Are You 18? Age Verification and Adult-only Consumption in the US *

Saharsh Agarwal[†] Uttara M. Ananthakrishnan[‡] Leonardo Madio[§]
Matthew Mitchell[¶] Martin Quinn[‖] Carlo Reggiani[**]

April, 20 2026

Several U.S. states now require age verification for accessing online pornography. We compile domain–state–week traffic for major adult platforms across fifteen states and estimate the causal impact of age verification laws using a difference-in-differences design with staggered adoption. In treated states, platforms that implemented age-verification gates experienced large declines in observed visit (e.g., xHamster), while large free-access sites that remained ungated during the sample period (e.g., XNXX, XVideos) gained 35–40%, consistent with substantial cross-platform substitution. By contrast, registration- and payment-based platforms (e.g., OnlyFans, Stripchat) exhibit smaller changes, consistent with lower marginal frictions when users are already authenticated. Overall, the evidence suggests that age-verification laws can reallocate observed traffic across major platforms in our sample while reducing visits to adopting and blocked domains.

Keywords: Age Restriction, Platforms, Adult Websites, Search.
JEL Code: D83; K42; L82; O39.

*We thank Maximilian Schaefer, Mark Tremblay and seminar participants at the Paris Digital Conference and UK Ofcom for helpful comments and discussion. The usual disclaimer applies. Leonardo acknowledges financial support under the project "SAFE-net: Online Safety and Moderation in Digital Platforms" under the "Supporting Talent in Research @ University of Padua (STARS@UNIPD) programme, 2025". The views and opinions expressed in this paper are the authors' and do not necessarily reflect those of the Joint Research Centre or the European Commission.

[†] Indian School of Business, India. Email: saharsh_agarwal@isb.edu.
[‡] University of Washington, United States. Email: uttara@uw.edu.
[§] University of Padua, Italy. Email: leonardo.madio@unipd.it.
[¶] European University Institute, Italy. Email: matthew.mitchell@eui.eu.
[‖] Erasmus University Rotterdam, Netherlands. Email: quinn@rsm.nl.
[**] European Commission's Joint Research Centre, Spain and University of Manchester, UK. Email: carlo.reggiani@ec.europa.eu.

1 Introduction

The scale of online content markets generates substantial value (Brynjolfsson et al., 2019) but also exposes users, especially minors, to misinformation, harmful speech, cyberbullying, and potentially addictive overuse (Allcott et al., 2020, 2022), with consequences that may be amplified during childhood development (Manolios et al., 2025). Digital policy is increasingly oriented toward ensuring online safety as witnessed, for example, by the EU's Audiovisual Media Services Directive (AVMSD) and Digital Services Act (DSA) frameworks, the U.S. Children's Online Privacy Protection Act (COPPA), and the UK and Australian Online Safety Acts.[1] These efforts reflect a broader turn from regulating online content to regulating online access.

A central manifestation of this shift is the rapid adoption of age assurance requirements—rules that condition access to entire classes of online services on a credible signal of a user's age. Age assurance can take place through various mechanisms, including age verification (such as government ID checks) and age estimation (such as facial age analysis). Such measures are now proposed or implemented across a range of settings — from social networks and user-generated content to online gambling, app ecosystems, and adult-only material — with recent interventions increasingly targeting minors' access to social media in countries such as Australia, France, and Spain.[2]

In this paper, we focus on the effects of age verification, a mechanism to implement age assurance, in the adult entertainment industry. As of mid-2025, 19 U.S. states have enacted laws requiring age verification to access adult content, with a total of 24 states having passed such legislation, including those with pending implementation (for simplicity, we will refer to them as "age verification laws").[3] The stated goal of these policies is straightforward: age verification promises targeted protection of minors while preserving lawful access for adults. Yet it also changes the economics of online markets by introducing new frictions, including adoption costs and data-collection requirements for platforms and hassle and privacy costs for users, thereby shaping platform strategies and potentially triggering user substitution across services. Understanding these equilibrium responses is essential for evaluating both the effectiveness of child-protection regulation and its possible implications for competition. We leverage the

[1] More specifically, in the EU, the DSA strengthens protections for minors (including restrictions on profiling-based advertising to minors and, for very large platforms, systemic risk assessment and mitigation duties) and complements the AVMSD, which requires video-sharing platforms to adopt measures protecting minors from harmful or age-inappropriate audiovisual content; the DSA is broadly content-neutral and does not itself define illegality, leaving such determinations largely to other Union or Member-State law. In the U.S., the COPPA requires verifiable parental consent before operators of child-directed services—or services with actual knowledge—collect, use, or disclose personal information from children under 13, motivating age-screening/parental-consent workflows or minimum-age thresholds on general-audience social media. In the UK, the Online Safety Act imposes child-safety duties and, for services that allow pornography, requires "highly effective" age assurance (in force from 25 July 2025); Australia's Online Safety Act similarly strengthens platform duties and enforcement regarding online harms to children.

[2] France has pursued restrictions for under-15s (including parental-consent requirements and more recent legislative proposals), while Spain has advanced proposals to bar social media access for under-16s and to require age-verification systems.

[3] These state-level laws differ in enforcement and scope but broadly mandate either government ID checks or third-party verifications for access to pornographic sites. See, e.g., https://web.archive.org/web/20250628105357/https://www.axios.com/2025/01/16/adult-website-age-verification-states.

staggered implementation of different age verification laws in the United States to estimate the causal impact of these laws on traffic to pornographic sites.

Studying this industry is important for several reasons. First, despite its substantial economic scale and exceptionally high traffic, with a few oligopolistic firms and a fringe of smaller sites, it is underexplored in the economics literature. Second, it is a domain in which the protection of minors is a critical policy objective and, even beyond minors, exposure may generate negative externalities. Third, platforms' responses to age verification laws in the US are quite heterogeneous both among firms with similar business models and across distinct models, thus providing a rare opportunity to study the resulting traffic patterns. Fourth, age verification inherently raises privacy and security trade-offs, because effective enforcement may require the collection, processing, or retention of sensitive user data.

We use web traffic data from Similarweb, a data analytics company, covering fifteen U.S. states — including five (Florida, Georgia, North Carolina, Texas, and Virginia) that adopted age verification laws between July 2023 and July 2025 — and a difference-in-differences (DID) design to estimate the impact of these laws on traffic in the adult entertainment sector.

We categorize platforms into three groups according to their response to the laws: those that are *adopting* age verification, those that *exit* the market by blocking access to consumers in the state, and those that remain *non-adopting* in our sample period. Among the four dominant free-streaming platforms ("tubes") in the sector, XVideos and XNXX remained accessible in the affected jurisdictions and do not appear to have implemented discernible changes to their user interfaces or access protocols, a pattern that also characterizes many smaller sites. We refer to these sites as "non-adopters" of age-verification tools. By contrast, Pornhub, the largest platform globally, restricted access to users located in adopting states by shutting down its portal and explicitly cited the lack of privacy-preserving age-verification technologies and the risk of identity theft as reasons for not complying with the statutes.[4] We refer to these platforms as "exiters". xHamster, instead reported introducing a registration-based verification mechanism to satisfy regulatory requirements, while live-streaming cam sites and membership platforms such as OnlyFans implemented age verification within their existing account-based access models.[5] The latter sites are "adopters" of age verification tools.

Leveraging time variation in the enactment of age verification laws across states, we estimate the effect of the age verification laws on domain-state online traffic (i.e., visit and bounce rates) using recently developed estimation techniques that appropriately identify the average treatment effects in presence of staggered adoption (Callaway and Sant'Anna, 2021).

The results are as follows. First, the overall traffic to adult websites has decreased 23% in the treated states compared to control ones following the implementation of age verification laws. Second, platforms that exited a state lost about 55–56% of their traffic. Perhaps not surprisingly, this is a sizable impact, although visits did not go to zero. On the one hand, users still tried to

[4] See https://www.pornhub.com/blog/age-verification-in-the-news.

