Competition and Privacy

Yifei Wang[1]

[1]Graduate School of Business, University of Pittsburgh

January 17, 2026

Abstract

The question of whether competition enhances or impedes a firm's commitment to protecting consumer privacy has been debated in competition policy circles. This paper studies how a shift in competition affects the nature and type of a firm's intrusion on consumers' privacy. I study this question in the context of Android app markets in China, and measure privacy by examining apps' permission requests. I investigate a 2017 regulation in China that reduced competition in censored app categories by prohibiting censorship-circumvention tools commonly used to access foreign apps banned by the government. This regulatory change made banned apps less accessible and reduced competition faced by permitted apps in censored categories, but did not affect apps in uncensored categories. I documented empirically that the regulation substantially reduced the availability of VPNs and banned apps in China while increased market concentration in censored categories.

I use a synthetic differences-in-differences approach and compare permissions requested by permitted apps in censored categories to apps in uncensored categories before and after the regulatory change. By analyzing over 300,000 historical app installation packages, I show that reducing competition led to a significant increase in the number of permissions requested by apps. The increase is contributed by a wide range of permissions that have serious privacy implications. I show empirically that the removal of competition increases privacy-intrusive behavior by altering firms' incentives to differentiate and compete along the privacy dimension. The increase in privacy-intrusive behavior is due to treated apps' efforts to improve consumer engagement and monetize attention. It does not reflect the development of additional functions into their products or strategic response to entrants.

I am grateful to Catherine Tucker, Duncan Simester, and Birger Wernerfelt for their invaluable support and insightful comments on this paper. I would like to thank John Hauser, Juanjuan Zhang, Drazen Prelec, Dean Eckles, Rahul Bhui, Anthony Dukes, Garret Johnson, Lei Huang, Madhav Kumar, Jiaheng Yu, and Jerry Zhang for helpful suggestions.

1 Introduction

One of the frontier topics in competition policy is whether increased competition leads to more or less privacy for consumers. This has been argued in two landmark antitrust cases against two of the largest tech platforms – against Google and Meta. In both cases, the antitrust authority argued that their large size has led them to reduce privacy for consumers.[1] The U.S. Federal Trade Commission argued in its (amended) complaint against Facebook in 2021 that "without meaningful competition, Facebook has been able to provide lower levels of service quality on privacy and data protection than it would have to provide in a competitive market".[2] However, such claims are not supported by a marketing literature showing that larger firms collect less sensitive data than smaller firms (Kummer and Schulte, 2019), and that larger firms might be better positioned to follow procedures of informed consent than smaller firms (Campbell et al., 2015).

This paper attempts to provide some of the first empirical evidence to investigate whether a shift in competition affects firms' privacy intrusion on consumers. Empirically answering this question is challenging. The degree of competition in a market is often simultaneously influenced by factors that also affect how intrusively firms collect consumer data in that market, such as product features, market size, and policy regulations. Changes in competitive structures are infrequent in many markets and often coincide with other time-varying factors, making it challenging to isolate their causal effects.

In this paper, I investigate how market competition influences firms' privacy intrusion by studying an unexpected competition shock to Chinese app markets, introduced a major change in China's censorship circumvention policy. Due to Internet censorship, many popular apps in the international market, such as Google, Facebook, and Dropbox, are banned in mainland China. However, despite the censorship, Chinese users can still access banned apps using censorship circumvention tools, primarily virtual private networks (VPNs). Prior to 2017, using VPNs and similar tools for censorship circumvention was not formally regulated in China. However, on January 22nd, 2017, as part of its effort to strengthen Internet control, the Chinese government issued an official announcement restricting the provision and use of VPNs and similar cross-border Internet connection services. The new regulation targeted the use of such tools for bypassing censorship, prohibited those not approved by the government, and restricted their use to intra-firm communications only. The regulation was followed by an immediate enforcement campaign.

[1] https://www.justice.gov/opa/press-release/file/1328941/download

[2] https://www.ftc.gov/system/files/documents/cases/ecf_75-1_ftc_v_facebook_public_redacted_fac.pdf

Though the new regulation on censorship circumvention tools was not intended to change the competitive environment of Chinese app markets, it did so unexpectedly by further restricting users from accessing foreign products that were previously banned but still accessible through VPNs. In censored markets like news and chat apps, where many foreign players were banned, the number of available products suddenly declined. Foreign apps that were previously accessible through VPNs or similar tools became unavailable as the regulatory change prohibited the use of many such tools. Apps in these markets that were not banned by the government (henceforth 'permitted apps') therefore experienced an exogenous removal of competition. On the other hand, uncensored markets, such as banking apps and weather apps, were unaffected by the regulation. These markets were not subjected to extensive censorship as of January 2017, and did not have major foreign players that required circumvention tools to access. The difference in censorship across markets provides an opportunity to identify the causal effect of the competition change introduced by restrictions on censorship circumvention tools.

I document that the regulatory restrictions on censorship circumvention tools resulted in a notable increase in market concentration in censored app markets. The regulation led to a surge in VPN exit and increased difficulty to access previously censored apps. The average Herfindahl-Hirschman index in censored markets also increased substantially relative to uncensored markets, reflecting an increase in market power in censored markets following the regulatory change.

I then use a synthetic differences-in-differences (SDID) approach (Arkhangelsky et al., 2021) to identify the causal effect of reduced competition on the degree of privacy intrusion by apps in censored markets. I collect a unique dataset of over 300,000 historical app installation packages and examine the Android permissions requests made by permitted apps in censored markets ('treated apps') compared to those by apps in uncensored markets ('control apps') before and after the regulatory change. I find that, relative to the (synthetic) control group, treated apps request significantly more permissions, including more privacy-sensitive permissions, following the reduction in market competition. This difference is contributed by the increase in a wide range of permissions with serious privacy implications. The treatment effect is heterogeneous across markets and is the largest among communication, social, and video apps. Collectively, the findings suggest that the removal of competition resulted in a substantial increase in privacy-intrusive behavior among apps.

I investigate the mechanisms behind these findings and show that reduced competitive pressure limits consumer choice and weakens switching as a disciplining device. As switching becomes more constrained, firms have stronger incentives to aggressively collect data and request device access to

engage and monetize users. I find that the removal of competition had no substantial impact on the functional complexity of treated apps compared to synthetic control apps,

2 Literature Review

The findings of this paper build on a small but growing literature that investigates the effect of market competition on consumer privacy protection. Much of this literature takes a legal perspective and examines conceptually whether insufficient competition hurts consumer privacy protection and whether antitrust laws should be used to address this concern (Ohlhausen and Okuliar, 2015; Wasastjerna, 2018; Day and Stemler, 2019; Economides and Lianos, 2019). Complementing this legal perspective, a small number of economics and managerial research investigates how market structure affects firms' strategic decisions that have privacy implications, such as whether to adopt data-based personalization and price discrimination (Acquisti and Varian, 2005; Chen et al., 2022). The majority of this literature so far has been conceptual or theoretical in nature. One exception is perhaps Marotta-Wurgler (2016), who examines the privacy policies of 249 firms in seven markets. They conclude that firms in more competitive markets follow stricter data security standards. Despite the importance of their findings, Marotta-Wurgler (2016) relies on cross-market comparisons and does not address the endogeneity of cross-sectional differences in competition. This paper adds to the existing literature by empirically investigating the *causal* effect of market competition on privacy protection. Unlike previous studies that are either theoretical or rely on cross-market comparisons, this paper offers one of the initial empirical studies that identify the causal effect of competition by examining an exogenous competition shock introduced by a regulatory change.

This paper is related to, but different from, the literature on how privacy considerations and policies shape market competition. Theoretically, Casadesus-Masanell and Hervas-Drane (2015) show that consumers' privacy considerations can soften competition. Campbell et al. (2015) show that privacy regulations may intensify competition and are more likely to hurt small and new firms. Empirically, Johnson et al. (2023) and Peukert et al. (2022) find that the General Data Protection Regulation (GDPR) increased the concentration of the web-vendor market in favor of the big players. Janssen et al. (2022) find that the GDPR led to more exits and fewer entries in app markets. This paper also investigates the relationship between privacy and market competition, but from a different perspective. The key question of this paper is how competition changes affect privacy, instead of how privacy concerns or policies affect the competitive environment.

Finally, this paper also adds to the broader literature on the non-price effects of changes in competition. Past research has investigated the effect of changes in competition on a wide range of non-price outcomes, such as product quality (Mussa and Rosen, 1978; Banker et al., 1998; Goolsbee and Petrin, 2004; Matsa, 2011), product variety (Berry and Waldfogel, 2001; Sweeting, 2010; Fan, 2013), investment (Pindyck, 2007), innovation (Igami and Uetake, 2020; Cunningham et al., 2021), and risk taking (Hellmann et al., 2000; Jiang et al., 2018). This paper adds to the existing literature by addressing a previously unexplored aspect: the provision of privacy. The findings of this paper contribute to the understanding of the various consequences of changes in market competition.

Privacy is sometimes thought of as a dimension of product quality, which is extensively studied in the literature on non-price effects of competition. However, privacy differs from typical dimensions of product quality, such as battery life, in at least three ways. First, privacy is often seen as a fundamental human right (Warren and Brandeis, 1989; Economides and Lianos, 2019; Lin, 2022), rather than as a quality dimension that is determined solely by a producer's profit considerations. Second, determining the degree of privacy provision involves an inherent trade-off between protecting privacy and providing better personalization, both of which are valued by consumers. The costs associated with improving quality (such as extending battery life), on the other hand, are typically not directly recognized or valued by consumers. Third, while consumers can often observe and evaluate product quality, privacy is hard to evaluate (Goldfarb and Que, 2023). Consumers are often unaware of how their data is used and shared (Tsai et al., 2011; Kehr et al., 2015), and their privacy preferences are often uncertain and context-dependent (Spiekermann et al., 2001; Acquisti et al., 2013; Athey et al., 2017). These distinctions emphasize that privacy is an important outcome in its own right, instead of a straightforward extension of product quality.

3 Institutional Background

The landscape of Internet products and services in China is heavily influenced by governmental regulations. In this section, I discuss the development of Internet censorship within China and how it affects the competitive environment of the Chinese mobile app markets.

3.1 Internet Censorship in China

In 2003, China launched the Golden Shield Project, a nation-wide project for network security and information infrastructure.[3] The purpose of the project is "to improve the efficiency and crime investigation of the police system", and to "provide information sharing and support for public security services".[4] As part of the project, a set of legislative actions and information technologies, often referred to as 'the Great Firewall', was introduced and has since been used to block websites and content from foreign providers that are illegal or prohibited within mainland China (Clayton et al., 2006; Chen and Yang, 2019; Kerner, 2022). The blockage of websites and other digital content is operated through a set of sophisticated technologies, including Internet Protocol (IP) address blocking, Domain Name System (DNS) hijacking, and the Transmission Protocol (TCP) content filtering (Mou et al., 2016).

Internet censoring in China affects firms in the digital space. Some of the most popular digital products in the global market are inaccessible in China, such as Google, Facebook, YouTube, and Dropbox. In its report on foreign trade barriers in 2016, the U.S. Trade Representative listed China's Internet censorship as a trade barrier and reported that "eight of the top 25 most trafficked global sites" were blocked in China.[5,6] As of 2016, it was reported that 161 out of the top 1000 websites globally (according to the ranking by Alexa) were blocked in China by the Great Firewall (Chen and Yang, 2019). The impact of Internet censorship also extends to app markets. Apps for blocked websites are usually unavailable on domestic app stores, and cannot be used without censorship circumvention tools.

Notably, the degree of Internet censorship varies across markets. While some markets like communication and media apps are heavily censored with many major foreign players banned, other markets such as banking apps and weather apps are not subject to extensive Internet censorship during the period of this study. For simplicity, I refer to app categories that were subject to extensive Internet censorship before the regulatory change in 2017 as 'censored markets'and those where censorship does not impact major competitors—either due to the absence of major foreign players or because such players are permitted—as 'uncensored markets. I refer to apps in censored markets and blocked by the Great Firewall as the 'banned apps', and apps in censored markets but not

[3] http://www.gov.cn/zfjs/2006-11/17/content_445189.htm

[4] https://www.zgbk.com/ecph/words?SiteID=1&ID=43557&Type=bkzyb

[5] https://www.reuters.com/article/us-usa-china-trade-internet/u-s-says-china-internet-censorship-a-burden-for-businesses-idUSKCN0X50RD

[6] https://www.nytimes.com/2016/04/08/business/international/china-internet-controls-us.html

blocked as the 'permitted apps' throughout the rest of the paper.

Despite the Internet censorship, users in China can still access banned content and services with censorship circumvention tools. The predominant way to bypass Internet censorship is to use a virtual private network (VPN), which provides Internet traffic encryption and allows users within mainland China to tunnel to the Internet as if they were in a different country. VPNs can be used both on computers and on mobile phones. The easiest and most common way for a user to use a VPN is to obtain it from a VPN service provider.[7] The cost of VPN services varies widely, with options available from free providers like Green VPN to premium services such as Astrill VPN which costs around $12.5 per month. The demand to access banned websites and apps has led China to become one of the largest markets for VPNs globally. An 2014 report estimated that over 90 million Internet users in China have used VPNs to access restricted social media platforms.[8] More recently, a 2017 report by Statista showed that 31% of the surveyed Internet users in China reported to use a VPN in the past month.[9,10] Consistent with these statistics, my analysis in Section 3.3 finds that that 49 VPN apps placed in the top 1,000 of China's iOS apps in 2016 alone, which suggests a large and active VPN market among Chinese users.

