health
纵横大历史	Grand History in Perspective	2020-03-11 to 2025-05-07	to 客普拉斯 x 透客	Hakka Plus x Spoiler Hakka	2020-03-24 to 2025-04-26	to TV/film, performing arts
王丹学堂	Wang Dan Academy	2020-10-10 to 2022-01-25	to 秘呀	Mom's Secrets	2020-10-05 to 2022-01-28	to Parenting, family
要闻解说	Top News Analysis	2026-01-05 to 2026-01-28	to 你本就值得	You Were Always Worthy	2026-01-05 to 2026-01-28	to Mental health, self-worth
Dennis 的全球政治	Dennis's Global Politics Notes	2020-10-16 to 2026-01-28	to 迷品	Fans of Eslite	2020-10-15 to 2026-01-27	to Arts, books
角新 Daily Podcast	Corner International News Daily Podcast	2022-03-01 to 2026-01-28	to 基础颜究编辑部	Basic Beauty Research Editorial Department	2022-03-05 to 2026-01-26	to Skincare, beauty/lifestyle

References

- Blaydes, Lisa. 2018. *State of Repression: Iraq under Saddam Hussein*. Princeton University Press.
- Bradshaw, Samantha and Philip N Howard. 2018. “The global organization of social media disinformation campaigns.” *Journal of International Affairs* (1.5):23–32.
- Byford, Sam. 2020. “Apple pulls podcast apps in China after government pressure.” *The Verge* . Accessed May 22, 2026.
URL: <https://www.theverge.com/2020/6/11/21287436/pocket-casts-castro-china-apple-government-pressure>
- Chan, Kelvin and Barbara Ortutay. 2025. “Iran asks its people to delete WhatsApp from their devices.” AP News. Accessed May 20, 2026.
URL: <https://apnews.com/article/iran-whatsapp-meta-israel-d9e6fe43280123c9963802e6f10ac8d1>
- Chang, Charles. 2021. “Information credibility under authoritarian rule: Evidence from China.” *Political Communication* 38(6):793–813.
- 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.
- China Cyberspace Network. 2019. “Cyberspace Administration of China Launches Special Rectification Campaign on Online Audio.” WeChat public account post. Source: China Cyberspace Network. Accessed 2026-05-17.
URL: <https://mp.weixin.qq.com/s/tRN1-hoDesI5nWTInWTw8Q>

- Corduneanu-Huci, Cristina and Alexander Hamilton. 2022. “Selective control: the political economy of censorship.” *Political Communication* 39(4):517–538.
- Creemers, Rogier. 2019. “Cyber China: Upgrading propaganda, public opinion work and social management for the twenty-first century.” *Chinese authoritarianism in the Information Age* pp. 15–30.
- Cyberspace Administration of China. 2017. “[Provisions on the Administration of Internet News Information Services].”. Effective June 1, 2017.
URL: https://www.cac.gov.cn/2017-05/02/c_1120902760.htm
- Danescu-Niculescu-Mizil, Cristian, Robert West, Dan Jurafsky, Jure Leskovec and Christopher Potts. 2013. No country for old members: User lifecycle and linguistic change in online communities. In *Proceedings of the 22nd international conference on World Wide Web*. pp. 307–318.
- DeButts, Matthew and Jennifer Pan. 2026. “Media Patronage: Selection and Survival to Shape Foreign Media.” *Work in progress* .
- DeMets, Sydney A and Emma S Spiro. 2025. “Podcasts in the periphery: Tracing guest trajectories in political podcasts.” *Social Networks* 82:65–79.
- Dimitrov, Martin K. 2023. *Dictatorship and information: Authoritarian regime resilience in communist Europe and China*. Oxford University Press.
- DiResta, Renee, Kris Shaffer, Becky Ruppel, David Sullivan, Robert Matney, Ryan Fox, Jonathan Albright and Ben Johnson. 2019. “The tactics & tropes of the Internet Research Agency.” .
- Dukalskis, Alexander, Saipira Furstenberg, Yana Gorokhovskaia, John Heathershaw, Edward Lemon and Nate Schenckan. 2022. “Transnational repression: Data advances, comparisons, and challenges.” *Political Research Exchange* 4(1):2104651.