[5] Most U.S. state laws permit credit card checks as a compliant age assurance tools. Subscription-based platforms that already collect payment information may therefore have been compliant without material changes to their user interface, which could attenuate the adoption costs, as well as the hassle costs for their users.

connect to the platforms and they were probably bounced back by the exit announcement. On the other hand, they may have found alternative access methods such as VPNs (Lang et al., 2025).[6] A substantial negative effect is also identified on the traffic of adopting platforms, which lose about 36% of their visits. These results are consistent with a number of factors. For example, users may perceive high privacy[7] costs and stop visiting the website rather than submitting to age verification processes, which in most adopting states primarily rely on government-issued ID checks, though other methods such as facial age estimation are also permitted. Moreover, it is possible that the introduction of age verification could have created a "hassle cost" for users, linked to the time involved to verify their age, inducing some users to move to alternative non-adopting sites. Finally, and consistently with the legislation's goal, minors may have also been blocked from browsing adult websites.[8]

Third, and perhaps surprisingly, visits to non-adopting sites did not seem to change substantially on average. Heterogeneity analysis among non-adopting platforms, however, indicates very large traffic gains concentrated among two mainstream sites (i.e., XNXX and XVideos), which gained 40% of traffic (relative to their counterfactuals), reflecting diversion from adopting competitors. Fringe platforms do not show robust evidence of positive diversion; the estimated effect is small, negative, and statistically significant only under the never-treated counterfactual.[9]

Age verification laws involve an important trade-off. They aim to protect minors from legal but age-inappropriate sexual content, yet implementing them often requires platforms or third parties to collect, and in some cases store, users' identity data— with retention practices varying across implementations — raising potential privacy and security concerns.[10] The risks associated with storing sensitive user data are not merely theoretical: data breaches affecting platforms holding identity or behavioral data can expose users to significant harm.[11]

[6]Note that while users may activate VPNs to circumvent access restrictions, our data provider associates their traffic with their initial daily location connection, i.e., the location recorded before VPN activation. As a result, users who activate a VPN after their initial connection will still appear in our data as traffic originating from their actual state, meaning VPN usage does not systematically distort our state-level traffic measurements. For more details, see: https://support.similarweb.com/hc/en-us/articles/6315589042077-Country-Filter-FAQs.

[7]The term 'privacy' is used here in the sense common in U.S. legal and policy discourse, referring broadly to concerns about unauthorized collection and use of personal data. This differs from the EU framework, where the right to private and family life (Article 7 EUCFR) and the right to personal data protection (Article 8 EUCFR) are treated as distinct fundamental rights.

[8]Our data provider states that "our anonymous behavioral data includes people of all ages, but our estimates are based on statistical models about users, not demographics about each user, to maintain our privacy standards". For details, see: https://support.similarweb.com/hc/en-us/articles/360001253797-Website-Demographic-Data-FAQs#h_01H9TQ6K23PJ7TVX0QB8ZMGB92.

[9]If any, fringe sites experienced a reduction of approximately 6% of traffic relative to never treated sites, but this coefficient is only weakly significantly different from zero.

[10]The extent of these concerns depends significantly on the technological solution adopted, as emerging privacy-preserving approaches may reduce but not eliminate the tension between effective age-verification and user privacy. In jurisdictions such as the EU, the GDPR's data minimization principle would indeed constrain how age verification can be implemented, potentially pushing towards more privacy-preserving solutions.

[11]As an example, on 12 December 2025, following an extortion claim, Pornhub disclosed a third-party analytics incident that exposed a limited subset of Pornhub Premium user-activity records (possibly including email addresses and time-stamped viewing/search history). See: https://www.forbes.com/sites/zakdoffman/2025/12/17/pornhub-confirms-breach-user-search-watch-and-download-activity-stolen/. Furthermore, in 2025 a Discord breach exposed IDs submitted for age checks, showing how badly implemented age verification systems can themselves create new risks. See: https://www.pcmag.com/news/discord-70k-govt-ids-exposed-breach-hackers-posting-age-verification-selfies.

To assess whether the observed heterogeneity is consistent with hassle and privacy costs, we examine how the traffic effects of age verification laws vary across adopting platforms by business model. We find that platforms with pre-existing registration requirements—such as OnlyFans and live-cam sites, where users typically subscribe, tip creators, or otherwise maintain authenticated accounts—experienced smaller, and in some cases statistically insignificant, declines in traffic after the policy. A plausible interpretation is that formal age verification imposed relatively limited additional friction on these users, since they were already accustomed to disclosing personal information and payment details in order to access or transact on these platforms.

These findings have important implications, as they seem to point to a reduction in access to pornographic content following the implementation of age verification laws. This result can be contrasted with other policies, such as private enforcement of content moderation by Pornhub. According to Madio et al. (2025) this private enforcement, which led to the removal of unverified content from the platform, led to a reallocation of traffic from the dominant platform to other players, including small, less regulated, and arguably less safe platforms. However, the difference in the results is due to different underlying mechanisms. In the case of content moderation by Pornhub, content was removed, and consumers had to search for alternatives. In contrast, age verification leaves content accessible but places it behind an authentication wall, creating hassle and privacy cost. This difference may explain the incomplete diversion of users to non-adopting sites that we observe in the case of age verification. Moreover, in our context, the diversion is not driven by search results or referrals, unlike the case of content removal.

Notwithstanding cross-country variation in regulatory design and enforcement, our analysis can offer useful implications for the implementation of similar age-assurance measures in other contexts. In April 2025 France has enacted the Sécurité et Régulation de l'Espace Numérique (SREN) framework, then suspended by a court in June.[12] In Italy, age assurance has been implemented for Italian porn sites in November 2025, and foreign sites have to be compliant by 2026. In the UK, as of 25 July 2025, through changes to the Online Safety Act, all sites and apps that allow pornography need to implement highly effective age assurance measures, to ensure children cannot access pornography or other harmful content.[13]

Related Literature. Closely related to our work, there are a few recent studies focusing on age verification in the U.S. Unlike this paper, these studies focus on search indices data from Google Trends and infer patterns on user engagement. For example, both Lang et al. (2025) and Spencer (2025) provide evidence suggesting substantial switching in searches from adopting to non-adopting sites, as well as an increased interest in VPNs. Spencer (2025) also find that the laws did not affect search interest for hentai, the most popular porn category among youngsters. Although this early evidence is of great interest, it also has some limitations. Above all, a significant share of users, particularly on major platforms, accesses these websites directly by

[12] See, e.g., https://www.theguardian.com/world/2025/jun/20/pornhub-back-online-france. The European Commission itself is working with some EU member states to develop a privacy-protecting online age assurance tool. See: https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1820.

[13] https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-for-online-safety–what-you-need-to-know-as-a-user.

typing URLs or via bookmarks, bypassing search engines altogether (as we show in Section 6.3). Our work goes beyond search behavior and interest in these platforms by analyzing the actual consumption and exposure to adult content, as well as the substitution patterns across platforms. On top of that, we also provide evidence that heterogeneity in business models can mitigate traffic diversion: in particular, sites that rely on memberships and direct payments to creators do not experience, or experience only much smaller, traffic losses to rival platforms.

While age verification laws are in this case a demand-side intervention[14]—where content is considered moderated for some groups of users—our analysis complements previous works that have focused on supply-side content moderation in the adult entertainment market. Madio et al. (2025) document that self-enforcement of content moderation, through the removal of non-verified content by Pornhub (part of the Aylo group, previously named MindGeek) and its sister sites, led to a substantial demand reaction, with a staggered shift to alternative sites with minimal oversight. The present paper further highlights that demand-side interventions can induce a reallocation of traffic—albeit partial—driven by potential privacy concerns, hassle costs, and the local shutdown of dominant platforms. Notably, the traffic reallocation benefited large mainstream sites that resisted the adoption of the age verification laws.

This is not the first study to examine the effects of laws aimed at protecting minors. Johnson et al. (2026) and Kircher and Foerderer (2024) investigate the unintended consequences of legislation banning targeted ads to children, focusing respectively on YouTube content supply and app development. In our context, age verification laws are shown to have heterogeneous effects depending on the business model of the firm and can lead to a reallocation of traffic to alternative sites that do not implement such age-verification tools.

Our work also relates to a recent stream of literature on platform governance and content moderation, which include both theory (see, e.g., Liu et al. 2022; Kominers and Shapiro 2024; Bar-Isaac et al. 2025; Beknazar-Yuzbashev et al. 2024; Rendo 2024; Madio and Quinn 2025) and empirical works (see, e.g., Agarwal et al. 2022; Beknazar-Yuzbashev et al. 2022, 2026; Kumarswamy et al. 2023; Zeng et al. 2022). We contribute to this literature by providing first systematic evidence of the effects of interventions intended to target a selected group of users (e.g., minors).

Finally, our work also relates to the literature on the role of online platforms in the sex industry, which has mostly focused on the effects of the entry of Craigslist in the market and the effects on prostitution and HIV (Chan et al., 2019; Zeng et al., 2022). More recently Zeng et al. (2022) has studied the impact of the shutdown of Backpage and Craigslist personal ad on prostitution arrests and violence against women.

Structure. The rest of the paper is structured as follows. Section 2 provides the background about the industry and the regulatory scenario. Section 3 introduces our data and preliminary

[14]Note also that in some jurisdictions, internet service providers for mobile devices also perform the age check when the contract is concluded and filters are applied by the internet service provider for the device if the user is underage.

descriptives. Section 4 outlines our empirical strategy. Section 5 presents our main findings. Section 6 presents additional heterogeneity results. Finally, Section 7 concludes.