3.2 The Regulation on Censorship Circumvention Tools

For many years, China did not formally regulate the use of VPNs and other circumvention tools, making it relatively easy for users to bypass Internet censorship and access banned apps and websites. However, on November 7, 2016, China's national legislature, the National People's Congress, passed the Cybersecurity Law of the People's Republic of China, signifying a major shift in the regulation of Internet access and control. The Cybersecurity Law is the first basic law that comprehensively regulates cyberspace security and related issues. While the law focused primarily on cybersecurity obligations and infrastructure, it also introduced "the principle of cyberspace sovereignty", which established a legal foundation for the state to exercise authority over China's cyberspace territory.

Following the Cybersecurity Law, on January 22, 2017, China's Ministry of Industry and Infor-

[7] A user can also set up their own VPN server, which is technically possible but requires a certain level of technological sophistication.

[8] https://blog.gwi.com/chart-of-the-day/90-million-vpn-users-in-china-have-accessed-restricted-social-networks/

[9] https://www.statista.com/statistics/301204/top-markets-vpn-proxy-usage/

[10] Notably, the estimate of VPN usage in China seems to be growing over time. For example, the 2012 U.S. congressional research report estimated only between 1% and 8% of Web users in China use proxy servers and virtual private networks to get around government-erected Internet firewalls. https://sgp.fas.org/crs/row/R42601.pdf

mation Technology (MIIT) issued the "Notification on Regulating the Market of Internet Connection Services".[11] The MIIT notification introduced a set of new regulations on international data corporation (IDC), Internet service provider (ISP), and content delivery network (CDN) markets. In particular, Section 2 Article 4 of the notification stated that self-constructed or leased cross-border Internet connections without the approval of China's telecommunication authorities would be prohibited. Notably, the notification explicitly mentioned VPNs as one (and in fact the only) example of such cross-border connection services. It further required telecommunication enterprises providing these services to maintain records of user profiles and to inform users that such services were restricted to intra-firm communication only. The MIIT announced an immediate campaign to enforce the new regulations and to "clean up" the market of Internet connection services. The enforcement campaign was originally expected to complete by March 31st, 2018, but was later extended to March 31st, 2019.[12]

The MIIT notification was the first legal document in China to explicitly regulate the provision and usage of VPNs. While the new regulation did not prohibit all VPNs, it significantly increased the difficulty of accessing and utilizing VPNs as a means of censorship circumvention. The new regulation was widely reported by domestic and International media, and had a substantial impact on Chinese VPN market.[13,14,15,16] For example, Bloomberg News reported that the government had directed China's three primary telecommunications providers to fully prohibit individual access to VPNs by February 2018.[17] Apple also announced it had removed nearly 700 VPN apps from the Chinese iOS App Store following the MIIT notification.[18] Additionally, since the introduction of the new regulation, several legal cases have been reported where individuals and firms were fined or jailed for selling VPN services.[19,20]

[11] http://www.cac.gov.cn/2017-01/23/c_1120366809.htm, last accessed on January 22th, 2023.

[12] https://www.miit.gov.cn/zwgk/zcwj/wjfb/txy/art/2020/art_f53db3415f884c038a77bb2754fb9ea2.html, the extension was announced on May 10th, 2018.

[13] https://www.sohu.com/a/124994917_126540

[14] https://www.scmp.com/news/china/policies-politics/article/2064587/chinas-move-clean-vpns-and-strengthen-great-firewall

[15] https://money.cnn.com/2017/01/23/technology/china-vpn-illegal-great-firewall/index.html

[16] https://www.science.org/content/article/science-suffers-china-s-internet-censors-plug-holes-great-firewall

[17] https://www.bloomberg.com/news/articles/2017-07-10/china-is-said-to-order-carriers-to-bar-personal-vpns-by-february#xj4y7vzkg

[18] https://cn.wsj.com/articles/CN-BGH-20180118112906

[19] https://www.theguardian.com/world/2017/dec/22/man-in-china-sentenced-to-five-years-jail-for-running-vpn

[20] http://www.jinmao.com.cn/cn/researchdetail.aspx?id=3406

Importantly, while VPNs also encrypt Internet traffic and protect online identity, their primary use for individuals in China is to bypass censorship and access restricted content.The censorship circumvention usage of VPNs is also the primary target of the 2017 regulation, as MIIT stated explicitly in the notification that VPNs were still allowed for (and only for) securing communications. This paper therefore takes the 2017 regulation on VPNs and similar tools as primarily a restriction on censorship circumvention, which affects the app markets through its impact on the competition between foreign and domestic players.

3.3 Regulatory Impact on Competition

The regulatory change on censorship circumvention tools limited access to banned foreign apps, unexpectedly reducing competition for permitted apps in censored categories. This paper investigates how this unexpected removal of competition affected privacy practice for apps in censored categories, compared to those in uncensored categories.

An important question is whether the new regulation was effectively enforced, and whether the enforcement meaningfully affected app market competition. Section 3.2 provided qualitative evidence for the MIIT enforcement campaign and the extensive media coverage on the removal of VPNs. The rest of this section discusses additional evidence that provides background information on how the regulatory change affected users' access to banned apps and thereby reshaped competition in censored markets.

The background evidence presented in this section used data on VPNs and banned apps listed on the Chinese iOS app store. This data was collected from an app analytics site that provided historical data on entry, exit, and reviews for iOS apps. Using iOS data to provide evidence for the competition shock was a practical choice due to data availability restrictions: the focal Android app store I study did not provide data on apps that exited the platform (such as many VPN apps), or apps that were banned by the government.[21] Although the iOS and Android app stores differ, this analysis focuses on the most popular VPNs and banned apps that are widely used by users across device types. Because both Android and iOS app stores operated under the same government regulations, an app banned on one platform would almost certainly be banned on the other. Therefore, availability on

[21]The Android app store I studied, as well as most other Chinese Android app stores, did not list apps that were banned by the government. App users needed to use VPNs to download these apps from non-Chinese app stores or websites. This meant that I could not directly observe whether users accessed these apps in data from Chinese Android app stores.

iOS reliably indicates availability on Android.

Figure 1 presents evidence on the accessibility of VPNs. In Figure 1a, I analyze a sample of 49 major VPN apps that appeared in the top 1000 China iOS app rank at least once in 2016. This sample represented the most popular VPNs among Chinese users prior to the regulatory change. The fact that 49 VPN apps placed in the top 1,000 of China's iOS apps in 2016 alone highlights a large and active VPN market among Chinese users, consistent with the statistics discussed in Section 3.1. Figure 1a shows the cumulative percentage of VPN apps that became inaccessible on the iOS App Store over time, removed either by their developers or by the app store. Figure 1a shows a sharp and discontinuous increase in VPN exit, starting from around five months after the regulatory announcement in January 2017. By March 2018, all VPN apps in this sample were unavailable.

The removal of leading incumbent VPNs would not necessarily affect consumers' access to banned apps if new VPN providers entered the market and survived. In Figure 1b, I further examine the impact of the regulatory change on new VPN entrants. Figure 1b plots the percentage of new entrants that exited within three and six months of their entry.[22] In 2016, most VPN entrants survived for a relatively long time - only about 6% exited within three months and 14% within six months of entry. After the regulatory change in 2017, exit rates jumped sharply to 50% at three months and 85% at six months, and have remained high ever since. Figure 1b indicates that regulatory enforcement displaced not only incumbent VPN providers but also new entrants soon after their entry. In fact, Figure A1 shows that over 95% of major VPNs in the market were inaccessible by September 2018, even after accounting for new entrants. Together, these evidence confirms that regulation was enforced and substantially affected the availability of censorship-circumvention tools.

To understand how the regulation on VPNs affected users' experiences with the banned apps, I analyze the performance of major banned apps in the Chinese iOS store. In Figure A2a, I analyze user reviews of a sample of 16 banned apps that ranked among the top 1000 on the China iOS app store in 2016.[23] These apps represent the most popular banned apps among Chinese users prior to the regulatory change, and were ranked highly despite being banned by the government. I use

[22]The sample comprises VPN apps ranked in the top 1,000 in China's iOS App Store at least once between 2016 and 2019, representing the most popular services during the study period. For the 2016 cohort, only apps entering before October 1, 2016 are included in the analysis to guarantee at least three months of pre-regulation observation for that cohort.

[23]Note that users did not need a VPN to review a banned app in the iOS App Store. To leave a review, users had to have previously downloaded the app, but downloading a banned app from the Chinese iOS app store also did not require a VPN. A VPN was only necessary when actually using the app after it had been downloaded.

Figure 1: Removal of Major Incumbent and Entrant VPNs

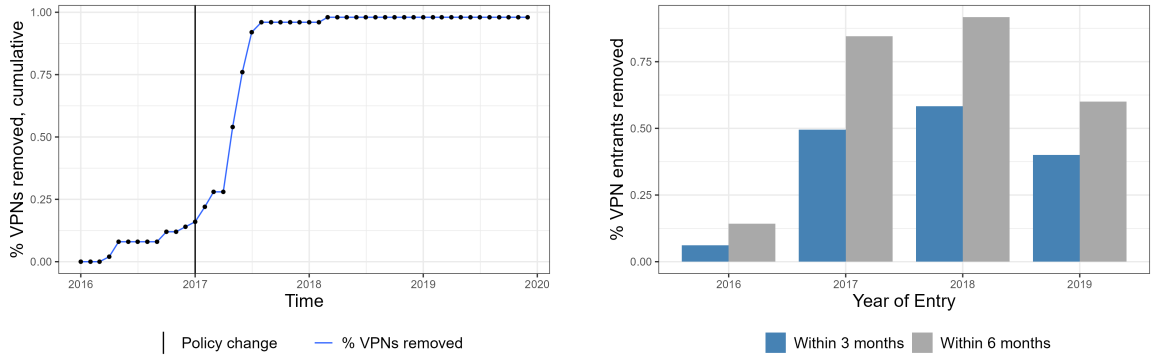

(a) Removal of Incumbent VPNs

(b) Removal of VPN Entrants

Notes: Panel (a) analyzes 49 VPN apps that appeared in the top 1000 China iOS app rank in 2016. This sample represents the most popular VPNs among Chinese users prior to the regulatory change. The plot shows the cumulative total of VPNs in this sample that became unavailable from the iOS app store over time, either because the app decided to exit the market or because it was removed by the platform. Panel (b) shows, for each year, the percentage of VPN entrants that were removed from the market shortly after entry. The sample used in this figure includes VPN apps that appeared at least once in the top 1000 China iOS app rank between between 2016 and 2019. This sample represents the most popular VPNs among Chinese users during this period. For each entry year, the figure plotted the percentage of new VPN entrants that became inaccessible - either due to removal by the app developer or by the app store - within 3 months and 6 months of their entrance date. To ensure at least a full three-month follow-up period for the 2016 cohort, only apps that entered before October 1, 2016 were included in this figure.

gpt-3.5-turbo to analyze whether a review indicates that the user had difficulty accessing or using the app (see details in Section A1.1). Figure A2a shows that the average percentage of user reviews mentioning inaccessibility rose substantially following the regulatory change, suggesting that the ban on VPNs created a significant barrier to accessing these apps. In Figure A2b, I further analyze the monthly rankings of banned apps within their relative category.[24] Figure A2b plots the share of banned apps that ranked in the top 300 of their category over time. The figure reveals a clear decline in their relative ranking performance compared to competitors within the same category following the regulatory change. The pattern is also robust to alternative cutoffs, such as the top 100 or top 200 in the respective category.

Finally, I check whether the regulatory change affected market concentration within censored categories. Although historical data on app downloads or usage were not directly observable, I collected the number of reviews[25] posted for each app that appeared in the top 300 of its iOS category each month from 2016 to 2020,[26] and used this as a proxy for monthly usage. I then calculated each app's market share within its category, as well as the market (category) level Herfindahl-Hirschman Index (HHI), based on this proxy measure. I classify categories into censored categories, where foreign apps were reported to be blocked by the Great Firewall, and uncensored categories, where no major foreign players were reported to be blocked, as detailed in Section 4.3. A differences-in-differences analysis shows that, after the regulatory change, the average HHI in censored markets rose by 146.3 points ($p = 0.016$) relative to uncensored markets. This estimate approaches the 200-point benchmark that the U.S. DOJ and FTC use to flag significant antitrust concerns, highlighting a substantial increase of market concentration in censored markets. In Figure 2, I plot the average Herfindahl–Hirschman Index over time for uncensored categories, and for censored categories with varying exposure to the regulatory change. Panel (a) of Figure 2 compares uncensored categories to more affected censored categories where multiple banned foreign apps appeared within the top 300 of their respective category rankings prior to 2017, and Panel (b) compares it to less affected censored categories where banned foreign apps did not appear in the top category rank. Figure 2 shows that the trend of the average HHI of the more affected categories closely resembled that

[24] This analysis includes banned apps that appeared at least once in the top 300 of their category during the study period, representing major players in their market.