- Earl, Jennifer, Thomas V Maher and Jennifer Pan. 2022. “The digital repression of social movements, protest, and activism: A synthetic review.” *Science Advances* 8(10):eabl8198.
- Egorov, Georgy, Sergei Guriev and Konstantin Sonin. 2009. “Why resource-poor dictators allow freer media: A theory and evidence from panel data.” *American political science Review* 103(4):645–668.
- Esberg, Jane and Alexandra A Siegel. 2021. “How Exile Shapes Online Opposition: Evidence from Venezuela.” *American Political Science Review* pp. 1–18.
- Fan, Yingjie, Jennifer Pan and Tongtong Zhang. 2025. “Gender and political compliance under authoritarian rule.” *Comparative Political Studies* 58(9):1890–1929.
- Fang, Kecheng. 2022. “What is Zimeiti? The commercial logic of content provision on China’s social media platforms.” *Chinese Journal of Communication* 15(1):75–94.
- Feldstein, Steven. 2021. *The rise of digital repression: How technology is reshaping power, politics, and resistance*. Oxford University Press.
- Frantz, Erica, Andrea Kendall-Taylor and Joseph Wright. 2020. “Digital repression in autocracies.” *Varieties of Democracy Institute Users Working Paper (27)* pp. 1–22.
- Freyburg, Tina and Lisa Garbe. 2018. “Blocking the bottleneck: Internet shutdowns and ownership at election times in sub-Saharan Africa.” *International Journal of Communication* 12:3896–3916.
- Gehlbach, Scott, Zhaotian Luo, Anton Shirikov and Dmitriy Vorobyev. 2026. “Is there really a dictator’s dilemma? Information and repression in autocracy.” *American Journal of Political Science* 70(1):381–395.
- Gillespie, Tarleton. 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.

- Gohdes, Anita R. 2015. "Pulling the plug: Network disruptions and violence in civil conflict." *Journal of Peace Research* 52(3):352–367.
- Gohdes, Anita R. 2020. "Repression technology: Internet accessibility and state violence." *American Journal of Political Science* 64(3):488–503.
- Gohdes, Anita R. 2023. *Repression in the digital age: Surveillance, censorship, and the dynamics of state violence*. Oxford University Press.
- Gueorguiev, Dimitar. 2021. *Retrofitting leninism: Participation without democracy in China*. Oxford University Press.
- Guriev, Sergei and Daniel Treisman. 2019. "Informational autocrats." *Journal of economic perspectives* 33(4):100–127.
- Hassan, Mai, Daniel Mattingly and Elizabeth R Nugent. 2022. "Political control." *Annual Review of Political Science* 25(1):155–174.
- Hassanpour, Navid. 2014. "Media disruption and revolutionary unrest: Evidence from Mubarak's quasi-experiment." *Political Communication* 31(1):1–24.
- Hern, Alex. 2020. "Apple removes two podcast apps from China store after censorship demands." *The Guardian* . Accessed 2026-05-17.
URL: <https://www.theguardian.com/technology/2020/jun/12/apple-removes-two-podcast-apps-from-china-store-after-censorship-demands>
- Hoang, NP., AA. Niaki, J. Dalek, J. Knockel, P. Lin, B. Marczak, M. Crete-Nishihata, P. Gill and M. Polychronakis. 2021. How Great is the Great Firewall? Measuring China's DNS Censorship. In *30th USENIX Security Symposium*. USENIX Association pp. 3381–3398.
URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/hoang>
- Hobbs, William R and Margaret E Roberts. 2018. "How sudden censorship can increase access to information." *American Political Science Review* 112(3):621–636.

- Horz, Carlo M. 2024. "Identity propaganda." *British Journal of Political Science* 54(2):313–338.
- Howard, Philip N, Sheetal D Agarwal and Muzammil M Hussain. 2011. "When do states disconnect their digital networks? Regime responses to the political uses of social media." *The Communication Review* 14(3):216–232.
- Huang, Vincent Guangsheng and Hongyu Wu. 2023. "Contingent symbiosis: News start-ups and local cyberspace administration in contemporary China." *Information, Communication & Society* 26(3):512–529.
- Jansen, Sue Curry and Brian Martin. 2015. "The Streisand effect and censorship backfire." *International Journal of Communication* 9:16–16.
- Kalla, Joshua L and David E Broockman. 2018. "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments." *American Political Science Review* 112(1):148–166.
- King, Gary, Jennifer Pan and Margaret E Roberts. 2013. "How censorship in China allows government criticism but silences collective expression." *American political science Review* 107(2):326–343.
- King, Gary, Jennifer Pan and Margaret E. Roberts. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review* 111(3):484–501.
- Kuo, Lily. 2018. "No joke: have China's censors gone too far with ban on humour app?" *The Guardian* . Accessed May 17, 2026.
URL: <https://www.theguardian.com/world/2018/apr/21/no-joke-have-chinas-censors-gone-too-far-with-ban-on-humour-app>
- Lakhani, Kalsoom. 2020. "Why Pakistan's TikTok ban is a bad sign for investors." Atlantic Council. Accessed May 20, 2026.

URL: <https://www.atlanticcouncil.org/blogs/new-atlanticist/why-pakistans-tiktok-ban-is-a-bad-sign-for-investors/>

Lazarsfeld, Paul F, Bernard Berelson and Hazel Gaudet. 1968. *The people's choice: How the voter makes up his mind in a presidential campaign*. Columbia University Press.

Link, Perry. 2002. "The anaconda in the chandelier: Chinese censorship today." *The New York Review of Books* 49(6):1230–1254.

Litterer, Benjamin Roger, David Jurgens and Dallas Card. 2025. Mapping the podcast ecosystem with the structured podcast research corpus. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 25132–25154.

Lorentzen, Peter. 2014. "China's strategic censorship." *American Journal of political science* 58(2):402–414.