2 Background

The age verification laws. Between 2023 and 2025, a growing number of U.S. states enacted legislation requiring adult-content websites to verify that users are at least 18 years old before granting access. As of mid-2025, at least 19 states have such laws *in force*. Table I summarizes enacted verification legislation.

While statutory language varies, most laws require platforms hosting a substantial share of content deemed "harmful to minors", typically defined as one-third or more, to implement *reasonable* age verification, often through government-issued ID or third-party credentialing services. Some states have gone further, granting private rights of action (e.g., lawsuits initiated by parents of minors) or empowering Attorneys General to impose civil fines for non-compliance.[15]

Additional states (including Arkansas, South Dakota, North Dakota, Georgia, and Arizona) have passed similar laws scheduled to take effect in late 2025. In some cases, enforcement has been delayed or temporarily blocked by judicial injunctions.

The market. The online adult-content market is highly concentrated and organized around distinct business models. "Tube" sites operate largely on a freemium model, offering large catalogs of videos that are directly accessible, with monetization primarily through advertising and optional premium subscriptions. "Cam sites" are live-streaming platforms where users can passively watch for free, but the core economic activity is performer–user interaction via tips, tokens, or private shows, which typically requires registration and payment credentials. Membership platforms such as OnlyFans instead rely on subscription and pay-per-view relationships between users and individual creators, with access conditional on account creation and recurring payments.

In what follows, we provide details about our main market players and how they responded to state interventions. A summary is also provided in Table I

2.1 Free Tubes

A handful of platforms capture the majority of global traffic among the free tubes, most notably those operated by *Aylo* (formerly MindGeek), including *Pornhub*, *YouPorn*, and *RedTube*. Aylo is headquartered in Montreal, with affiliates in Luxembourg, it is estimated to generate several hundred million dollars annually, and ranks among the world's most visited web properties.

XVideos, XNXX, and xHamster, like Pornhub, are large "tube" platforms that operate as multi-sided, user-generated content sites in the online adult industry. Their core model is to provide

[15] `https://www.washingtonpost.com/technology/2023/05/03/pornhub-blocks-utah-vpn-searches/`.

Table I: Platform responses to age-verification laws (as of Aug. 2025)

State	Law / Bill	Effective date	PH	XV	XN	XH
Louisiana	Act 440 / HB 77	Jan 1, 2023 (HB 77: Aug 1, 2023)	A/B	NA	NA	A
Utah	SB 287	May 3, 2023	B	NA	NA	A
Mississippi	SB 2346	July 1, 2023	B	NA	NA	A
Arkansas	SB 66	Aug 1, 2023	B	NA	NA	A
Virginia	SB 1515	July 1, 2023	B	NA	NA	A
Texas	HB 1181	Sept 1, 2023	B	NA	NA	A
Montana	SB 544	Jan 1, 2024	B	NA	NA	A
North Carolina	HB 8	Jan 1, 2024	B	NA	NA	A
Idaho	H 498	July 2024	B	NA	NA	A
Kansas	SB 394	July 2024	B	NA	NA	A
Kentucky	HB 278	July 2024	B	NA	NA	A
Nebraska	LB 1092	July 2024	B	NA	NA	A
Indiana	SB 17	Aug 2024	B	NA	NA	A
Oklahoma	SB 1959	Nov 1, 2024	B	NA	NA	A
Alabama	–	Oct 1, 2024	B	NA	NA	A
Florida	HB 3	Jan 1, 2025	B	NA	NA	A
South Carolina	HB 3424	Jan 1, 2025	B	NA	NA	A
Tennessee	SB 1792	Jan 1, 2025	B	NA	NA	A
Wyoming	HB 43	Mar 13, 2025	B	NA	NA	A
Missouri	Rule 15 CSR 60–18	May 7, 2025	B	NA	NA	A
South Dakota	HB 1053	July 2025	B	NA	NA	A
Georgia	–	July 1, 2025	B	NA	NA	A
North Dakota	HB 1561	Aug 2025	B	NA	NA	A

Notes. PH = Pornhub, XV = XVideos, XN = XNXX, XH = xHamster. **A** = Adopting; **B** = Blocked; **NC** = Not Adopting.

free, streaming access to short- and long-form video clips, typically financed by advertising (display, pre-roll, and native ads), traffic resale, and data monetization, with optional paid tiers (e.g., premium subscriptions) offering higher quality, exclusive content, and reduced advertising.

Despite this concentration at the top, the broader ecosystem comprises thousands of small, independent, and often anonymous websites operating with little oversight or regulatory compliance. For large platforms, integrating verification systems, while costly, is technically and legally feasible. In contrast, smaller sites (fringe sites) often lack the incentive to invest in compliant systems for data security, user authentication, or legal risk management, as their competitive positioning may depend precisely on avoiding such restrictions. Many fringe sites, including those hosting pirated or extreme content, frequently flout these obligations, raising concerns that stricter regulations could unintentionally drive traffic toward unregulated or illegal online spaces.

Aylo's platform withdrawal strategy. Aylo has publicly endorsed the principle of age verification but has consistently criticized the patchwork of state-level mandates, which expose platforms to privacy and liability risks.[16] In Louisiana, Pornhub chose to comply using the LA Wallet system, a state-authorized digital ID app, becoming the only major site to maintain access while verifying user age. Nonetheless, traffic from Louisiana reportedly dropped by as

[16] `https://www.pornhub.com/insights/statement-on-utah-age-verification-law`.

much as 80% following the law's implementation.[17] Based on this outcome, Aylo blocked access entirely in other states, including Utah, Mississippi, and Virginia, displaying protest messages to visitors. Similar action was taken in the UK in February 2026.[18] Figure I, panel (a), displays the message explicitly stating that the company could not implement age verification without jeopardizing data security.

Age verification by other platforms Adoption patterns among major adult sites have been mixed. xHamster implemented age-verification tools to comply with the applicable laws. Yet, Texas authorities filed suit against the company for alleged non-compliance, signaling both enforcement pressure and partial efforts by the platform to operate within the legal framework.[19]

By contrast, XVideos and XNXX have continued to operate in treated states without implementing visible age-verification tools, and they have been part of the broader industry coalition challenging these statutes in court. In Texas, the Free Speech Coalition (FSC) represented major adult sites, including XNXX, in *Free Speech Coalition, Inc. v. Paxton*, where the U.S. Supreme Court ultimately upheld the state's verification law (HB 1181) in June 2025.[20] In Florida, FSC and co-plaintiffs initially filed a federal challenge against HB 3's age verification mandate but later moved to dismiss the case in light of the Supreme Court's ruling in *Paxton*. Following this dismissal, state authorities signaled their intention to pursue civil enforcement against sites allegedly serving Florida users without adopting age checks.[21]

Figure I, panel (b), presents the homepage of xHamster in Texas, where it instructs users to submit ID or biometric data for confirmation. These screenshots provide direct evidence of the heterogeneous adoption responses observed across major platforms: while Aylo platforms withdrew from regulated jurisdictions, xHamster adapted by implementing verification technology and maintaining access within the legal framework.

2.2 Live-cam sites and membership platforms

Free tubes also co-exist with other types of platforms. In the live-cam segment, Stripchat, Chaturbate, and xHamsterLive (the live-cam service of xHamster) present a different adoption path from large "tube" sites. Stripchat was aligning with the laws and required to complete an ID-based "Verify Age" process or equivalent third-party check to access streams. Chaturbate and xHamsterLive also complied, although they were explicitly targeted by the Texas Attorney

[17] `https://www.theverge.com/2023/1/10/23547863/pornhub-age-verification-law-louisiana-traffic-drop`.

[18] https://www.aylo.com/newsroom/aylo-upgrades-age-assurance-methods-in-united-kingdom/

[19] `https://www.texasattorneygeneral.gov/news/releases/attorney-general-paxton-files-lawsuit-against-pornographic-website-xhamster-violating-texas-law`.

[20] See, for example, the Supreme Court's opinion in *Free Speech Coalition, Inc. v. Paxton*, No. 23–1122 (June 27, 2025), available at `https://www.supremecourt.gov/opinions/24pdf/23-1122_3e04.pdf`, and contemporary coverage in, e.g., Reuters and *The Guardian* on the decision and its implications for major adult sites such as Pornhub and XNXX.

[21] On the Florida litigation and subsequent dismissal, see the FSC's announcement of its lawsuit against HB 3 (`https://www.freespeechcoalition.com/blog/free-speech-coalition-sues-florida-to-block-age-verification-law`) and later reporting on the withdrawal of the case, e.g., WUSF ("Plaintiffs drop challenge to Florida's age verification law for porn sites," July 10, 2025) and *Florida Politics*/*USA Today* coverage of the same development.