[25] As ratings typically require less effort than reviews, the number of ratings is likely a better proxy for usage. However, the date when a rating was posted was not observable for ratings not associated with a review. I therefore use the number of ratings posted each month as the proxy for usage to estimate historical trends.

[26] The monthly Chinese iOS category ranking data were collected for the first day of each month from January 2016 to January 2020.

of the uncensored categories in 2016, but increased substantially after the regulatory change and diverged from the latter. This divergence did not occur for the less affected censored categories. The contrast between Panels (a) and (b) of Figure 2 confirms that the overall increase in HHI for censored categories reflects reduced competitive pressure from major foreign competitors, primarily in categories where those competitors had been more important players.

The analysis of market concentration provides qualitative evidence that the regulation on VPNs had a meaningful effect on competition in censored categories. However, the number of reviews is only an imperfect proxy for usage, and the relationship between usage and reviews is not explicitly modeled in this analysis. Therefore, this evidence should be interpreted as qualitative and suggestive rather than as a precise estimate of the regulatory change's impact on market concentration.

Figure 2: Policy Change Affected Market Concentration

(a) More Affected Categories

(b) Less Affected Categories

This figure shows the monthly Herfindal Hirschman Index (HHI) for each market (app category), calculated with the number of iOS user reviews as a proxy for app usage. A difference-in-differences regression analysis shows that the average HHI for censored markets increased by 146.3 points (p = 0.016) following the regulatory change, compared to uncensored markets, suggesting that the regulatory restrictions on censorship circumvention tools led to a notable increase in market concentration in censored markets.

3.3.1 Discussion

The background evidence discussed in this section confirms that the regulatory change on censorship circumvention tools significantly affected the accessibility of these tools and, consequently, the accessibility of banned apps. Since many of these banned apps were major competitors in the market, the regulatory shift reduced competition faced by permitted apps in censored markets and led to

increased market concentration.

The evidence presented in this section does not imply that access to banned apps was entirely eliminated for Chinese consumers. Although evidence suggested that access to major VPN apps in China was severely restricted, some smaller providers might not be monitored as closely and could have remained available for download, though likely with limited functionality. In addition, technically sophisticated users might still be able to bypass censorship by developing their own circumvention tools, which could be less detectable by regulators. The 2017 regulatory change should therefore be viewed as a policy intervention that substantially raised the barriers to accessing such tools, thereby significantly reducing - rather than completely eliminating - competitive pressure from banned foreign apps.

4 Data and Measurement

4.1 Data

I obtained the app data from a major third-party Android app store in China. As Google Play was banned in China, third-party app stores are a major source for Android users to access and download apps. The app store I studied is widely recognized as one of the top ten third-party Android app stores in China in terms of market share. The apps listed on this app store are therefore a representative sample of the Android apps available to Chinese users. My data includes apps listed under each category defined by the app store. A unique feature of this app store is that it provides historical versions of an app in addition to the latest version. My data includes descriptive information for the most recent version of each app at the end of every month during my study period, including its name, developer, size, and release date.[27] The data contains 17,001 apps listed on the Android app store and their 327,734 historical versions from September 2014 to August 2022.

Tracking privacy changes also requires historical data on the permissions requested by apps. Because the app store I studied provided historical versions of apps, I was able to collect permission used in these versions by analyzing their Android Package Kit (APK) files. The Android Package Kit is a downloadable package that includes the code and assets needed to install and run an app on an Android device. In particular, it includes the AndroidManifest.xml file (refered to as the Manifest file hereafter), which is a mandatory.xml file for Android apps to communicate information such as the permissions requested and the structure of basic app components to the Android operating

[27]The dataset contains the latest version of a month if an app released multiple versions in that month.

system.[28] I obtained 327,734 historical APK files, and after removing those that were broken or could not be properly processed, I extracted Manifest data for 327,146 app versions. On average, each app version requested 35.41 permissions during the period from September 2014 to August 2022.

I also complemented the permissions data with information on the permission group and protection level associated with each Android system-defined permission, collected from Android's developer guide.[29] A permission group refers to a category of functionally-related permissions.[30] The protection level is a categorical variable that reflects the potential privacy or security risk a permission poses to the user.[31] The Android system defines four main protection levels for permissions: normal, dangerous, signature, and signatureOrSystem.[32] Of the four types, 'normal' permissions are the only group Android considers to have minimal privacy risk, though it is widely recognized among developers that even 'normal' permissions can sometimes involve significant privacy risk. These permissions will also be granted automatically to apps when requested, while permissions of other protection levels will only be granted to apps after certification or user authorization. In the analysis that follows, I adhere to the classification by Android and use 'normal' vs. 'sensitive' permissions as the primary way to classify permissions.

I collected data on banned foreign apps from the GreatFire Analyzer. Operated by GreatFire.org, the GreatFire Analyzer is a web tool that has been monitoring websites and keywords censored by the Great Firewall of China since 2011. In particular, it tracks the censorship status of the Alexa top 1000 domains in China.[33] I collected the list of Alexa top 1000 domains that were banned prior to January 1st, 2017, as reported by the GreatFire Analyzer.[34] I matched the banned domains

[28] See here for the official description for the AndroidManifest file: https://developer.android.com/guide/topics/manifest/manifest-intro

[29] Source: https://developer.android.com/reference/android/Manifest.permission, accessed on December 30th, 2022

[30] https://developer.android.com/guide/topics/manifest/permission-group-element

[31] The protection level of a permission also affects the procedure the Android system follows when an app requests the permission.

[32] See here for an explanation of each permission level: https://developer.android.com/guide/topics/manifest/permission-element

[33] Source: https://en.greatfire.org/search/alexa-top-1000-domains, accessed on 31st October, 2022.

[34] The GreatFire Analyzer tests each of the target domains on a regular basis and reports whether the domain was banned at the time of the test. Most domains were either completely banned or not banned in all the tests, but a small number of domains had inconsistent test results. For those domains, I classify them as being 'banned' during the sample period if the domain was banned in five consecutive tests after the first blockage, and 'not banned' if otherwise.

with the names and categories of their associated Android apps by searching for and matching the keyword of each banned domain name in Google Play. I also supplement this list with websites and apps reported as being blocked in various media sources.[35, 36, 37]

4.2 Measuring Privacy Intrusion

4.2.1 Definition and the Primary Measure

I follow the definition of privacy from the seminal paper by Warren and Brandeis (1989), which defines it as "the right to be let alone." The right to privacy encompasses multiple dimensions. As discussed in Posner (1981), at least three key aspects are included within the general concept of privacy: (1) concealment of information, such as when individuals object to their personal information being shared in unwanted way; (2) freedom from intrusion or disturbance, such as when one complains about unsolicited phone calls or other disruptions violating their privacy; and (3) freedom and autonomy, which extends to issues such as the right to make personal decisions, including the legal classification of abortion rights under the right to privacy. This paper considers all three dimensions collectively, without explicitly distinguishing among them.

I measure privacy intrusion with app permission requests, and consider an increase in the number of permissions requested as the app becoming more privacy intrusive. Permission requests are widely used in the literature as a proxy of the degree of privacy intrusion by Android apps (Mylonas et al., 2014; Krafft et al., 2017; Kesler et al., 2017; Kummer and Schulte, 2019). In the Android environment, apps request permissions from the operating system in order to access users' data or control device functions. The number of permissions therefore represents the level of data access and device control requested by an app.

Requesting access to data or device controls does not necessarily mean that such access is always unwanted by the consumer. However, it generally increases the risk of unwanted access, as firms can potentially use the permissions in ways that consumers may not fully understand or agree with. For example, an app might collect location data for navigation purposes but also use it for ad targeting, or it may send an excessive number of push notifications, which can feel as intrusive as advertising or phone solicitations. This privacy risk is particularly important in the Android system, where

[35] https://zh.wikipedia.org/wiki/, accessed on January 29 2023

[36] https://www.vpnmentor.com/blog/the-complete-list-of-blocked-websites-in-china-how-to-access-them/, accessed on January 29 2023

[37] https://www.saporedicina.com/english/list-of-blocked-websites-in-china/, accessed on September 10, 2024

permission requests are one-off: during the period of my study, once a permission was granted, the app could access resources protected by that permission as much as it wished, unless consumers manually revoke that permission.[38] In fact, recent reports have revealed that well-known Chinese apps accessed users' location data up to 75 times per hour, even when the app wasn't actively in use.[39] Similar patterns have been observed and reported for other permissions as well.

A common misconception is that apps requesting more permissions are being more respectful of user privacy rather than intrusive, since they appear to seek consent before accessing data. This misconception likely comes from the assumption that some apps may exploit data without issuing any explicit permission request. However, this is not the case within the Android system. Permission requests are technical commands submitted to the Android system and are distinct from obtaining explicit user consent through privacy agreements or other forms of legal contracts. While apps can use data without obtaining explicit user consent through privacy agreements, they cannot bypass the system's permission requests to access the data. Most permissions are automatically certified and granted by the Android system without notifying the user. Only in rare cases, where permissions are deemed "dangerous" by Android (as was the case with 41 out of 754 permissions requested over 50 times in my sample), are users prompted to grant or deny a permission request. Therefore, the alternative to requesting a permission is not using data without informing the user but rather not accessing the data associated with the permission at all. Thus, requesting more permissions should be understood as an app accessing more data, which increases privacy risks for users, rather than indicating a greater respect for user privacy.

4.2.2 Alternative Measures

I also presented two alternative measures for privacy intrusion. The first measure counts only the 'sensitive' permissions and excludes the 'normal' permissions labelled by Android. It serves as a robustness check to the main metric by measuring the changes in permission requests that impose substantial privacy risks to users. Notably, the 'sensitive' permissions cover only a small subset of Android permissions, as the Android permission classification system tends to be lenient on privacy concerns and often overlooks the potential privacy risks associated with 'normal' permissions, particularly when misused or exploited by malicious parties.

The second measure divides the number of permissions by a proxy of the number of functions

[38] Consumers can only manually grant and revoke dangerous permissions but not other types.

[39] http://www.xinhuanet.com/2021-11/03/c_1128024150.htm

an app provides. Though requesting more permissions always imposes higher privacy risks to users regardless of the number or the type of functions an app provides, some might argue that higher levels of privacy risks are justified when privacy-intrusive permissions are necessary for an app's functionality. The main metric assesses only the level of potential privacy risks associated with permissions; it does not consider the extent to which consumers might be willing to accept such risks when these permission requests can be justified by functionality requirements. In Section 7.3.1, I present an alternative measure that takes functionality into consideration and captures the number of permissions requested *per function*, addressing the concern that increases in functional complexity may justify additional permissions requests. My findings are consistent across different measures for privacy intrusion.

4.3 Classifying Censored vs. Uncensored Markets

I match the Google Play category of each banned foreign app, as identified by the GreatFire Analyzer or reported elsewhere, with the category tag(s) used by the focal Android app store. I classify an app category (as defined by the Android app store) as a censored market if any non-Chinese apps in that category was banned by the Great Firewall according to the list identified in Section 4.1.[40,41] In Section 6.2.1, I address the possibility that treated categories likely differed in the number or influence of banned apps, and thus in their exposure to the treatment. 22 app categories were classified as 'censored' (also referred to as the 'treatment' group) and 66 categories as 'uncensored' (the 'control' group). I summarized the names and frequencies of the censored app categories alongside the 22 largest uncensored categories in Table A1. I present summary statistics for treated and control apps both before and after the regulatory change in Table 1. I check the robustness of the results to different classifications of the treatment and control group in Section 6.2.3.

[40] A few banned apps can be matched to multiple category tags. For example, Facebook has the category tag 'social' on Google play, which has three matching categories on the app store I analyze: 'friend-making', 'communities', and 'chat'. I classified all three categories as in the treated group in this case.

[41] I excluded the category that contains VPN apps, the 'optimization' category, from the analysis. This is because many of the VPN apps were removed from app stores at a faster rate than other apps as a result of the policy (see the discussion in the identification section), making it harder to collect a representative sample of apps in this category during the period of this study (2016-2020). However, including it as a treated category doesn't affect the findings or conclusions of this paper qualitatively.

Table 1: Summary Statistics

Treatment	After	Variable	Mean	SD	N
0	0	N Permissions	26.90	14.92	23148
0	0	N Sensitive Permissions	8.35	3.79	23148
0	0	N Activities	76.58	74.14	23148
0	0	N Ad Activities	3.82	15.86	23148
0	0	N Payment Activities	0.86	1.20	23148
0	1	N Permissions	32.72	19.51	71373
0	1	N Sensitive Permissions	9.72	4.37	71373
0	1	N Activities	121.69	112.40	71373
0	1	N Ad Activities	6.23	26.33	71373
0	1	N Payment Activities	1.41	2.48	71373
1	0	N Permissions	30.12	14.54	10596
1	0	N Sensitive Permissions	9.07	3.88	10596
1	0	N Activities	84.61	74.30	10596
1	0	N Ad Activities	4.99	17.75	10596
1	0	N Payment Activities	0.96	1.39	10596
1	1	N Permissions	37.87	19.70	32671
1	1	N Sensitive Permissions	10.80	4.18	32671
1	1	N Activities	130.65	114.25	32671
1	1	N Ad Activities	7.64	24.17	32671
1	1	N Payment Activities	1.51	1.92	32671

This table presents summary statistics for apps in censored markets (where Treatment = 1) and uncensored markets (where Treatment = 0), both before and after the regulatory shock in January 2017. Each observation represents a monthly snapshot of an individual app.