Lu, Yingdan, Jack Schaefer, Kunwoo Park, Jungseock Joo and Jennifer Pan. 2022. "How Information Flows from the World to China." *The International Journal of Press/Politics* p. 19401612221117470.

Lu, Yingdan and Jennifer Pan. 2022. "The pervasive presence of Chinese government content on Douyin trending videos." *Computational Communication Research* 4(1).

Lu, Yingdan, Jennifer Pan, Xu Xu and Yiqing Xu. 2025. "Decentralized Propaganda in the Era of Digital Media: The Massive Presence of the Chinese State on Douyin." *Available at SSRN 5145803*.

Luqiu, Luwei Rose. 2024. Unleashing voices: How uncensored feminist podcasts broaden the discourse on gender issues in Mainland China. In *New Digital Feminist Interventions*. Routledge pp. 114–131.

- Meng, Jing and Haiyan Wang. 2025. “Podcasting as alternative media: Navigating politics, platform power, and journalistic professionalism in China.” *Digital Journalism* pp. 1–20.
- Nichter, Simeon. 2008. “Vote buying or turnout buying? Machine politics and the secret ballot.” *American political science review* 102(1):19–31.
- Ong, Jonathan Corpus and Jason Vincent Cabañes. 2019. “Politics and profit in the fake news factory: Four work models of political trolling in the Philippines.”
- Pan, Jennifer. 2017. “How market dynamics of domestic and foreign social media firms shape strategies of internet censorship.” *Problems of Post-Communism* 64(3-4):167–188.
- PodFest China. 2020. 2020 Survey: A Look at Chinese Podcast Listeners. Technical report PodFest China.
- Qin, Amy. 2019. “In China, a Podcast Inspired by This American Life Gives Voice to the Real.” *The New York Times* .
URL: <https://www.nytimes.com/2019/05/12/world/asia/china-podcast-gushi-fm.html>
- Rasul, Muhammad Ehab, Saifuddin Ahmed, Jaeho Cho and Teresa Gil-Lopez. 2025. “From podcasts to protests: Examining the influence of podcasts and misinformation on contentious political participation.” *Journal of Broadcasting & Electronic Media* 69(3):273–293.
- Roberts, Margaret E. 2018. Censored. In *Censored*. Princeton University Press.
- Rodman, Emma. 2020. “A timely intervention: Tracking the changing meanings of political concepts with word vectors.” *Political Analysis* 28(1):87–111.
- Rydzak, Jan, Moses Karanja and Nicholas Opiyo. 2020. “Internet shutdowns in Africa—Dissent does not die in darkness: Network shutdowns and collective action in African countries.” *International Journal of Communication* 14:24–24.
- Sauer, Pjotr. 2022. “Russia bans Facebook and Instagram under ‘extremism’ law.” The Guardian. Accessed May 20, 2026.

URL: <https://www.theguardian.com/world/2022/mar/21/russia-bans-facebook-and-instagram-under-extremism-law>

- Shirikov, Anton. 2024. "Rethinking propaganda: how state media build trust through belief affirmation." *The Journal of Politics* 86(4):1319–1332.
- Stern, Rachel E and Jonathan Hassid. 2012. "Amplifying silence: uncertainty and control parables in contemporary China." *Comparative Political Studies* 45(10):1230–1254.
- Stockmann, Daniela and Mary E Gallagher. 2011. "Remote control: How the media sustain authoritarian rule in China." *Comparative political studies* 44(4):436–467.
- Stukal, Denis, Sergey Sanovich, Richard Bonneau and Joshua A Tucker. 2017. "Detecting bots on Russian political Twitter." *Big data* 5(4):310–324.
- Sullivan, John L. 2019. "The platforms of podcasting: Past and present." *Social media+ society* 5(4):2056305119880002.
- Truex, Rory. 2017. "Consultative authoritarianism and its limits." *Comparative political studies* 50(3):329–361.
- Truex, Rory. 2019. "Focal points, dissident calendars, and preemptive repression." *Journal of Conflict Resolution* 63(4):1032–1052.
- Wang, Yun and Xiaotong Lin. 2025. "The silent public: reconsidering online participation and public concern in a context of Chinese podcast community." *Chinese Journal of Communication* 18(1):21–38.
- Warner, Benjamin R, Qin Li, Guadalupe Madrigal, Breck Gillespie, Julia de Braganca Peixoto and Kelly Winfrey. 2026. "Podcasts, conservative identity politics, and the 2024 election." *Journal of Radio & Audio Media* 33(1):80–96.
- Wintrobe, Ronald. 2000. *The political economy of dictatorship*. Cambridge University Press.

- Woolley, Samuel C. 2020. "Bots and computational propaganda: Automation for communication and control." *Social Media and Democracy: The State of the Field, Prospects for Reform*. Cambridge pp. 89–110.
- Yang, Fan. 2023. "Feminist podcasting: A new discursive intervention on gender in Mainland China." *Feminist Media Studies* 23(7):3308–3323.
- Zheng, Wang. 2005. "'State feminism'? Gender and socialist state formation in Maoist China." *Feminist studies* 31(3):519–551.