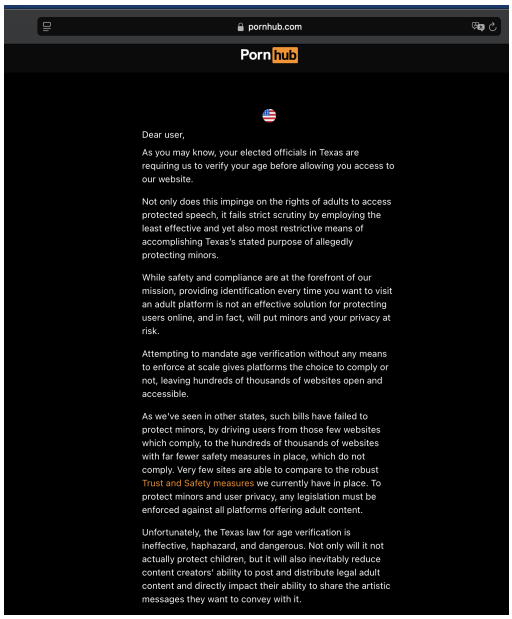

(a) Pornhub (Aylo) shut down in Texas.

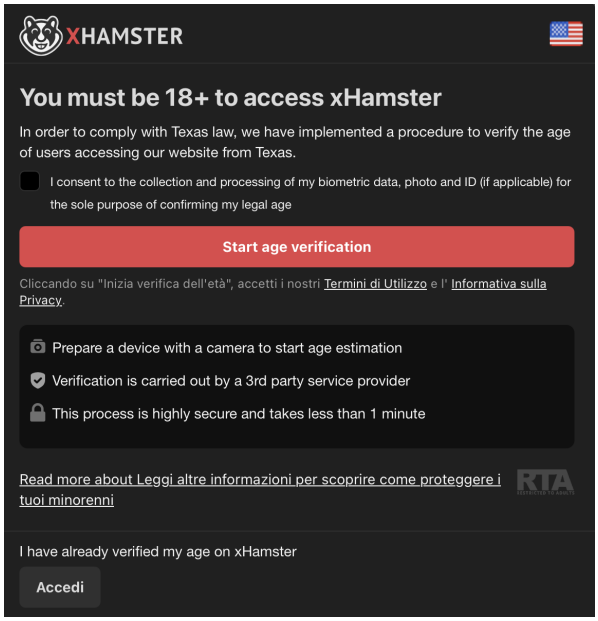

(b) xHamster's adoption in Texas.

Figure I: Adoption responses to the Texas age-verification law. Pornhub (panel a) exited the state by blocking access to the video-library and showing an open letter to viewers, whereas xHamster (panel b) introduced an age-verification gate.

General under HB 1181 for allegedly failing to implement "reasonable" age-verification safeguards. For the purpose of this study, we will consider these platforms "adopting',' as they implemented some changes in the platform design, even though not fully compliant with state-specific regulation.[22]

OnlyFans operates under a rather different business model, requiring membership payments and subscriptions. As a result, access already requires registration and creators must pass a relatively stringent verification process (government ID, selfie, 18+ requirement), so the impact of new age verification laws on its user flow may be negligible or even zero, as the platform already meets or closely approximates the requirements such laws impose.

Unlike Aylo's explicit open letters and state-specific shutdown announcements, there do not appear to be high-profile public statements by OnlyFans or Stripchat specifically addressing individual U.S. state pornography age verification laws; their responses are instead observable via on-site verification flows and regulatory proceedings.

3 Data

As a main source for our analysis, we use data from a data aggregation and software development company specializing in web analytics, web traffic, and performance, Similarweb

[22]Texas sued the parent companies Multi Media LLC (Chaturbate) and Hammy Media (xHamster), seeking substantial civil penalties for past non-compliance. Chaturbate's operator subsequently entered into a settlement with Texas, paying a $675,000 penalty and committing to deploy age-verification software for Texas users, which court filings state was implemented promptly after the lawsuit.

(similarweb.com). The provider collects information on website traffic volume, as the total number of visits to each domain in each location, as well as several measures of user online behavior.

The provider collects data along a number of relevant metrics, including[23] *Visits*, that is the total visits to a domain through all traffic channels. A visit is a user's access to one or more pages of a website. All subsequent page views belong to the same visit until the visitor is inactive for 30 minutes.

Our unit of observation is a state-domain pair at the week level. In the paper, we use the terms "domain" and "website" interchangeably. Our main outcome is *visits*, capturing the number of visits in the state-domain pair on a specific week.

We focus on the top 24 adult domains according to our provider. We collect audience data per domain-state in 15 states between June 2022 and July 2025 at the week level. Following Table I, we define a period *After* equal to 1 after age verification went into application for each treated state.

Adult domains respond heterogeneously to the laws. Empirically, we observe three mutually exclusive adoption strategies at the domain–state level: (i) adoption, where the site remains accessible in the state behind an age-verification gate, (ii) exit/block, where the site becomes inaccessible from that state, and (iii) non–adopting, where the site remains accessible without implementing age-verification tools.

Table II summarizes how web traffic changed around the implementation of age verification laws across five adopting states (Florida, Georgia, North Carolina, Texas, and Virginia).[24] Visits declined sharply for both adopting and exit platforms, while traffic to non-adopting domains remained largely unchanged. Adopting sites lost roughly one-third of their pre-treatment visits, and exit platforms more than three-quarters, indicating substantial deterrence where users faced new verification barriers or outright access blocks. Non-adopting domains, which continued to operate without restrictions, show a small increase in visits, consistent with a minor reallocation of users avoiding verification procedures.

Table II also reports average bounce rates, defined as the share of visits in which users leave the site after viewing only the landing page. For all domains combined, the bounce rate increases from 0.25 before to 0.28 after implementation. The most pronounced change occurs for exit platforms, whose bounce rate rises from 0.34 to 0.52, consistent with users reaching a static shutdown or information page and then immediately leaving without further navigation. By contrast, adopting sites exhibit only a small increase (from 0.29 to 0.30), and non-adopting sites see a modest decline (from 0.20 to 0.18), suggesting slightly more within-site browsing among users who continue to access those domains.

Figure II plots unadjusted average weekly visits for adult domains in treated states, aligned to each state's law implementation week (event time $t = 0$; red dashed line). The three panels

[23] For more information about the variables available through Similarweb, one can refer to the API documentation at https://developers.similarweb.com/docs/similarweb-web-traffic-api.

[24] We do not report data about other states in our dataset since these states are not treated and, therefore, any distinction between before and after policy would not be informative.

Table II: Average visits and per-visit values by compliance status

	Before		After			
Group	Visits	B. rate	Visits	B. rate	States	Domains
Adopting	231,948.29	0.29	154,720.01	0.30	5	5
All domains	130,307.90	0.25	92,437.04	0.28	5	24
Exit	153,124.49	0.34	37,757.52	0.52	5	5
Non-adopting	85,858.97	0.20	89,721.52	0.18	5	14

Notes: States and Domains report the number of unique states and domains contributing to each compliance-status cell.

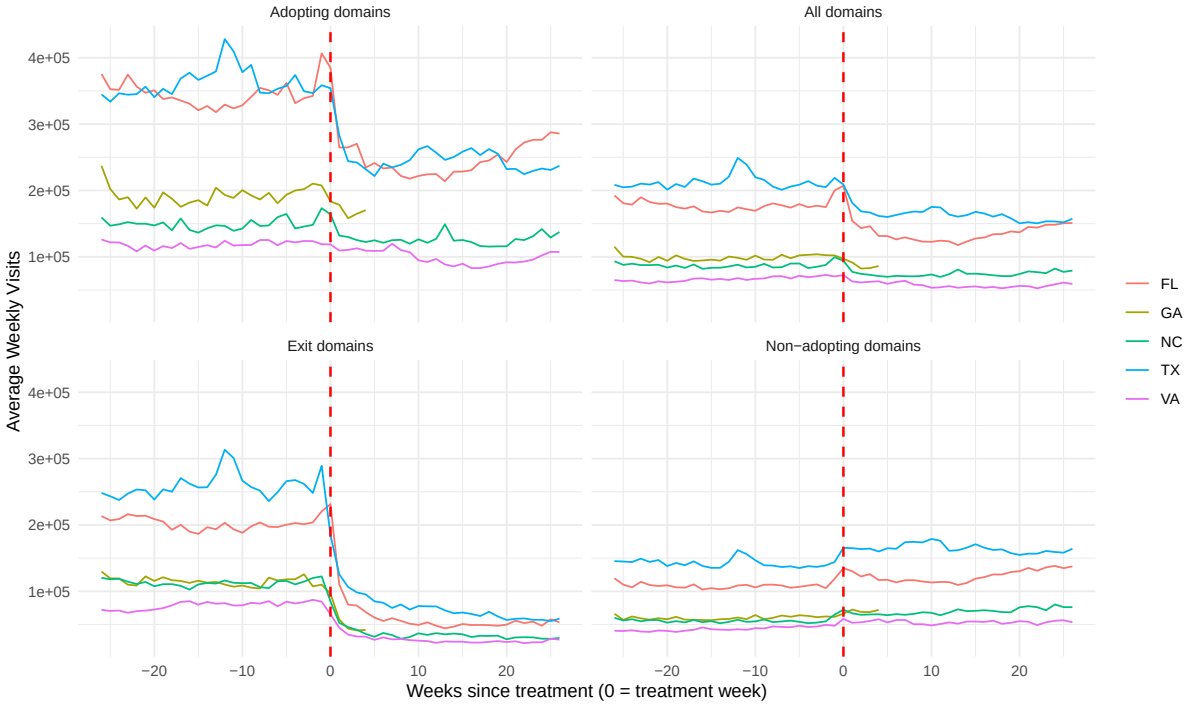

Figure II: Average weekly visits by adoption status (Centered at Treatment Date).