5 Identification

The key challenge to establish a causal interpretation between changes in competition and changes in firms' privacy-intrusive behavior is that competition in a market is often determined endogenously by factors that could simultaneously affect whether firms engage in privacy-intrusive business practices. To address this challenge, I observed that the levels of Internet censorship varied across markets before the regulatory change. These differences in censorship led to varying degrees of competitive disruption in different app markets, particularly between censored and uncensored markets, after the 2017 regulation on censorship circumvention tools. In censored app markets, the regulation unexpectedly reduced the number of available products, since banned foreign apps that were previously accessible via VPNs (or similar tools) were no longer easily accessible. However, uncensored app

markets, such as banking or weather apps, were less affected by the regulatory change because they were not subject to intensive Internet censorship as of January 2017 and did not have important foreign players in the market that were banned. This variation in treatment exposure, along with the (modified) parallel trend assumption, enables me to identify the causal effect of the change in market competition induced by the 2017 regulation. I do this by comparing apps that were not banned in censored markets with those in uncensored markets, both before and after the regulatory change.

More specifically, I use the synthetic differences-in-differences (SDID) approach proposed by Arkhangelsky et al. (2021). Similar to the conventional differences-in-differences (DID) method, the SDID method includes both a time and a unit fixed effect. However, the two methods differ in the weights they assign to individual control units and pre-treatment time periods. While the conventional DID weights each unit and peroid equally, the SDID method reweights the control units and the pre-treatment periods to construct a 'synthetic' control group that better matches the treated units' average pre-treatment outcome. This method has been applied in several recent studies (Berman and Israeli, 2022; Lambrecht et al., 2023) and has shown superior performance compared to conventional DID and synthetic control methods in applications where these traditional approaches are commonly used.

For a panel dataset consisting of N units over T time intervals, the SDID method estimates the average treatment effect on the treated (ATT) by solving the following equation (Equations (1) to (4) follow from the equations developed in Arkhangelsky et al. (2021)):

$$
\begin{aligned}
&(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) \\
&= \underset{\tau,\mu,\alpha,\beta}{argmin} \Bigg\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (Y_{it} - \mu - \alpha_i - \\
&\beta_t - W_{it}\tau)^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \Bigg\}
\end{aligned} \tag{1}
$$

where Y_{it} is the outcome variable for unit i at time t, W_{it} is the treatment indicator, and τ is the average treatment effect on the treated. The terms α_i and β_t are the units and time fixed effects, similar to the ones used in a conventional DID model. Similar to the synthetic control method, SDID uses unit-specific weights $\hat{\omega}_i^{sdid}$ to align the trends in pre-treatment outcomes of control units with those of treated units. The unit-specific weights $\hat{\omega}^{sdid}$ are chosen in the following way:

$$
\begin{aligned}
&(\hat{\omega}_0, \hat{\omega}^{sdid}) \\
&= \underset{\omega_0 \in \mathbb{R}, \omega \in \Omega}{argmin} \left\{ \sum_{t=1}^{T_{pre}} (\omega_0 + \sum_{i=1}^{N_{co}} \omega_i Y_{it} - \frac{1}{N_{tr}} \sum_{i=N_{co}+1}^{N} Y_{it})^2 \right. \\
&\left. + \zeta^2 T_{pre} \|\omega\|_2^2 \right\},
\end{aligned} \tag{2}
$$

$$
\begin{aligned}
\Omega = \Bigg\{ &\omega \in \mathbb{R}_+^N : \sum_{i=1}^{N_{co}} \omega_i = 1, \ \omega_i = N_{tr}^{-1} \\
&for \ all \ i = N_{co} + 1, ..., N \Bigg\}.
\end{aligned} \tag{3}
$$

where units 1, 2, ... N_{co} are the control units, and units $N_{co}+1$, $N_{co}+2$, ..., N are the treated units, which are subject to the treatment after T_{pre}. The term ζ is a regularization parameter (see Arkhangelsky et al. (2021) for further details). The SDID method also introduces time weights, $\hat{\lambda}_t^{sdid}$, to balance pre- and post-treatment periods. The time weights are selected in such a way that pre-treatment periods more similar to the post-treatment periods are assigned higher weights in the estimation:

$$
\begin{aligned}
&(\hat{\lambda}_0, \hat{\lambda}^{sdid}) \\
&= \underset{\lambda_0 \in \mathbb{R}, \lambda \in \Lambda}{argmin} \left\{ \sum_{i=1}^{N_{co}} (\lambda_0 + \sum_{t=1}^{T_{pre}} \lambda_t Y_{it} \right. \\
&\left. - \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^{T} Y_{it})^2 \right\},
\end{aligned} \tag{4}
$$

where

$$
\begin{aligned}
\Lambda = \Bigg\{ &\lambda \in \mathbb{R}_+^T : \sum_{t=1}^{T_{pre}} \lambda_t = 1, \ \lambda_t = T_{post}^{-1} \\
&for \ all \ t = T_{pre} + 1, ..., T \Bigg\}.
\end{aligned} \tag{5}
$$

In the context of this paper, each unit i in Equation (1) represents an app in the sample, and each time period t represents a month. Y_{it} is a measure for privacy intrusion of the latest version of

app i in month t, such as the number of permissions or sensitive permissions requested by the app. The treatment indicator variable W_{it} takes the value of 1 for observations from permitted apps in censored markets after January 2017, when the new VPN regulation took effect, and 0 otherwise. The coefficient for W_{it}, τ, is the key quantity of interest. It captures the treatment effect of the removal of competition introduced by the regulatory change on apps' permission requests. In my discussion that follows, I refer to permitted apps (those not banned by the Great Firewall) within censored categories as the "treated group", and apps in uncensored categories (all of which were permitted) as the "control group". [42]

6 Empirical Results

Most of my empirical analysis focuses on the subset of apps that were listed on the Android platform *before* January 2016 (henceforth 'incumbent apps'), for which I have at least one year of pre-treatment observations. Using this subset of apps makes intuitive sense as they were in the market prior to the regulatory change, and therefore were directly affected by the removal of competition induced by the new regulation. Using incumbent apps also guarantees a balanced panel throughout the sample period I study, which is required by the SDID method. Throughout most of the paper, I analyze this subset of incumbent apps for the period from January 2016 to January 2020, except for in section 7.3.2 where I discuss entry.

6.1 Main Effect

Figure **??** presents the event study plot from a simple difference-in-differences model. Here, I use a standard two-way fixed effects model to estimate the relative difference in permission requests between censored and uncensored categories each month, compared to their baseline difference in January 2017, conditional on app and year×month fixed effects. Panels (a) and (b) use the log of the total number of permissions and log of the number of sensitive permissions requested by an app as the outcome variable, respectively. The solid vertical line indicates the pre-treatment period $t = -1$. Figure 3 shows that apps in censored categories ('treated' apps) exhibited similar trends from those in uncensored categories prior to the regulatory change, but requested significantly more permissions after the change. While Figure 3 provides reassuring evidence that the parallel trends assumption is plausible in this setting, a closer look at Panel (a) also reveals that the difference in

[42] My empirical analysis did not include banned apps in permitted categories.

total permissions requested between the two groups slightly widened over time prior to the treatment, though this difference was not statistically significant. To provide further reassurance, I next use synthetic difference-in-differences to estimate the treatment effect and to account for any remaining time-varying differences between the two groups.

I present graphical evidence for the synthetic differences-in-differences estimates in Figure 4. To better assess the resemblance of pre-treatment trends between groups after reweighting, I use the 'overlay' algorithm from 'synthdid' R package, which adjusts the trajectory of the synthetic control group by subtracting its (average) pre-treatment difference from the treatment group. The solid vertical line indicates the pre-treatment period $t = -1$, and the dashed line indicates five months after the initial regulatory announcement, which was when the removal of VPNs substantially increased in Figure 1. Figure 4 provides reassuring evidence that the SDID method generates a synthetic control group that matches well with the treatment group on pre-treatment trend. The divergence between treated and synthetic control apps emerged around five to six months after the initial policy announcement, which aligned with the timing of the large-scale regulatory enforcement of VPN removal shown in Figure 1. The difference in permission requests across groups increased gradually over time until early 2019 and remained relatively stable afterwards, although the difference in sensitive permissions continued to expand.

Table 2 presents the numerical estimates. Columns (1) of Table 2 take the log of the number of permission requests as the outcome variable and estimates a conventional DID model, which resembles Equation (1) but does not have the unit and time weights $\hat{\omega}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$. The results show that the restrictions on censorship circumvention tools significantly increased permissions requested by treated apps. I then estimate a synthetic diff-in-diff model in Column (3) using the exact specification of Equation (1). The results shows that an average treated app requested 3.05% ($exp(0.030) - 1 = 0.0305$) additional permissions than a synthetic control app after the regulatory change. This suggests the decrease in competitive pressure in censored categories increased the degree of privacy intrusion by treated apps. I report the jackknife standard errors for the SDID estimator, as suggested by Arkhangelsky et al. (2021). Notably, the SDID estimate is qualitatively similar, but quantitatively slightly smaller than the conventional DID estimate. This suggests that the DID estimator likely overestimates the treatment effect when the raw control data is only a noisy match to the treated units.

The previous analysis uses the number of permissions requested as a proxy for the degree of privacy intrusion by an app. While the total number of permissions indicates the extent of data

Figure 3: Regulatory Change Increased Permission Requests, Simple Difference-in-Difference

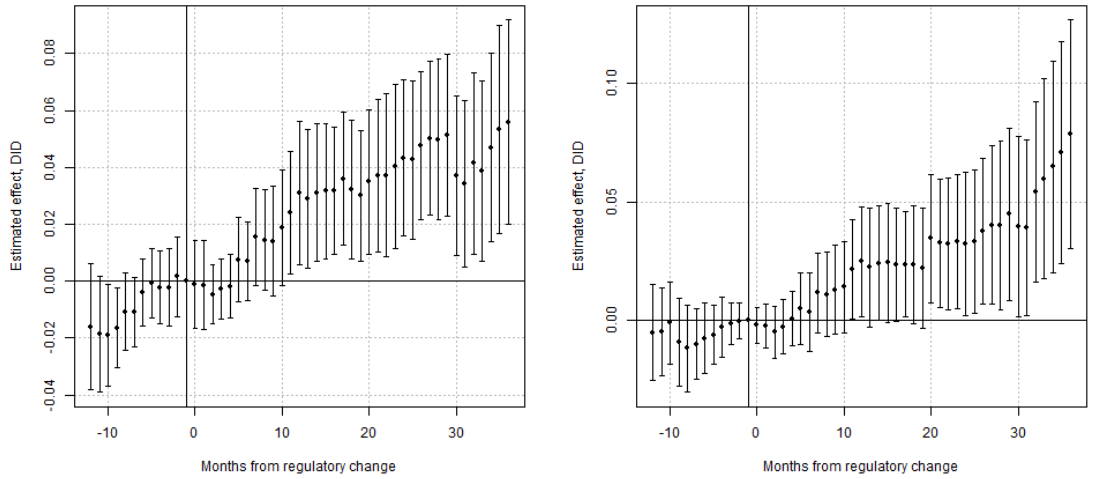

(a) All Permissions

(b) Sensitive Permissions

Notes: This figure presents the event study plot from a simple DID model that estimates the relative difference between censored and uncensored categories at each month, compared to their baseline difference in January 2017, controlling for app-specific and year-month fixed effects. Panel (a) plots the average number of permissions requested by an app. Panel (b) plots the average number of 'sensitive' permissions requested by an app.

Figure 4: Number of Permissions and Sensitive Permissions, SDID

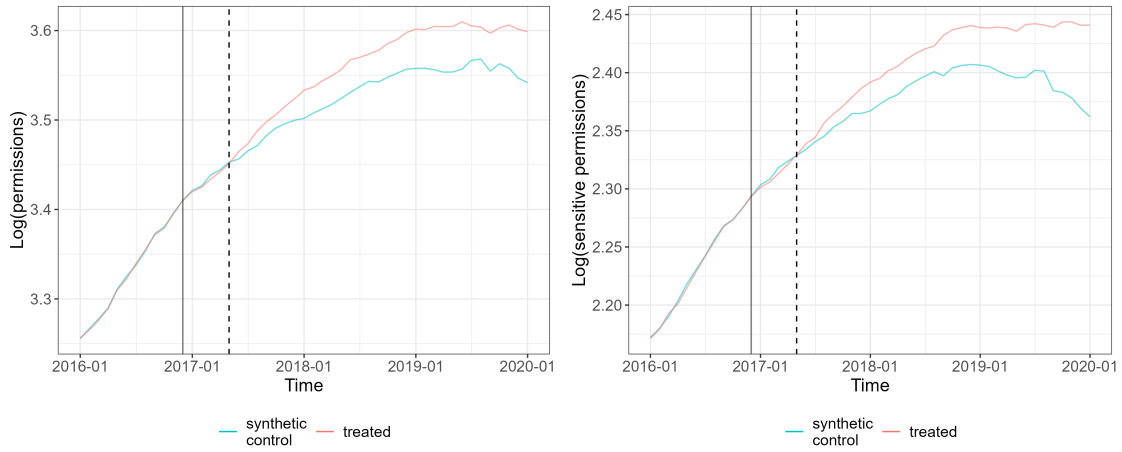

Notes: This figure shows the permission requests from the treated and the synthetic control groups between January 2016 and January 2020. I use the 'overlay' algorithm from 'synthdid' R package to adjust the trajectory of the synthetic control group by subtracting its difference from the pre-treatment average of the treated group. Panel (a) plots the average number of permissions requested by an app. Panel (b) plots the average number of 'sensitive' permissions requested by an app.

access and device control requested by an app, permissions vary in their degree of privacy risks to users. As discussed in Section 4.2, I check the robustness of the findings using the set of permissions that are explicitly labeled as having higher than average privacy risks by Android in Columns (2) and (4) of Table 2. Android's developer guide describes dangerous permissions as "a higher-risk permission that would give a requesting application access to private user data or control over the device that can negatively impact the user".[43] I also include custom permissions whose protection levels are higher than 'normal' in the outcome variable for the analysis in Columns (2) and (4).[44] These permissions are developed by third-party apps and are typically used to share sensitive data or functionality across apps. I refer to the permissions with higher levels of privacy risks as 'sensitive permissions' throughout the rest of the paper. The estimated effects on requests of sensitive permission are qualitatively and numerically similar to those on all permission requests. In particular, the SDID estimate indicates an average of 2.84% ($exp(0.028) - 1 = 0.0284$) additional sensitive permissions requested by treated apps as a result of the removal of competition.