group domains by their observed response to the policy. Within each panel, each colored line corresponds to a treated state (FL, GA, NC, TX, VA) and reports the average of weekly visits across domains in that response group.[25]

Three patterns are evident. (i) For *Adopting* domains, visits drop markedly at $t = 0$ and then partially recover over subsequent weeks, stabilizing below pre-implementation levels. (ii) For *Exit* domains, visits fall sharply at $t = 0$ and remain low thereafter; the residual post-period traffic likely reflects attempts to access blocked sites and/or measurement frictions (e.g., VPN use or imperfect geolocation). (iii) For *Non–adopting* domains, there is no comparable discontinuity at $t = 0$ on average, with trajectories that are flat or mildly increasing before and after implementation. These plots are descriptive averages (not seasonally adjusted and without fixed effects) and motivate our staggered difference–in–differences/event–study analysis below.

4 Empirical strategy

In this section, we describe our empirical strategy. As discussed, our unit of observation is a domain–state–week triple (i, s, t) and we use as outcome variable Y_{ist} the (log) weekly visits during week t to domain i from state s.

Our identification leverages variation in (i) the timing of the law implementation across states and (ii) the adoption strategy chosen by each domain in a given state. We assume that, absent age verification law in state s, within-state traffic to adult domains in s would have followed parallel trends to adult domains in not-yet-treated states and/or never-treated states. Furthermore, users in treated states are assumed not to change their behavior in anticipation of the policy change. Finally, we assume no spillovers between treated and control states—that is, the law in state s does not affect user behavior in other states. A potential threat to this assumption is VPN usage to circumvent restrictions. However, our data provider tracks users through installed software, meaning VPN users attempting to access content by appearing to be in a control state would still be counted as visits from their treated state of origin.[26] Because adoption is staggered, we follow the state-of-the-art difference-in-differences estimator of Callaway and Sant'Anna (2021), which is robust to heterogeneous treatment effects across cohorts.

As a primary specification, we implement an intent-to-treat (ITT) design with staggered adoption. Let G_{is} denote the first week in which domain–state pair (i, s) is exposed to age verification, with $G_{is} = 0$ for never-treated observations. On the adult-domain sample, we estimate the group–time average treatment effect

$$ATT(g, t) \;=\; \mathbb{E}\big[Y_{ist}(1) - Y_{ist}(0) \,\big|\, G_{is} = g\big], \tag{1}$$

where g indexes the treatment cohort and t indexes calendar week. In our baseline specification, treated observations are compared to domain–state pairs in not-yet-treated states; in robustness

[25] The data extend through the end of July 2025, which explains the shorter post-implementation window for Georgia.

[26] Other spillover channels, such as users physically relocating or cross-state information diffusion, are assumed to be negligible given the short time horizon and state-specific nature of the policies.

exercises, we instead use never-treated states as the comparison group. Standard errors are clustered at the state and domain-state levels, allowing for dependence within domain–state pairs over time and within states across domains.[27] Under our identification assumptions, the aggregated post-treatment effect captures the average effect of age verification implementation on within-state traffic to adult domains relative to the selected control group.

Second, let $\mathbb{1}\{\text{Strat}_{is} = m\}$ indicate the strategy $m \in \{\text{Adopting, Exit, Non-adopting}\}$ that domain i adopts in state s after implementation.[28] We estimate strategy-specific staggered-adoption effects of the form

$$ATT_m(g,t) \;=\; \mathbb{E}\big[Y_{ist}(1) - Y_{ist}(0) \,\big|\, G_{is} = g,\ \text{Strat}_{is} = m\big]. \tag{2}$$

We then aggregate these group–time treatment effects across cohorts and post-treatment periods to obtain an overall post-implementation effect for each strategy arm, denoted by $\overline{ATT}_m$. Here, $\overline{ATT}_m$ captures the average post-implementation traffic change for domains in strategy arm m relative to the relevant control group, namely not-yet-treated or never-treated states depending on the specification. Consistent with economic intuition and Table II, we expect

$$\overline{ATT}_{\text{Exit}} \ll \overline{ATT}_{\text{Adopting}} < 0, \qquad \overline{ATT}_{\text{NonAdopting}} \geq 0$$

on average.

Finally, to assess pre-trends and dynamic treatment effects, we estimate a cohort-robust event study following Callaway and Sant'Anna (2021). Let $k = t - g$ denote event time, i.e. weeks relative to the state's implementation date. We report dynamic effects obtained by aggregating the estimated group–time treatment effects by event time:

$$\theta_k \;=\; \sum_g w_{g,k}\, ATT(g, g+k), \tag{3}$$

where $w_{g,k}$ are the cohort-specific aggregation weights and $k = -1$ is the omitted reference period. We report Callaway and Sant'Anna (2021) event-study estimators, with simultaneous confidence bands conveying uncertainty across cohorts. We repeat this specification across the adoption strategies laid out above. Moreover, to allay concerns about violations of the parallel trends assumption, we provide event-study plots showing the effects for websites in treated states relative both to never-treated states and to not-yet-treated states, conditional on their status.

5 Results

We can now report our main results. We proceed first with showing the event-study estimates in Figure III, which support the validity of the identifying assumptions. Across all adoption

[27] Standard errors are computed using the multiplier bootstrap, which allows us to construct simultaneous confidence bands for the event-study estimates.

[28] We classify strategies by automated checks (HTTP status, age verification-gate presence) in the first post-implementation week and validate it with manual audits and press coverage.

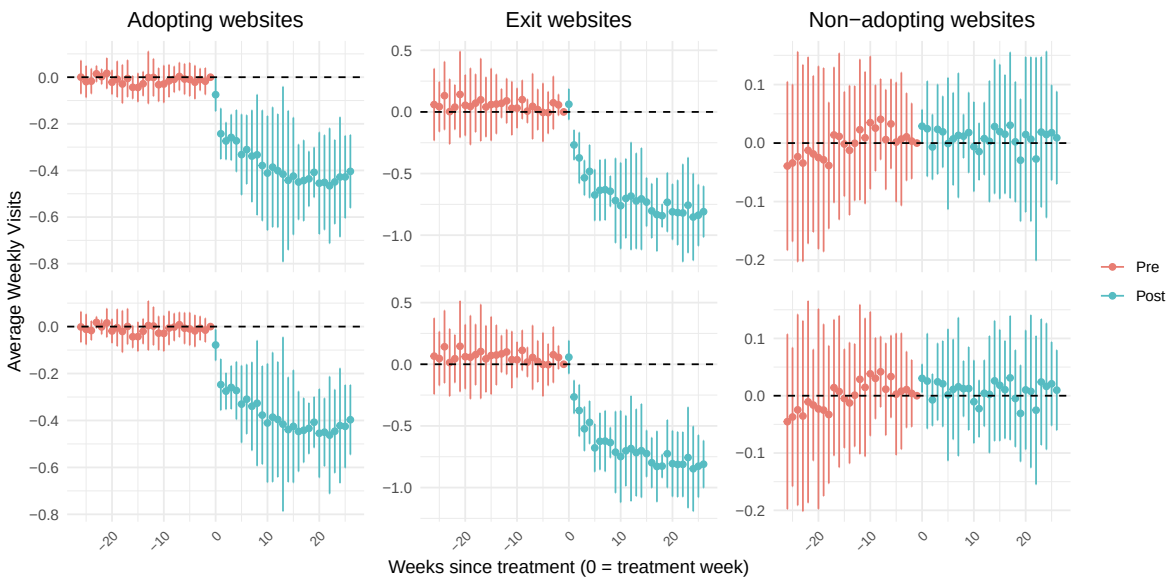

Figure III: Event study graph by adoption status (centered at treatment week). The top panel shows websites in treated state relative to never-treated states, conditional on their status. The bottom panel shows websites in treated state relative to not-yet-treated states.

groups, pre-treatment coefficients are statistically indistinguishable from zero, providing no evidence of anticipatory behavioral changes prior to policy implementation. The timing of effects aligns sharply with each state's enforcement week (denoted as period 0). For domains that were shut down, e.g., those belonging to Aylo such as Pornhub, the decline is immediate and persistent, with no subsequent recovery as the websites never came back visible to the public. For adopting domains, on average, traffic drops discretely at implementation, followed by a modest partial rebound that stabilizes below pre-law levels. Non-adopting domains display almost flat trajectories both before and after the law.

Table III reports the intent-to-treat (ITT) estimates of age verification on logged website visits and bounce rate. Consistently with the event study analysis, there is a substantial and heterogeneous impact of the policies on user traffic across platforms. Overall, visits decrease approximately by 23% relative to the counterfactual of never-treated or not-yet-treated states.[29]

Looking at the effect of the laws conditional on the platform strategies, we observe different patterns. Domains that withdrew access to users in affected states, notably Pornhub, experienced the largest reduction in traffic, with visits declining by approximately 55–56%. This near-elimination of observed traffic is consistent with the exit from the local market. It is important to note that exit platforms do not exhibit a complete (i.e., 100%) decline in measured traffic because their homepages often remain publicly accessible to display informational or advocacy messages opposing the legislation. As a result, these static pages can still attract limited user traffic recorded by the provider, even though access to the main content is effectively blocked. Likewise, this traffic may also represent part of the traffic that is generated via VPNs, where the

[29]We compute the implied effects by transforming the estimated coefficients into percentage changes, e.g., $\exp(\beta) - 1$, where β is obtained from estimating equation (1) and reported in Table III.

Table III: Intent-to-Treat (ITT) estimates of age-verification on log visits and Bounce rate

	log(Visits)		Bounce rate	
	(1)	(2)	(3)	(4)
All	-0.27***	-0.26***	0.02*	0.02*
	(0.03)	(0.03)	(0.01)	(0.01)
Adopting	-0.45***	-0.45***	0.01	0.01
	(0.03)	(0.03)	(0.01)	(0.01)
Non-adopting	-0.01	-0.00	-0.02**	-0.02***
	(0.03)	(0.02)	(0.01)	(0.01)
Exit	-0.81***	-0.81***	0.17***	0.17***
	(0.09)	(0.09)	(0.02)	(0.02)
Counterfactual	Never-treated	Not-yet-treated	Never-treated	Not-yet-treated
Observations	57,960	19,320	57,960	19,320

Note: This table reports intent-to-treat (ITT) estimates of the effect of age verification on log visits and bounce rate. Columns (1) and (2) use log(Visits) as the dependent variable, while columns (3) and (4) use bounce rate. The rows *All*, *Adopting*, *Non-adopting*, and *Exit* report estimates for the full sample, adopting sites, non-adopting sites, and exiting sites, respectively. The counterfactual group is never-treated sites in columns (1) and (3), and not-yet-treated sites in columns (2) and (4). Standard errors are clustered at the state and domain-state levels. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

user has been recorded by the data provider as resident in a treated state, despite pretending to connect from another country or state after the law is implemented.[30]

Platforms that chose to comply with the legislation by introducing verification gates also faced substantial losses, albeit of smaller magnitude. For these domains, traffic fell by approximately 36%. This finding suggests that while adoption preserved formal accessibility, the requirement to disclose identity information introduced significant frictions that deterred a significant share of users from accessing the sites.

In contrast, non-adopting sites continued to operate without verification requirements and exhibited no statistically significant change in traffic, on average. This heterogeneous group includes major platforms such as XNXX and XVideos, as well as a number of fringe ones. This pattern suggests that age verification laws did not trigger systematic diversion of users from adopting or exiting platforms toward non-adopting ones within treated states. The absence of meaningful substitution stands in contrast to previous supply-side moderation episodes (e.g., Madio et al. 2025), where the removal of content by major platforms produced a pronounced reallocation toward fringe or unregulated outlets. This difference is intuitive in the present context: users facing verification requirements or temporary shutdowns can still access the same platforms through alternative channels—such as connecting from non-regulated states or employing virtual private networks (VPNs)—without bearing substantial additional search costs.[31] From a search-cost perspective, the availability of simple circumvention tools implies

[30] As explained above, even if a user changes to use a VPN, the data provider would not record that user to the state in which the user appears to the platforms to reside, but to the state of the first connection in a given day.

[31] It is important to note, however, that this limited additional search cost can come at a price, given the ability

that users do not need to engage in a costly search for substitutes; rather, they can maintain access to their preferred sites by adjusting their access method.

As for the results on the bounce rate, the pattern is consistent with the fact that exit sites only display a shutdown message on their homepage once the law is in force. For these platforms, the bounce rate (which captures whether users navigate beyond the landing page) increases by about 17 percentage points relative to the control group (columns (3) and (4)), indicating that users briefly arrive, see the closure notice, and immediately leave. By contrast, non-adopting sites experience a reduction in the bounce rate of approximately 2 percentage points to their counterfactual. This may indicate that any diverted users who reach these platforms are then more likely to stay and browse, consistent with a more engaged or better "matched" residual user base.

6 Heterogeneity analysis

In this section, we provide more details on the effect of the policies on user traffic depending on different platform characteristics such as the way the platforms collect money before the policies (namely their business model) and size, distinguishing between large and fringe sites.

6.1 Heterogeneity among adopting sites

We first examine heterogeneity in treatment effects across *adopting* platforms, focusing on those that already operated with mandatory user registration or payment-based interactions before age verification laws. The intuition is that such sites, whose business model depends on user authentication and monetized participation, likely faced weaker marginal frictions from age gates relative to "free tube" platforms where users are accustomed to anonymous access.

From a user's perspective, the age-verification adoption costs can be both technological and behavioral. For registration-based or payment-dependent platforms, these costs are largely sunk: users are already accustomed to providing identifying information, and the additional privacy burden imposed by government ID verification is incremental rather than new. Platforms like Stripchat, Chaturbate, OnlyFans, and xHamsterLive rely on tipping, paid subscriptions, or interactive live sessions, all of which require a registered account or payment credentials. Consequently, their user base is composed of individuals who have already traded off some degree of anonymity for participation or monetary exchange.

By contrast, open-access "tube" sites (e.g., xHamster) authentication step introduces a larger marginal deterrent, potentially leading to sharper reductions in post-policy traffic. Hence, among adopting sites, we expect smaller proportional traffic losses for registration-based platforms than for non-transactional ones.

Figure IV shows average weekly visits by a group of adopting sites, distinguishing between registration-based sites. We find that registration-based sites experienced attenuated declines in

of VPN providers to track users.

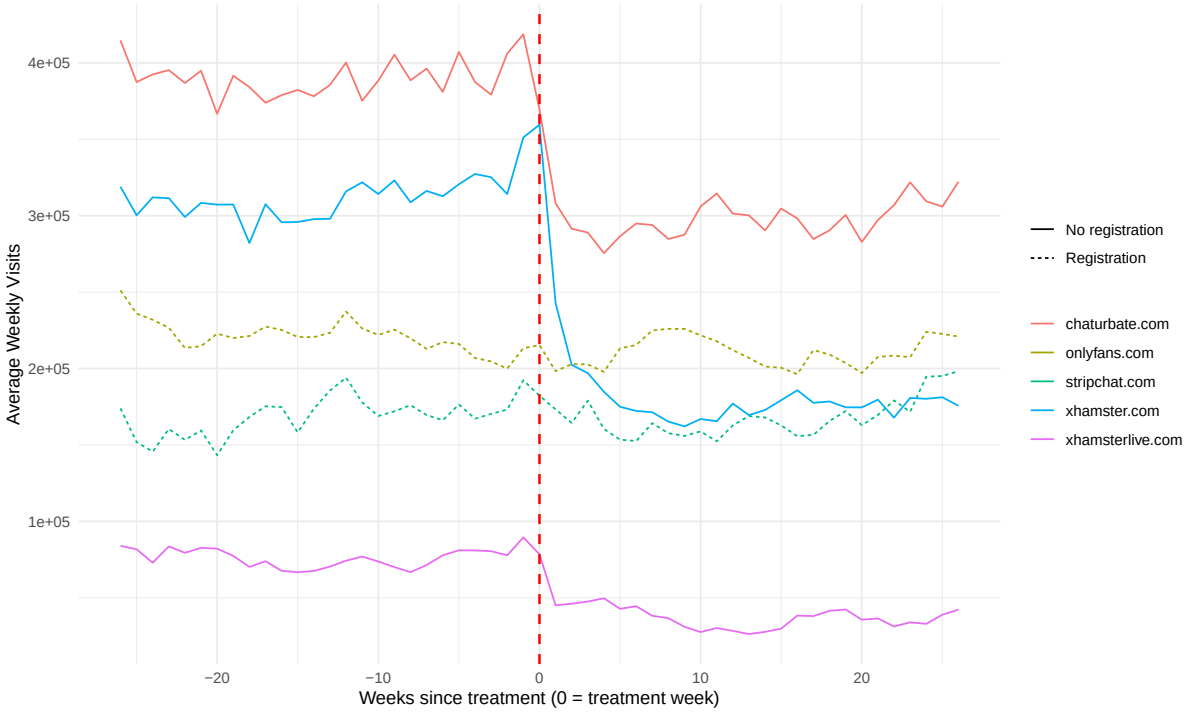

Figure IV: Average weekly visits by adoption status (Centered at Treatment Date).

visits, consistent with lower incremental privacy and hassle costs for users already operating in authenticated environments. In other words, xHamster loses the most among the adopting sites, followed by Chaturbate and xHamsterLive. At the domain level, some registration-based platforms no little visible contraction in the descriptive figures. This is the case of live-cam site Stripchat and membership-based site like OnlyFans did not display traffic contractions after implementation. The result suggests that hassle and privacy frictions, rather than the age verification requirement per se, drive the bulk of user deterrence in this market segment.