Table 2: The Removal of Competition Led to Increased Permission Requests

	DID		*SDID*	
	(1)	(2)	(3)	(4)
	Log(permissions)	Log(sensitive permissions)	Log(permissions)	Log(sensitive permissions)
Treatment Effect	0.038**	0.033**	0.030***	0.028***
	(0.014)	(0.013)	(0.011)	(0.008)
App FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	137,494	137,494	137,494	137,494

Notes: The unit of observation is an app × month. Cluster robust standard errors at the market and year × month level are reported for DID estimates (Columns (1) - (2)). Jackknife standard errors are reported for SDID estimates (Columns (3) - (4)). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.2 Robustness Checks

I did a battery of additional analysis to check the robustness of the main finding. I discuss the results of these analysis below.

[43] https://developer.android.com/guide/topics/manifest/permission-element

[44] This includes custom permissions whose protection levels are 'dangerous' or 'signature'. Notably, I did not include the Android-defined signature permissions because these permissions are designed for "apps that are signed with the same digital signature as the Android operating system or are included as part of the system firmware" based on the Android developer manual.

6.2.1 Effect is Due to Shift in Competition

I first provide evidence that the regulatory shock affected permission requests through competition rather than through alternative channels. In Table 3, I stratify treated categories based on how much the competition in the category was affected by the regulatory change. Although all treated categories had at least one internationally popular app (that appeared in the Alexa Top 1000 ranking) blocked by the Great Firewall, banned foreign apps had larger market shares among Chinese users in some categories than in others, making these markets more affected by the regulatory change. To capture this difference in treatment exposure, I classify treated categories into three groups. A category is classified as 'less affected' by the regulatory change if none of the banned foreign apps in that category once appeared in the top iOS China category ranking prior to the regulatory change.[45] A category is classified as 'more affected' if at least one banned foreign app appeared in the top iOS China category ranking, and as 'most affected' if at least two such apps appeared, suggesting that banned foreign apps were significant competitors in these markets.

Table 3 reports the results for each stratum. The effect of the regulatory change is larger and statistically significant for treated apps in 'more affected' or 'most affected' categories. The effect is smaller and insignificant for treated apps in 'less affected' categories, where none of the banned foreign apps were among the top choices of Chinese users. Compared to the synthetic control group, treated apps in categories where competition was more affected requested 3.46% more permissions and 2.94% more sensitive permissions; treated apps in categories where competition was most substantially affected requested 5.44% more permissions and 4.39% more sensitive permissions. These findings align with expectations and confirm that the regulatory change on circumvention tools affected apps' behavior through the shift in competition.

As the 2016 Chinese Cybersecurity Law formalized state regulation and control of cyberspace, one potential concern is whether treated categories were disproportionately subject to governmental control and, if so, whether their behavior reflected increased governmental oversight during this period rather than shifts in competition. Though the extent to which governmental internet monitoring varied across categories is unclear and unobservable, I provide indirect evidence supporting the robustness of my findings against this alternative explanation. Specifically, I analyze the cat-

[45] My data collection covers the top 300 apps in each iOS category on the first day of each month in 2016. I map the iOS categories to the categories of the focal Android app store analyzed in this paper. I use the Chinese iOS category ranking in this analysis because there was no widely used app ranking that covered all major Chinese domestic Android app stores during this period to the best of my knowledge.

Table 3: Treatment Effect Is Larger for Categories where Competition Was More Affected

	Competition Less Affected		*Competition More Affected*		*Competition Most Affected*	
	(1)	(2)	(3)	(4)	(5)	(6)
	Log(permissions)	Log(sensitive perm.)	Log(permissions)	Log(sensitive perm.)	Log(permissions)	Log(sensitive perm.)
Treatment Effect	0.018	0.023	0.034***	0.029***	0.053***	0.043***
	(0.018)	(0.014)	(0.012)	(0.009)	(0.014)	(0.010)
App FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	101,381	101,381	127,449	127,449	116,277	116,277

Notes: Synthetic differences-in-differences estimates reported. The unit of observation is an app × month. Cluster robust standard errors at the market and year - month level are reported for DID estimates (Columns (1) - (2)). Jackknife standard errors are reported for SDID estimates (Columns (3) - (4)). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

egory of news apps, which was among the most affected treated categories, as access to multiple influential international newspapers – such as the New York Times and the Wall Street Journal – was impacted by the regulatory change. Although all news apps in the Chinese market faced the same governmental regulation, they varied in their primary content focus. Some reported on national and global news and competed directly with (banned) international newspapers, while others focused on regional news at the city or province level and thus were not close competitors to the banned foreign products. This variation in product positioning within the same industry provides a natural opportunity to test the influence of the competition shock while holding constant the degree of governmental regulation.

In Table 4, I group news apps into national and global news apps, and regional news apps. I then estimate the effect of the regulatory change to each group separately. If the change in permission requests reflects a decrease in competitive pressure from major international newspapers, the effect should be larger for global and national news apps and smaller for regional ones. By contrast, if it was due to differences in governmental influence between treated and control categories, different news apps should exhibit similar effects as they faced the same regulatory environment. The results in Table 4 show that the treatment effect is much larger for news apps that reported on national and global news, compared to those focused on regional news. This difference is pronounced in requests for sensitive permissions, where national and global news apps increased their requests of sensitive permissions by 7.57%, whereas regional news apps only had an (insignificant) 1.92% increase in such requests. This is consistent with the hypothesis that changes in permission requests reflected shifts in competitive pressure rather than differential governmental influence.

Table 4: Treatment Effect Is Larger for National News Apps Compared to Regional News Apps

	National & Global News		*Regional News*	
	(1)	(2)	(3)	(4)
	Log(permissions)	Log(sensitive perm.)	Log(permissions)	Log(sensitive perm.)
Treatment Effect	0.061	0.073**	−0.006	0.019
	(0.054)	(0.028)	(0.032)	(0.024)
App FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	92,855	92,855	94,570	94,570

Notes: Synthetic differences-in-differences estimates reported. The unit of observation is an app × month. Jackknife standard errors are reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.2.2 Alternative Dependent Variables

I next check the robustness of the results to alternative specifications of the dependent variable. In Columns (1) and (2) of Table 5, I replicate the analysis in Columns (3) and (4) of Table 2 using the number of permission requests and sensitive permission requests (without the log transformation) as the outcome variables in the SDID regressions. The results show that the regulatory change led to an average increase of 1.34 permissions and 0.27 sensitive permissions, which is consistent with the conclusions from the analysis using log-transformed outcome variables.

One question is whether treated apps requested more permissions because they added additional functions to their products after the removal of competition. If so, the quality improvement from these added features might outweigh the privacy risks of additional permission requests, leaving the overall welfare implications unclear. To empirically test this hypothesis, I leverage a unique feature of my data, which is the record all Android activities each app performed. An 'activity' in Android represents a single screen in an app's user interface, such as 'select photo' or 'send email'. The number of activities of an app represents the number of screens that users can directly interact with and can serve as a proxy for the number of functions provided by the app. In Columns (3) and (4) of Table 5, I redefine the outcome variable in Equation (1) as the number of permissions and sensitive permissions requested by an app, divided by the number of activities performed by the app. As discussed in Section 4.2, this dependent variable captures the number of permissions requested per app function, taking into account that apps with more functions might naturally request more permissions. The result suggests that even on a per-function basis, treated apps requested more

permissions and more sensitive permissions compared to synthetic control apps after the regulatory change.

6.2.3 Alternative Treatment Classification

In Columns (5)–(8) of Table 5, I test the robustness of the results to alternative classifications of treatment and control apps. While most categories fall clearly into one group or the other based on the censorship status of major foreign players, the classification of map and travel apps is less straightforward. In particular, I classify map and travel apps as part of the treated category in the main analysis since Google Maps, the leading map service in the global market, was blocked in China. However, some user-generated comments[46] suggested that offline versions of Google Maps could still be used if maps were downloaded outside the country, though this option was unlikely to be accessible to the majority of users within China. Google Maps also appeared to be outdated in its geographical information in China since Google's exit, which cast doubt on whether it was a meaningful competitor to other map and travel apps even if users accessed it via VPNs. To assess how these classification challenges affected the results, I re-estimate the model after removing the map and travel categories from the analysis in Columns (5) and (6) of Table 5. The findings remain robust and numerically similar, indicating that the inclusion or exclusion of these categories in the treatment group does not affect the main findings.

Another challenge involves the classification of lottery apps. GreatFire.org reported no records of blockages of major non-Chinese lottery services, though it did test blockages of sports-betting sites. In China, betting and lottery are treated separately. The former is considered gambling and is prohibited by law in mainland China,[47] whereas the latter is legally permitted with its revenue often serving public welfare functions for the state.[48] Because I was not aware of any major non-Chinese lottery providers being blocked, and given the clear legal distinction between lottery and betting services, I classify lottery apps as part of the control group in the main analysis. However, since both lottery and betting products involve uncertain rewards, it is possible that some consumers may view them as substitutes. To assess how the classification of lottery apps affects the findings, I rerun the main analysis in Columns (7) and (8) of Table 5 but exclude the category of lottery apps. The

[46] https://www.reddit.com/r/chinalife/comments/1bsf4go/first_timer_is_it_ok_to_use_google_maps_with_a/

[47] The special administrative regions of Macau and Hong Kong have distinct legal systems. In particular, gambling is legal in Macau and is a major source of revenue for the region.

[48] https://legalpilot.com/country/china/, accessed September 6, 2025

results remain robust.

6.2.4 Alternative Window of Analysis

The window of the main analysis covers the enactment of the European Union (EU)'s General Data Protection Regulation (GDPR) in 2018. The GDPR is a flagship privacy regulation that protects the personal data of individuals in the EU. It affects organizations worldwide if they collect or target such data. Although the GDPR does not directly regulate the Chinese market, it may still influence Chinese apps indirectly by shaping the behavior of their foreign competitors that also serve the EU market, or by raising general awareness of privacy-related issues and thereby changing business practices. Because the treated apps faced less competitive pressure due to the regulatory change, they might have been affected to a different degree compared to the control apps. To check the robustness of the results to differential impacts of the GDPR, I re-estimate the model using only pre-GDPR data in Columns (9) and (10) of Table 5. The results indicate that the regulatory change had a significant impact on the treatment group relative to the synthetic control group, even prior to the launch of the GDPR.

6.2.5 VPN-related Permissions

Because VPNs mask users' real IP addresses with those of VPN servers, and IP addresses can be used for geo-location inference, some users also use VPNs to obscure their locations. As the ban on VPNs made hiding IP addresses much harder, it might affect consumers' tolerance of location data collection by apps in general, which might lead to a rise in permission requests related to such data. Note that under the DID/SDID framework, this mechanism would affect the findings only if the inability to hide IP-based location *differentially* affected how consumers interact with apps with and without banned foreign competitors. To test this possibility, I re-estimate the model excluding all permissions directly related to location data. Though an app does not need any additional permission to access a device's IP address once it connects to the Internet, it needs permissions to access more accurate GPS-based location data. If the change in consumers' general attitude toward location data collection is the main driver of the finding, we should expect the treatment effect to go away after excluding all location-related permissions from the analysis. However, the results remain, suggesting that the alternative mechanism is unlikely the key driver of the findings.