Figure V presents an event-study analysis showing that, among adopting platforms, xHamster experienced substantial traffic losses, whereas sites requiring registration for payments do not display any clear decline. The figure also shows that this pattern is robust across alternative control groups: in the top panel, we use websites in never-treated states, conditional on their adoption status, while in the bottom panel, we compare websites in treated states to those in not-yet-treated states.

Table IV quantifies the effect of the policy among adopting platforms, distinguishing between those that do and do not require user registration for payment purposes. xHamster, the only non-registration site, experiences a decline in visits of about 48% relative to its control group (regardless of the chosen counterfactual), whereas registration-based platforms show a much smaller average reduction of roughly 14—15%, across specifications.

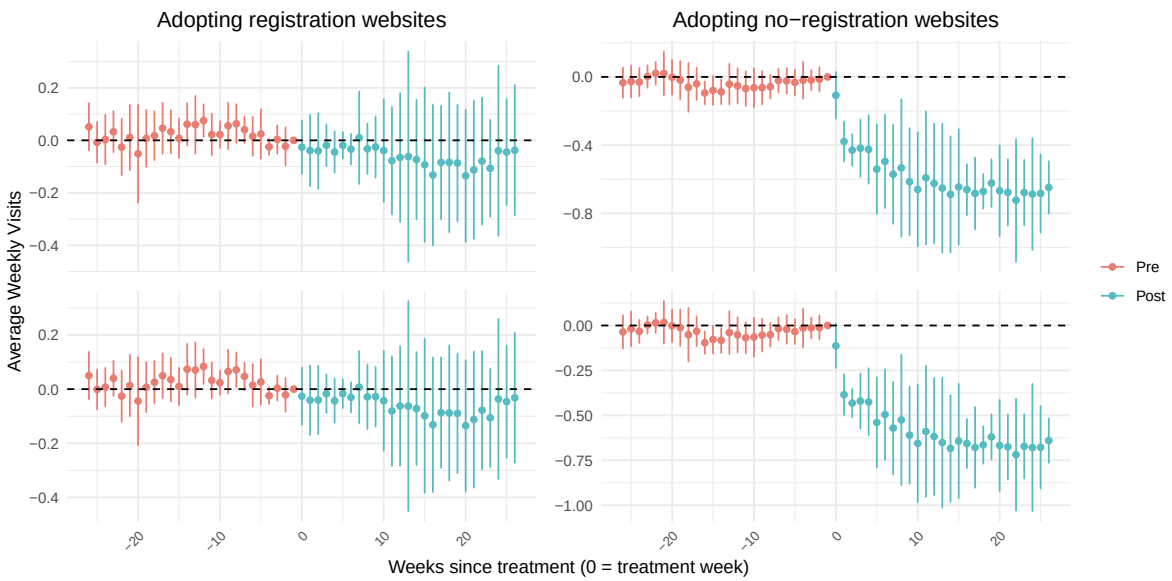

Figure V: Event study graph for adopting platforms (centered at treatment week) comparing xHamster (no registration) with registration/membership-based platforms (i.e., Chaturbate, OnlyFans, xHamsterLive, Stripchat). The top panel shows websites in treated state relative to never-treated states, conditional on their status. The bottom panel shows websites in treated state relative to not-yet- treated states.

Table IV: Intent-to-Treat (ITT) estimates of age-verification on log visits and Bounce rate

	log(Visits)		Bounce rate	
	(1)	(2)	(3)	(4)
Adopting	-0.45***	-0.45***	0.01	0.01
	(0.03)	(0.03)	(0.01)	(0.01)
Registration	-0.15**	-0.16***	0.01	0.01
	(0.05)	(0.05)	(0.01)	(0.01)
No-registration	-0.65***	-0.65***	0.01	0.01
	(0.05)	(0.05)	(0.02)	(0.02)
Counterfactual	Never-treated	Not-yet-treated	Never-treated	Not-yet-treated
Observations	12,075	4,025	12,075	4,025

Note: This table reports intent-to-treat (ITT) estimates of the effect of age verification on log visits and bounce rate for adopting sites, separately by whether sites require registration. Columns (1) and (2) use log(Visits) as the dependent variable, while columns (3) and (4) use bounce rate. The rows *Adopting*, *Registration*, and *No-registration* report estimates for all adopting sites, adopting sites with registration requirements, and adopting sites without registration requirements, respectively. The counterfactual group is never-treated sites in columns (1) and (3), and not-yet-treated sites in columns (2) and (4). Standard errors are reported in parentheses. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

6.2 Mainstream vs. fringe sites

Large and small adult platforms differ systematically in both market weight and the amount of content that they carry out. In the baseline specification, each domain contributes equally

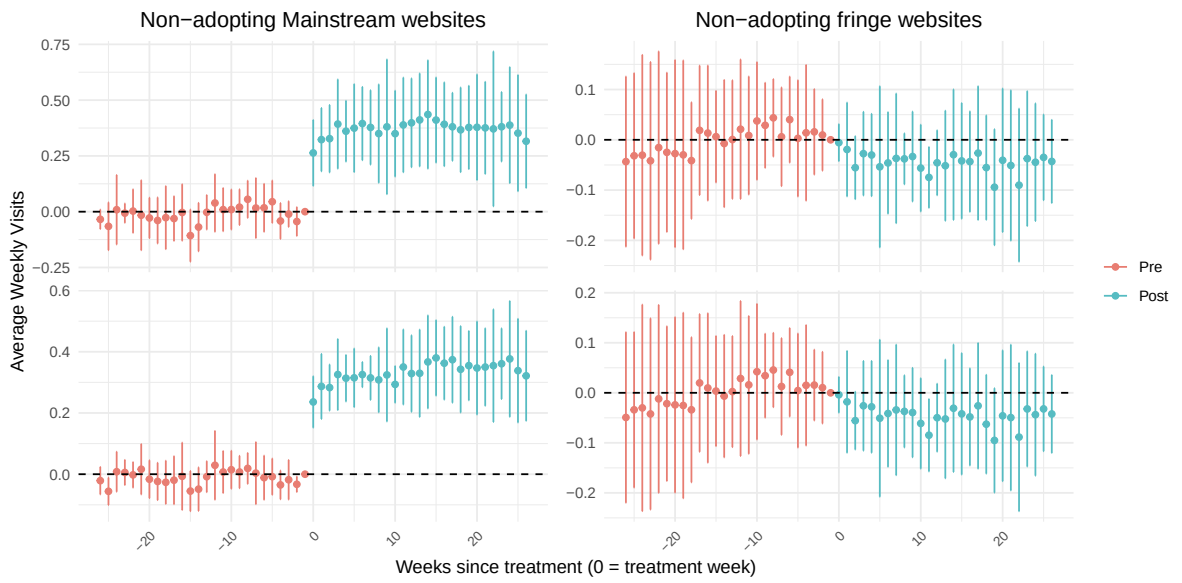

Figure VI: Event study graph for non-adopting platforms (centered at treatment week) comparing mainstream (e.g., XNXX and XVideos) with fringe. The top panel shows websites in treated state relative to never-treated states, conditional on their status. The bottom panel shows websites in treated state relative to not-yet- treated states.

to the estimation of the ITT. This approach risks over-weighting small, low-traffic websites, whose higher volatility may mechanically drive aggregate effects. To mitigate this concern, we decompose non-adopting sites by size, distinguishing between the two dominant "tube" platforms—*XNXX* and *XVideos*—and all other non-adopting domains and re-estimate our main analysis.

Figure VI presents an event study analysis showing that among non-adopting platforms, there is substantial heterogeneity between mainstream players, which gained substantial amounts of traffic, and fringe sites whose traffic remained unaffected. The figure shows that this result is robust to both when the control group is made of websites in never-treated states, conditional on their status (top panel) and websites in treated state relative to not-yet- treated states (bottom panel). Moreover, the pre-treatment period shows no significant difference between the treated websites and the control group, reassuring about the validity of the parallel trend assumption.