Table 5: Additional Robustness Checks

	Alterntive DV 1		*Alterntive DV 2*		*Alternative Treatment*		*Alternative Control*		*Alternative Window*		*Excluding Location*	
	(1) N perm.	(2) N sen. perm.	(3) Perm. per activity	(4) Sen. perm. per activity	(5) Log- (perm.)	(6) Log- (sen. perm.)	(7) Log- (perm.)	(8) Log- (sen. perm.)	(9) Log- (perm.)	(10) Log- (sen. perm.)	(11) Log- (perm.)	(12) Log- (sen. perm.)
Treatment Effect	1.340***	0.270***	0.093***	0.040***	0.035***	0.033***	0.029***	0.028***	0.015*	0.011*	0.033***	0.038***
	(0.471)	(0.076)	(0.028)	(0.009)	(0.011)	(0.008)	(0.011)	(0.008)	(0.008)	(0.006)	(0.011)	(0.008)
App FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	137,494	137,494	137,494	137,494	134,407	134,407	137,151	137,151	81,374	81,374	137,494	137,494

Notes: Synthetic differences-in-differences estimates reported. The unit of observation is an app × month. Cluster robust standard errors at the market and year - month level are reported for DID estimates (Columns (1) - (2)). Jackknife standard errors are reported for SDID estimates (Columns (3) - (4)). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Heterogeneous Treatment Effects

I investigate the heterogeneity in the effects of reduced competition across app genres. Genres are defined by the app store and consist of several functionally related app categories.[49] I compare treated apps in each genre to a synthetic control group constructed using the entire set of control apps.[50] Figure 5 presents the estimated treatment effects on the log of total and sensitive permission requests across genres. Figure A3 presents the estimated effects on the number of each type of requests. The point estimate of the treatment effect is positive for most genres. The percentage increase in total requests is largest for treated apps in the audio/video genre, followed by those in the photography and image-sharing genre. Treated apps in the system tools genre (which includes browser and file-management apps) and the communication and social genre are estimated to have insignificant percentage increases. However, Figure A3 shows that communication and social apps actually experienced large and significant increases in the number of requests, suggesting that the small percentage change is likely due to the large number of baseline permission requests in this genre. Travel apps are the only genre that experienced a significant percentage decrease in permission requests, though Figure A3 suggests this corresponds to relatively small changes in absolute levels.

[49] For example, 'Communication & social' contains four treated app categories: chat, friending, telecommunication, and communities.

[50] If a genre contains both treated and control categories, I use the treated app categories in this genre and all control apps categories in the sample, including control app in the same genre and in different genres, in the synthetic diff-in-diff analysis for this genre.

Figure 5: Treatment Effect by App Genre

(a) Estimated effect on log(number of permissions)

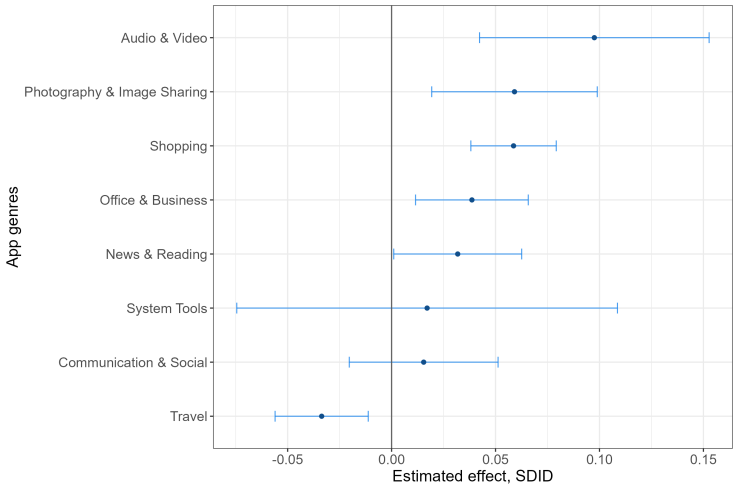

(b) Estimated effect on log(sensitive permissions)

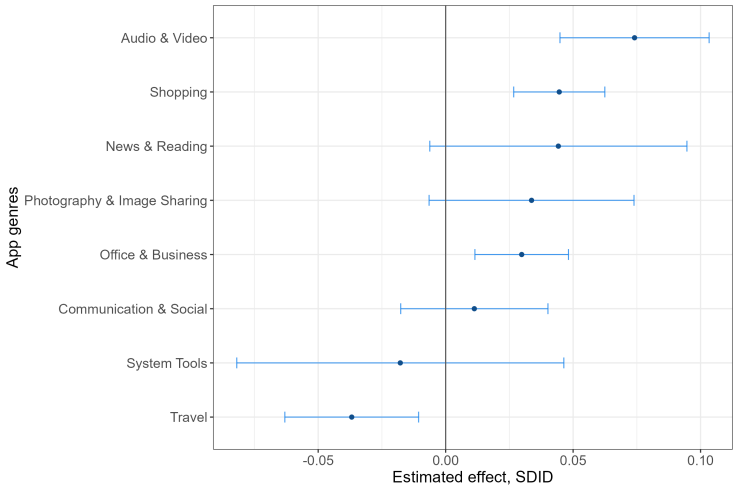

Notes: Panel (a) presents SDID estimates of the effect on the log of permission requests for each app genre. Panel (b) presents the SDID estimates of effect on the log of sensitive permissions requests for each app genre. Genres are defined by the app store I study and contain several app categories (markets) that are functionally related. If a genre has both treated and control categories, I use all control apps, including those in the same genre and those in different genres, to construct the synthetic control group for treated apps in this genre.

7 Mechanism

Why does the shift in competitive pressure change firms' privacy practice? One explanation is that lower competitive pressure means consumers have fewer choices and are therefore less able to discipline firms' intrusive behavior through switching. When limited consumer choice constrains switching, firms have stronger incentives to aggressively collect data and request device access in order to better engage and monetize consumers. In this section, I first document that increases in privacy-intrusive practice are systematically associated with lower firm incentives to compete on privacy due to limited (privacy-friendly) outside options for consumers. I then provide evidence that intrusive practices are driven by incentives and practices related to capturing and monetizing user engagement.

7.1 Regulatory Change Reshaped Firms' Incentives to Compete on Privacy

I first investigate whether the removal of competitive pressure effectively limits consumers' privacy choices and thereby reduces firms' incentives to compete on privacy. To do so, I leverage the fact that the degree of domestic differentiation in privacy protection varies across markets. In markets where domestic firms differentiate on privacy protection, consumers retain privacy options even when foreign competitors are removed. Constrained by domestic competition, firms in these markets are less likely to substantially increase privacy-intrusive data collection practices. In contrast, in markets where domestic firms do not differentiate on privacy, the removal of foreign competitors - who are generally more privacy-friendly - substantially reduces consumers' privacy options and further weakens firms' incentives to protect privacy. Therefore, we should expect larger treatment effects in markets with lower baseline domestic differentiation in privacy, where the removal of competition leads to a significant reduction in competitive pressure on privacy.

Table 6 tests this prediction. In Columns (1) and (2) of Table 6, I stratify markets based on how much domestic firms differentiate on privacy. I measure domestic differentiation on privacy as the standard deviation of the number of dangerous permissions requested across Chinese apps within a market prior to the regulatory change. I define high differentiation markets as those with a standard deviation above the median, and low differentiation markets as those below the median. comparing Columns (1) and (2) confirms that the effect of the removal of competition is indeed larger in markets with low prior differentiation on privacy among domestic firms.

Permission requests differ in their salience to consumers. If firms' post-treatment increases in permission requests are driven by reduced consumer switching following the removal of competition, we should observe smaller increases in more salient permissions, which are more likely to trigger switching. This should be particularly so in markets with greater privacy differentiation, where switching remains easier compared to in low-differentiation markets. I investigate how effects of the regulatory change vary with permission salience in Columns (3) - (6) of Table 6.

In Android (post Android 6.0), requests for dangerous permissions are more salient to consumers both because they govern access to more sensitive data and device features, and because the Android system prompts users to explicitly grant or deny such permissions when an app first attempts to access a protected resource.[51] For this reason, dangerous permissions are also commonly referred to as runtime permissions. In contrast, install-time permissions are less salient because they are granted when the user installs the app and do not trigger additional prompts during use. These permissions are disclosed at installation, typically as part of the app's permission list, and access to the associated resources is granted immediately if the user proceeds, without further user interaction at runtime.

Columns (3)–(6) of Table 6 examine the effects of the regulatory change on low- versus high-salience permissions across markets with different degrees of privacy differentiation. Comparing Columns (3)–(4) with Columns (5)–(6) shows that, across market types, treated apps have smaller percentage increases in high-salience permissions than in low-salience permissions. This pattern is consistent with the hypothesis that firms face stronger constraints when intrusive permission requests are more likely to attract consumer attention and induce switching. Notably, while markets with high and low privacy differentiation display similar increases in low-salience permissions, increases in high-salience permissions are concentrated in markets with low domestic privacy differentiation. This suggests that firms act more intrusively when consumers have fewer privacy-friendly outside options.

I next exploit within-market variation in baseline privacy protection by comparing treated Chinese apps that were more versus less privacy-intrusive prior to treatment. In Table 7, I classify apps as more or less privacy intrusive based on a median split of their percentile rank in the distribution of high-salience (dangerous) permissions within their market as of December 2016. Because foreign apps requested fewer permissions and offered stronger privacy protection on average (), their

[51] If granted, the app can subsequently access the protected resource without additional prompts; if denied, access is blocked, although the app may request the permission again at a later time.

Table 6: Stratification by the Degree of Market Differentiation

	All Permissions		*Low Salience Permissions*		*High Salience Permissions*	
	(1)	(2)	(3)	(4)	(5)	(6)
	High differentiation	Low differentiation	High differentiation	Low differentiation	High differentiation	Low differentiation
Treatment Effect	0.026*	0.036**	0.031*	0.032*	0.007	0.022**
	(0.014)	(0.016)	(0.017)	(0.019)	(0.010)	(0.011)
App FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	79,184	58,310	79,184	58,310	79,184	58,310

Notes: This table reports SDID estimates of the regulatory change on app permission requests, stratified by market-level privacy differentiation. Differentiation is measured as the pre-treatment (through December 2016) standard deviation of high salience (dangerous) permissions across apps within a market. Markets with above-median dispersion are classified as high differentiation and those below the median as low differentiation. Columns (1)–(2) report effects on log total permissions, columns (3)–(4) on log low-salience (install-time) permissions, and columns (5)–(6) on log high-salience (dangerous) permissions. The unit of observation is an app × month. Jackknife standard errors are reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

removal disproportionately affected the choice sets of consumers who preferred privacy-friendly Chinese apps, for whom foreign apps were close substitutes. As a result, the removal of foreign apps directly reduced switching opportunities for these consumers. By contrast, more privacy-intrusive Chinese apps continued to face competition on privacy from relatively more privacy-friendly domestic firms, and thus experienced a smaller change in competitive pressure along this dimension. Columns (1)–(2) of

This analysis predicts larger treatment effects among less intrusive apps, which were closer to foreign competitors in baseline privacy and therefore faced stronger incentives to adjust following the removal of competition. Columns (1) - (2) of Table 7 confirm this prediction. Columns (3)–(6) further show that, while treated apps increased low-salience permissions requests regardless of baseline intrusiveness, increases in high-salience permission requests are concentrated among apps that were less intrusive prior to treatment and thus closer substitutes to the banned foreign apps. This pattern suggests that when privacy is more salient, firms that were previously more privacy-friendly respond more strongly, as the removal of privacy-friendly competitors weakens their incentives to compete for marginal users who value privacy and had viable outside options before the regulatory change. This mechanism is consistent with standard models of horizontal differentiation, in which policy-induced changes have larger effects when competitive constraints bind more tightly.

Taken together, the evidence presented in this section supports a competing-with-privacy mechanism, in which the regulatory change operates by reshaping firms' incentives to differentiate on privacy. Firms adjust permission requests in response to reduced consumer choice and switching

Table 7: Stratification by the Pre-treatment App Intrusiveness

	All Permissions		*Low Salience Permissions*		*High Salience Permissions*	
	(1)	(2)	(3)	(4)	(5)	(6)
	More intrusive apps	Less intrusive apps	More intrusive apps	Less intrusive apps	More intrusive apps	Less intrusive apps
Treatment Effect	0.021	0.048***	0.031*	0.041**	−0.005	0.040***
	(0.014)	(0.016)	(0.017)	(0.018)	(0.008)	(0.012)
App FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	66,836	70,609	66,836	70,609	66,836	70,609

Notes: This table reports SDID estimates of the regulatory change on app permission requests, stratified by apps' pre-treatment privacy intrusiveness within their market. Apps are classified based on their percentile rank in the distribution of high-salience (dangerous) permissions within the market as of December 2016. Apps at or above the median are classified as more intrusive, while those below the median are classified as less intrusive. Columns (1)–(2) report effects on log total permissions, columns (3)–(4) on log low-salience (install-time) permissions, and columns (5)–(6) on log high-salience (dangerous) permissions. The unit of observation is an app × month. Jackknife standard errors are reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

due to changes in competitive pressure, with stronger effects in settings with limited baseline differentiation and among firms closest to removed privacy-friendly competitors.

7.2 Regulatory Change Increased Intrusive Practices to Capture and Monetize Engagement

In the previous section, I show that reduced competitive pressure weakens firms' incentives to protect consumer privacy, leading to more privacy-intrusive practices. This section investigates how firms benefit from such behavior. In particular, I study whether increased access to user data and device control reflects firms' efforts to enhance and monetize user engagement following the change in competitive pressure.

Testing this hypothesis directly requires data on developers' intention and effort to engage consumers, which not available in my data. However, it is possible to test this hypothesis indirectly by comparing cases where app developers are likely to have strong incentives to engage consumers with cases where such incentives are absent. The intrusion-for-engagement hypothesis would predict a larger effect on permission requests in the former case compared to the later case. In particular, when an app developer can monetize and profit from user engagement directly, they are likely to have a stronger incentive to improve engagement, compared to situations where the developer does not profit directly from app usage, such as when the app serves as a complement to an offline business.

To empirically test this idea, I collect additional data on app monetization models by examining the activities documented in their Android Manifest files prior to the regulatory change. I analyze

the names of app activities and identify those associated with two primary methods of monetizing engagement: displaying ads and processing payments. Based on the presence of such activities, I classify apps into four categories: ad-based models, in-app purchase models, a combination of both, or 'non-profit' apps. In Table 8, I compare the treatment effect between 'low incentive' apps, which did not directly monetize engagement via either advertising or in-app purchases, and 'high incentive' apps, which employed one of the two monetization methods throughout the treatment period. Table 8 supports the previous hypothesis: reduced competition led to more permission requests when app developers have incentives to engage consumers due to their monetization model. The results also show that treated apps using both monetization methods increased their permission requests, though the percentage increase is larger for apps that monetize with the advertising model.

Table 8: Stratification by Monetization Models

	Monetize with Ads		*Monetize with In-app Purchases*		*No Monetization*	
	(1)	(2)	(3)	(4)	(5)	(6)
	Log(permissions)	Log(sensitive permissions)	Log(permissions)	Log(sensitive permissions)	Log(permissions)	Log(sensitive permissions)
Treatment Effect	0.055***	0.068***	0.038	0.044**	0.038	0.005
	(0.018)	(0.014)	(0.031)	(0.022)	(0.034)	(0.026)
App FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45,913	45,913	19,012	19,012	22,736	22,736

Notes: Synthetic differences-in-differences estimates reported. The unit of observation is an app × month. Jackknife standard errors are reported for SDID estimates. Columns (1)–(2) analyze the subsample of apps that monetized through ads (but not in-app purchases) prior to the treatment. Columns (3)–(4) analyze the subsample of apps that monetized through in-app purchases (but not ads) prior to the treatment. Columns (5)–(6) analyze the subsample of apps that did not use either ads or in-app purchases prior to the treatment. Apps that used both ads and in-app purchases are dropped from this analysis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7.3 Ruling out Alternative Explanations

I now discuss alternative explanations for the effect that are ultimately not supported by the empirical evidence.

7.3.1 App Functionality

One possible explanation is that the increase in permission requests reflected changes in app functionality. For example, treated apps might have more profit as the competitive pressure decreased, which could then be invested in developing additional functions. This may naturally require additional permissions, and might not raise significant privacy concerns for consumers if those permissions are properly utilized to support the new features.

In Panel (a) of Table 9, I test this hypothesis empirically by combining the previous analysis with additional data on app functionality. A unique feature of my data is that it records all activities each app performs. An 'activity' is an Android terminology for a window in which the app draws its user-interface, such as 'select photo' or 'send email'. The number of activities of an app represents the number of windows that users can directly interact with, and is therefore a good measure for the number of functions provided by the app. In Column (1) of Table 9, I estimate an SDID model on whether the regulatory change significantly changed the number of activities provided by treated apps compared to control apps. The results show no significant differences in the number of activities between the treated and control apps. I conclude that the change in app functionality is not likely the main driver of the findings.

7.3.2 Strategic Differentiation from Entrants

Another potential explanation for the findings is that treated apps developed more privacy-intrusive behavior as a strategic response to new entrants. My analysis so far has focused on incumbent apps that were available to consumers before the regulatory change in 2017. One might wonder whether the removal of competition in treated markets also affected the rate of new entrants, which might in turn affect the behavior of incumbents. In the context of this paper, an entry into a market refers to the availability of a new app within the associated category on the app store. I observed the entry of a new app only if it survived until the point of my data collection in August 2022. If an app entered a market but later decided to exit and removed itself from the app store prior to my data collection, I would not be able to observe it. I therefore caution readers that the entry results presented in this section should be interpreted as conditional on survival.

The SDID result presented in the first column of Table 9 Panel (b) shows that the number of entrants to each market in the treated group was not significantly different from that of the synthetic control group. These results suggest that the removal of competition impacted treated markets mainly through changes in the behavior of incumbent apps, rather than by significantly affecting the rate of new entries into these markets. This is understandable, as the primary competitors of the banned foreign apps (such as Facebook and YouTube) in China are likely well-established apps developed by large firms that were already present in the market before the regulatory change. As a result, the demand affected by the removal of banned apps is more likely to be captured by these incumbent firms, rather than attracting new entrants. Columns (2) and (3) of Table 9 Panel (b) further show that the entrants in treated markets were not significantly different from those

in control markets in terms of the number of permissions or the number of sensitive permissions they requested at the time of entry. If anything, the results show that entrants to treated markets followed the behavior of the incumbents: they requested slightly more permissions and slightly more privacy-sensitive permissions compared to entrants to the control markets, though this difference is not statistically distinguishable from zero. These results suggest that the response of incumbent apps to a group of strategically different entrants is unlikely the main explanation of the findings.

Table 9: Ruling out Alternative Explanations

(a) App Functionality

	App Functionality		
	(1) Number of Activities	(2) Permissions/ Activity	(3) Sensitive Perms/ Activity
Treated	0.525	0.074***	0.037***
Std. Error	(2.190)	(0.028)	(0.008)

Notes: Synthetic diff-in-diff estimates are reported. Jackknife standard errors are in parentheses. The unit of observation for columns (1) - (3) is an app version × month, and the number of observations is 137,788. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(b) Entry & Differentiation

	Entry & Differentiation		
	(1) Number of Entrants	(2) All Permissions	(3) Sensitive Permissions
Treated	-0.166	0.489	0.138
Std. Error	(0.389)	(1.950)	(0.445)

Notes: Synthetic diff-in-diff estimates are reported. Jackknife standard errors are in parentheses. The unit of observation for columns (1) - (3) is a market (category) × month, and the number of observations is 4,312. The outcome variables in the three columns are: the number of new apps that entered a market at a given month, the average number of permissions requested by new apps that entered a market at a given month (measured at the first month of entrance), and the average number of sensitive permission requested by new apps that entered a market at a given month (measured at the first month of entrance). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

8 Conclusion

This paper provides one of the first empirical studies of the relationship between competition and firms' privacy intrusive behavior. I study this question in the context of Chinese Android app markets. In particular, I examine a regulation that prohibited censorship-circumvention tools commonly used by consumers to access apps banned by the government in censored app markets. This regulation reduced the competition faced by permitted apps in censored markets, but did not affect apps in uncensored markets. A synthetic differences-in-differences analysis comparing permitted apps in censored markets to those in uncensored markets reveals that the removal of competition resulted in a significant increase in permission requests by the former compared to the latter. This effect is driven by an increase in requests for a broad range of privacy-intrusive permissions.

I present evidence that the removal of competition affected firms' permission requests by changing their incentives to differentiate and compete for consumers on the privacy dimension. I further analyze app installation packages and find that treated apps seek more data and higher levels of device control via permissions to better monetize user engagement. The observed increase in permission requests cannot be explained by treated apps increasing their app functionality or engaging in strategic differentiation from new entrants.

The findings of this paper have important implications. The findings highlight that competition is an important factor that affects firms' privacy-intrusive behavior. The results show the removal of competition led to substantial increases in privacy-intrusive behavior across different markets and categories. The impact of insufficient competition on consumer privacy is likely to be particularly profound for digital products in data-rich settings such as communication, social, and media categories, and when a product profits from monetizing consumer engagement. The findings also reveal an additional mechanism beyond the commonly discussed privacy intrusion from ad technologies: the removal of competition may also impact consumer privacy by altering firms' incentives to engage and monetize user attention through intrusive methods. Finally, these findings provide an important perspective to understand the potential harm of insufficient competition. They suggest that regulators should carefully consider the potential consequences on consumer privacy when examining business practices that might reduce competition, such as data-driven mergers, highlighting the importance to coordinate regulatory efforts in protecting privacy and promoting market competition.

References

Acquisti, A., L. K. John, and G. Loewenstein (2013). What is privacy worth? *The Journal of Legal Studies 42*(2), 249–274.

Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science 24*(3), 367–381.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review 111*(12), 4088–4118.

Athey, S., C. Catalini, and C. Tucker (2017). The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.

Banker, R. D., I. Khosla, and K. K. Sinha (1998). Quality and competition. *Management science 44*(9), 1179–1192.

Berman, R. and A. Israeli (2022). The value of descriptive analytics: Evidence from online retailers. *Marketing Science 41*(6), 1074–1096.

Berry, S. T. and J. Waldfogel (2001). Do mergers increase product variety? evidence from radio broadcasting. *The Quarterly Journal of Economics 116*(3), 1009–1025.

Campbell, J., A. Goldfarb, and C. Tucker (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy 24*(1), 47–73.

Casadesus-Masanell, R. and A. Hervas-Drane (2015). Competing with privacy. *Management Science 61*(1), 229–246.

Chen, Y. and D. Y. Yang (2019). The impact of media censorship: 1984 or brave new world? *American Economic Review 109*(6), 2294–2332.

Chen, Z., C. Choe, J. Cong, and N. Matsushima (2022). Data-driven mergers and personalization. *The RAND Journal of Economics 53*(1), 3–31.

Clayton, R., S. J. Murdoch, and R. N. Watson (2006). Ignoring the great firewall of china. In *International workshop on privacy enhancing technologies*, pp. 20–35. Springer.

Cunningham, C., F. Ederer, and S. Ma (2021). Killer acquisitions. *Journal of Political Economy 129*(3), 649–702.

Day, G. and A. Stemler (2019). Infracompetitive privacy. *Iowa L. Rev. 105*, 61.

Economides, N. and I. Lianos (2019). Restrictions on privacy and exploitation in the digital economy: a competition law perspective. *Forthcoming, Journal of Competition Law and Economics, CLES Research Paper Series 5*(2019), 21–02.

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review 103*(5), 1598–1628.

Goldfarb, A. and V. F. Que (2023). The economics of digital privacy. *Annual Review of Economics 15*.

Goolsbee, A. and A. Petrin (2004). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica 72*(2), 351–381.

Hellmann, T. F., K. C. Murdock, and J. E. Stiglitz (2000). Liberalization, moral hazard in banking, and prudential regulation: Are capital requirements enough? *American economic review 91*(1), 147–165.

Igami, M. and K. Uetake (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies 87*(6), 2672–2702.

Janssen, R., R. Kesler, M. E. Kummer, and J. Waldfogel (2022). Gdpr and the lost generation of innovative apps. Technical report, National Bureau of Economic Research.

Jiang, L., R. Levine, and C. Lin (2018). Does competition affect bank risk? *Journal of Money, Credit and Banking*.

Johnson, G. A., S. K. Shriver, and S. G. Goldberg (2023). Privacy and market concentration: intended and unintended consequences of the gdpr. *Management Science*.

Kehr, F., T. Kowatsch, D. Wentzel, and E. Fleisch (2015). Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal 25*(6), 607–635.

Kerner, S. M. (2022). Great firewall of china. https://www.techtarget.com/whatis/definition/Great-Firewall-of-China. Accessed: 2023-01-20.

Kesler, R., M. E. Kummer, and P. Schulte (2017). Mobile applications and access to private data: The supply side of the android ecosystem. *ZEW-Centre for European Economic Research Discussion Paper* (17-075).

Krafft, M., C. M. Arden, and P. C. Verhoef (2017). Permission marketing and privacy concerns—why do customers (not) grant permissions? *Journal of interactive marketing 39*(1), 39–54.

Kummer, M. and P. Schulte (2019). When private information settles the bill: Money and privacy in google's market for smartphone applications. *Management Science 65*(8), 3470–3494.

Lambrecht, A., C. Tucker, and X. Zhang (2023). Tv advertising and online sales: a case study of intertemporal subsitution effects for an online travel platform. *Journal of Marketing Research*.

Lin, T. (2022). Valuing intrinsic and instrumental preferences for privacy. *Marketing Science 41*(4), 663–681.

Marotta-Wurgler, F. (2016). Self-regulation and competition in privacy policies. *The Journal of Legal Studies 45*(S2), S13–S39.

Matsa, D. A. (2011). Competition and product quality in the supermarket industry. *The Quarterly Journal of Economics 126*(3), 1539–1591.

Mou, Y., K. Wu, and D. Atkin (2016). Understanding the use of circumvention tools to bypass online censorship. *New Media & Society 18*(5), 837–856.

Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic theory 18*(2), 301–317.

Mylonas, A., M. Theoharidou, and D. Gritzalis (2014). Assessing privacy risks in android: A user-centric approach. In *Risk Assessment and Risk-Driven Testing: First International Workshop, RISK 2013, Held in Conjunction with ICTSS 2013, Istanbul, Turkey, November 12, 2013. Revised Selected Papers 1*, pp. 21–37. Springer.

Ohlhausen, M. K. and A. P. Okuliar (2015). Competition, consumer protection, and the right [approach] to privacy. *Antitrust LJ 80*, 121.

Peukert, C., S. Bechtold, M. Batikas, and T. Kretschmer (2022). Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science 41*(4), 746–768.

Pindyck, R. S. (2007). Mandatory unbundling and irreversible investment in telecom networks. *Review of Network Economics 6*(3).

Posner, R. A. (1981). The economics of privacy. *The American economic review 71*(2), 405–409.

Spiekermann, S., J. Grossklags, and B. Berendt (2001). E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pp. 38–47.

Sweeting, A. (2010). The effects of mergers on product positioning: evidence from the music radio industry. *The RAND Journal of Economics 41*(2), 372–397.

Tsai, J. Y., S. Egelman, L. Cranor, and A. Acquisti (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information systems research 22*(2), 254–268.

Warren, S. and L. Brandeis (1989). The right to privacy. In *Killing the Messenger*, pp. 1–21. Columbia University Press.

Wasastjerna, M. C. (2018). The role of big data and digital privacy in merger review. *European Competition Journal 14*(2-3), 417–444.

A1 Appendix

A1.1 Additional Evidence for Changes in Competitive Pressure

Figure A1: Removal of All VPNs

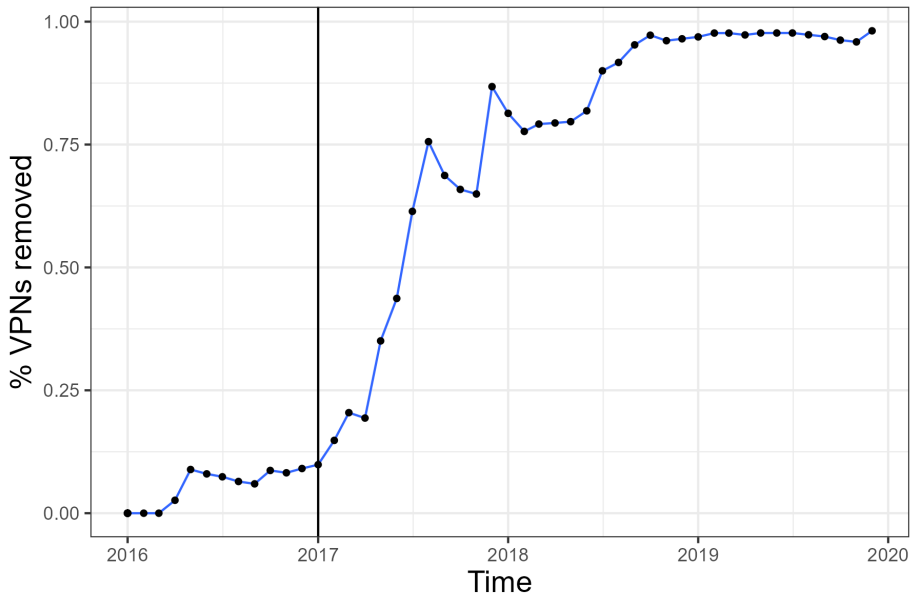

Notes: This figure analyzes 269 VPN apps that ranked in the top 1,000 of China's iOS App Store at least once between 2016 and 2019, including both incumbents and post-2016 entrants. It thus captures the most popular VPNs among Chinese users over this period. Each month, the figure plots the share of all VPNs which had ever entered the market by that month that were inaccessible. By September 2018, over 95% of these leading VPNs were inaccessible - either removed by their developers or by the platform.

I also investigate whether Chinese users experienced difficulty using banned foreign apps when VPNs became less accessible. I analyze a sample of 16 banned apps that ranked among the top 1000 on the China iOS app store in 2016 in this analysis. These apps represent the most popular banned apps among Chinese users prior to the regulatory change, and were ranked highly despite being banned by the government. To understand uses' experience with these apps, I randomly sampled 200 user reviews[52] of each app and used the gpt-3.5-turbo model with few-shot prompting to determine whether a review indicated that the user had difficulty accessing or using the app.

[52]Notably, in the iOS app store, writing reviews for banned apps does not require VPNs. Users can post reviews for an app directly via its product page in the iOS app store.

Figure A2a plots the average percentage of reviews complaining about inaccessibility over time. The share of user complaints about the inaccessibility of banned apps rose substantially following the regulatory change, indicating that the censorship introduced a significant barrier to accessing these apps.

To better understand the performance of banned apps relative to permitted ones after the regulatory change, I collected additional data on monthly category-specific ranks for Chinese iOS apps. I then analyzed the rankings of banned apps within the same category.[53] Figure A2b plots the share of banned apps that ranked in the top 300 of their category over time. The figure reveals a clear decline in their relative ranking performance compared to competitors within the same category following the regulatory change. This pattern is robust to alternative cutoffs, such as the top 100 or top 200.

Figure A2: Changes in Accessibility of Censored Apps

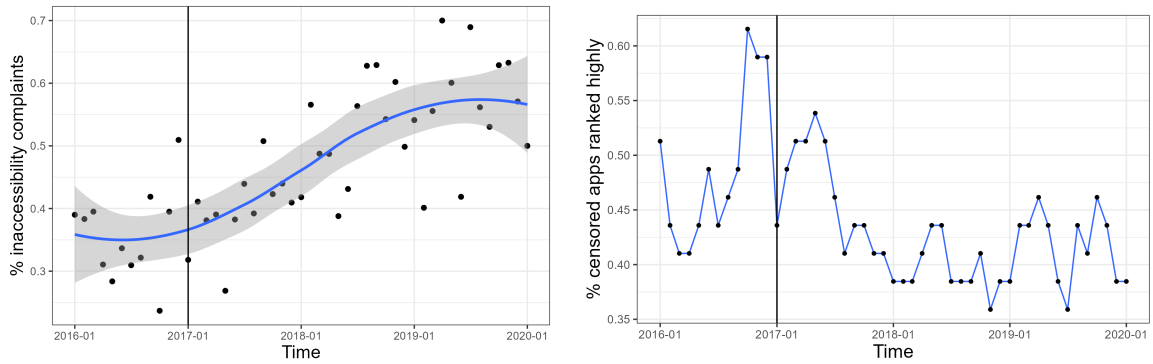

(a) Inaccessibility Complaints of Censored Apps

(b) Category Rank of Censored Apps

Notes: Panel (a) analyzes a random subset of 200 user reviews for each of the 16 banned apps that ranked among the top 1000 in the iOS app rank in 2016. These banned apps represent the most popular banned foreign products among Chinese users prior to the regulatory change. This figure shows the average percentage of user reviews complaining about these apps being inaccessible and shows an increase in the percentage of inaccessibility complaints after the regulatory change. Panel (b) plots the share of banned apps that ranked in the top 300 of their category over time. The figure reveals a clear decline in their relative ranking performance compared to competitors within the same category following the regulatory change. This pattern is robust to alternative cutoffs, such as the top 100 or top 200.

[53]The banned apps I analyze are those that appeared in the top 300 of their category at least once during the study period, representing major players in the market.

A1.2 Additional Definitions and Statistics

Table A1: Distribution of App Categories by Treatment Status

Group	Category	Observations
Treated	Productivity	41203
Treated	E-shopping	38409
Treated	Communities	25381
Treated	Information	20932
Treated	Videos	16271
Treated	Chat	16209
Treated	News	15905
Treated	Novel	9315
Treated	Music	8083
Treated	Livestream	7976
Treated	Friending	7050
Treated	File management	5301
Treated	Short videos	4900
Treated	Photo book	4450
Treated	E-book	4335
Treated	Audio-book	3970
Treated	Telecommunications	3723
Treated	Browsers	2784
Treated	Cloud storage	1657
Control	Maps	11090
Treated	Emails	1198
Control	Tools	41552
Control	Learning	40112
Control	Office software	29741
Control	Examination	25918
Control	Discounts	23353
Control	Medical	23055
Control	Early childhood education	22868
Control	Games (children)	20829
Control	Cars	17425
Control	Car renting	16052
Control	Diet and exercise	15928
Control	Entertainment	14845
Control	Childcare	14224
Control	Stocks	12948
Control	Housing	12470
Control	Career	12176
Control	English	11441
Control	Food	10780
Control	Door-to-door service	10674

Notes: This table lists the treated markets (20 app categories) and the top 20 control markets in terms of the number of observations. The percentages are the percentage of observations within the treated (or control) group from each app category.

Table A2: Android-defined Dangerous Permissions

Permission Type	Permission Name	Permission Description
Phone	Accept_Handover	Allows accepting ongoing handover calls
Phone	Answer_Phone_Calls	Grants the ability to answer incoming calls
Phone	Call_Phone	Enables initiating phone calls
Phone	Read_Phone_Numbers	Access to read phone numbers
Phone	Read_Phone_State	Read current phone state
Phone	Use_SIP	Use Session Initiation Protocol for VoIP calls
Phone	Add_Voicemail	Add voicemail messages
Storage	Access_Media_Location	Access media files with location information
Storage	Read_External_Storage	Read external storage contents
Storage	Read_Media_Audio	Read audio files
Storage	Read_Media_Images	Read image files
Storage	Read_Media_Video	Read video files
Storage	Write_External_Storage	Write to external storage
Microphone	Record_Audio	Record audio using microphone
Camera	Camera	Access the camera
Sensors	Body_Sensors	Access data from body sensors
Sensors	Body_Sensors_Background	Access body sensor data in the background
Calendar	Read_Calendar	Read calendar events and details
Calendar	Write_Calendar	Create, modify, or delete calendar events
Location	Access_Background_Location	Access device location in the background
Location	Access_Coarse_Location	Access approximate device location
Location	Access_Fine_Location	Access precise device location
Contacts	Get_Accounts	Access device accounts
Contacts	Read_Contacts	Read contacts
Contacts	Write_Contacts	Modify or delete contacts
Call log	Process_Outgoing_Calls	Monitor, modify, or abort outgoing calls
Call log	Read_Call_Log	Read call history/log
Call log	Write_Call_Log	Modify or delete call history/log
SMS	Read_SMS	Read SMS messages
SMS	Receive_MMS	Receive and process MMS messages
SMS	Receive_SMS	Receive and process SMS messages
SMS	Receive_WAP_Push	Receive and process WAP push messages
SMS	Send_SMS	Send SMS messages
SMS	Read_Cell_Broadcasts	Read cell broadcast messages

Notes: This table summarizes the permission types (groups) of Android system-defined dangerous permissions obtained from the Android developer manual on December 30th, 2022. At the time of our data collection, three additional permission types, 'Activity recognition', 'Nearby devices' and 'Notification' also contain dangerous permissions, but all the corresponding dangerous permissions were introduced after the time of the treatment (January 2017). I therefore exclude these permission types from the analysis in Figure **??**.

A1.3 Additional Evidence for Heterogeneity Analysis

Figure A3: Treatment Effect by App Genre

(a) Estimated effect on the total number of permissions

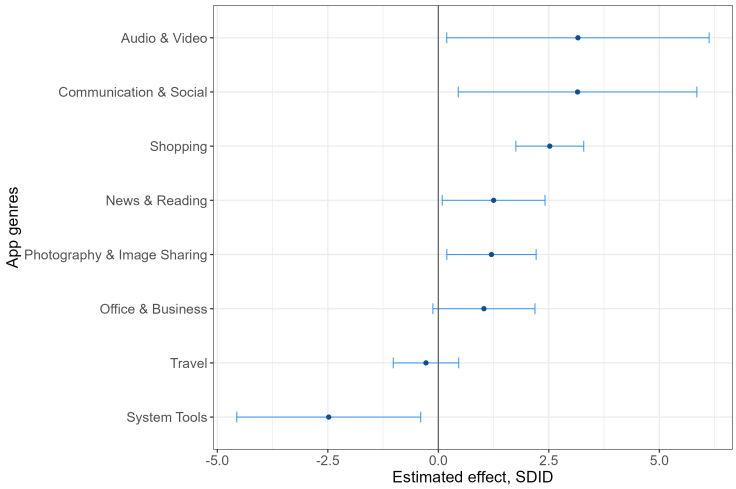

(b) Estimated effect on the number of sensitive permissions

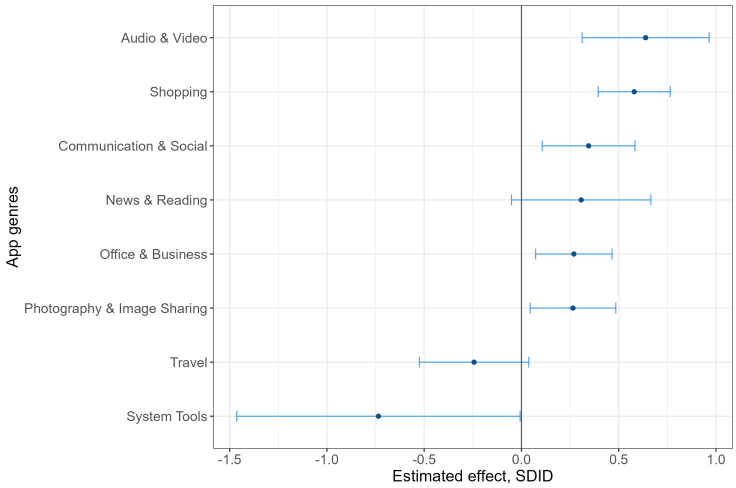

Notes: Panel (a) presents SDID estimates of the effect on the number of permission requests for each app genre. Panel (b) presents the SDID estimates of effect on the number of sensitive permissions requests for each app genre. Genres are defined by the app store I study and contain several app categories (markets) that are functionally related. If a genre has both treated and control categories, I use all control apps, including those in the same genre and those in different genres, to construct the synthetic control group for treated apps in this genre.