Table V quantifies the effect of the policy among non-adopting platforms. It shows that mainstream sites gain about 35–40% traffic relative to the control group, depending on the counterfactual specification, whereas fringe platforms do not exhibit any statistically significant change. In addition, mainstream sites improve their performance along the intensive margin: their bounce rate declines by 6 percentage poins, indicating that users who are diverted to these domains are more likely to continue browsing beyond the landing page.

6.3 Traffic source

Since traffic was partially reallocated from adopting and geo-blocked sites to non-adopting ones, we examine how age verification laws affect the composition of traffic sources across the different

Table V: Intent-to-Treat (ITT) estimates of age-verification on log visits and Bounce rate

	log(Visits)		Bounce rate	
	(1)	(2)	(3)	(4)
Non-adopting	-0.01	-0.00	-0.02**	-0.02***
	(0.03)	(0.02)	(0.01)	(0.01)
Mainstream	0.34***	0.30***	-0.06***	-0.06***
	(0.09)	(0.05)	(0.01)	(0.01)
Fringe	-0.06*	-0.05	-0.01	-0.01
	(0.03)	(0.03)	(0.01)	(0.01)
Counterfactual	Never-treated	Not-yet-treated	Never-treated	Not-yet-treated
Observations	33,810	11,270	33,810	11,270

Note: This table reports intent-to-treat (ITT) estimates of the effect of age verification on log visits and bounce rate for non-adopting sites, separately by mainstream and fringe sites. Columns (1) and (2) use log(Visits) as the dependent variable, while columns (3) and (4) use bounce rate. The rows *Non-adopting*, *Mainstream*, and *Fringe* report estimates for all non-adopting sites, mainstream non-adopting sites, and fringe non-adopting sites, respectively. The counterfactual group is never-treated sites in columns (1) and (3), and not-yet-treated sites in columns (2) and (4). Standard errors are reported in parentheses. *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

groups of sites. The provider distinguishes between direct visits (URL typing or bookmarks), referrals from other websites, search traffic (organic and paid), social media, display advertising, and email campaigns.

Figure VII reports the composition of traffic sources by adoption group around the implementation of age-verification laws. Several patterns emerge. First, direct traffic is the dominant source across adoption groups and states throughout the sample, whereas paid search, display advertising, email, and social media account for only small shares. Second, for adopting and non-adopting domains, the source composition changes little around the implementation date. This suggests that the law primarily affects aggregate traffic volumes, rather than materially altering the channels through which users reach these sites. Third, the most noticeable compositional shift occurs among exit domains. Following implementation, the share of referral traffic declines, while the share of organic search rises; direct traffic remains quantitatively important, although it appears to fall somewhat in several states. Because Figure VII is expressed in shares rather than levels, these patterns should be interpreted with caution. In particular, an increase in the share of search traffic does not imply an increase in search traffic in absolute terms, but only that search accounts for a larger fraction of the reduced post-treatment traffic. Taken together, the figure provides limited evidence that traffic reallocated toward non-adopting domains was driven primarily by discovery through search or advertising (unlike in Madio et al. (2025)). Instead, the continued predominance of direct traffic is more consistent with users accessing these sites through pre-existing navigation paths, such as typed URLs or bookmarks.

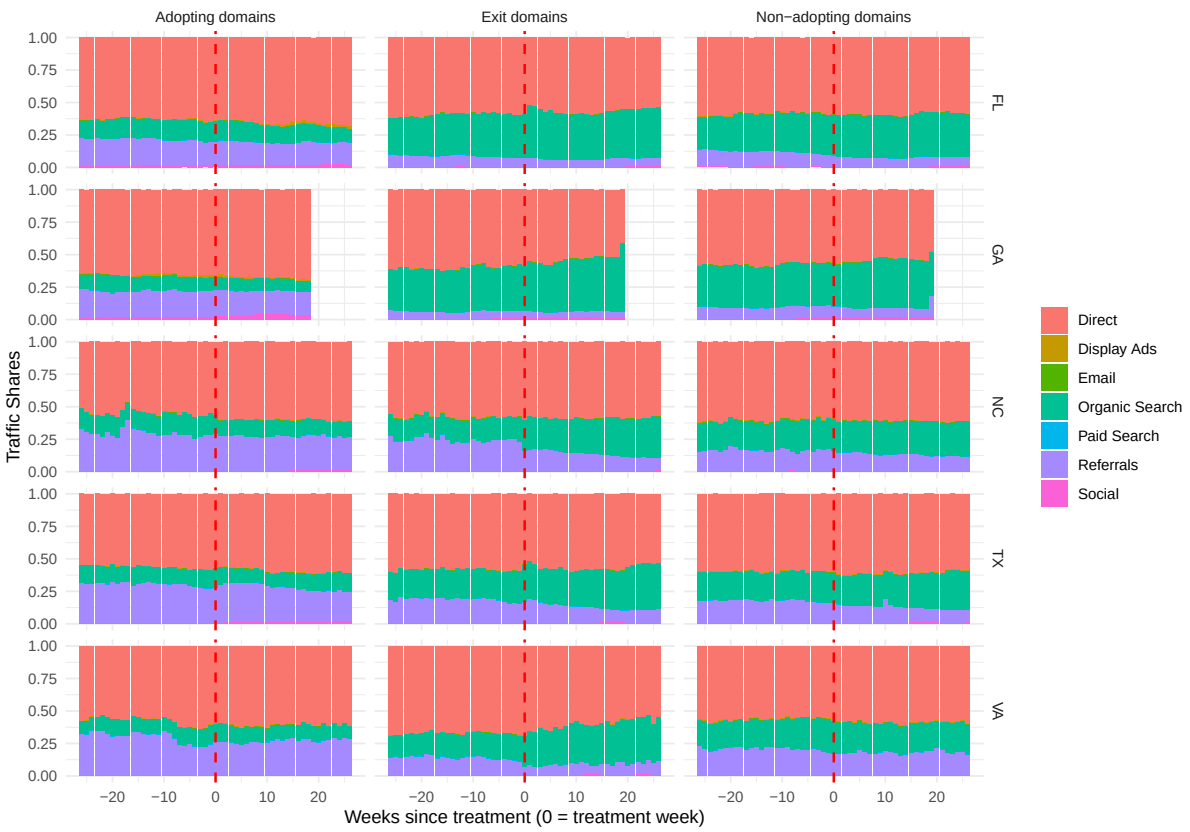

Figure VII: Sources of traffic share to adult domains, aggregated by adoption group.

7 Conclusions

In this work, we study the effects of age verification laws in the United States between 2023 and 2025 on traffic to adult sites, which are among the most visited sites worldwide. We leverage within-state adoption heterogeneity across platforms and across-state time variation in the adoption of the laws. Platforms respond along three margins—introducing ID-based (or similar) verification, blocking access from adopting states, or remaining non-adopting—and these choices map into markedly different demand responses.

Our estimates indicate that age verification laws reduced observed in-sample traffic to covered adult domains, but the welfare and child-protection implications remain uncertain because the data do not identify minors and do not measure substitution outside the observed platform set and via VPNs. The new requirements generate sizable effective reductions in observed consumption among platforms that geo-blocked access or became adopting. This traffic was diverted mostly towards non-adopting major platforms, with an increase by up to 35-40%.

The heterogeneity analysis highlights that these effects are not uniform across business models and platform sizes and may depend on users' heterogeneous hassle and privacy costs. Among adopting sites, open-access, advertising-funded tube platforms (such as xHamster) experience the largest proportional losses, while registration- and payment-based platforms—live-cam and membership sites where users already maintain accounts and store payment details—show little or no measurable decline in traffic following implementation. Likewise, within the non-adopting

segment, any traffic gains are concentrated among the largest mainstream platforms, whereas smaller domains do not experience meaningful increases.

These patterns are consistent with a mechanism in which age verification mainly deters users for whom the required disclosure of identity or other sensitive personal information—whether actual or merely perceived—represents a new and salient privacy and hassle cost, while leaving relatively unaffected those who had already incurred similar frictions prior to the law.

Taken together, our findings provide several implications for the design and evaluation of child-protection policies in digital markets. First, they suggest that age verification laws can materially reduce access to adult content in treated jurisdictions, even after accounting for partial circumvention via VPNs or cross-state access. However, our analysis cannot link traffic to age bands and, therefore, we cannot document whether the policy was effective in the targeted group of minors.[32] Second, the no diversion toward fringe or poorly supervised outlets contrasts with earlier episodes of supply-side content moderation and indicates that, in this context, stricter access controls need not mechanically push users into more harmful corners of the web. This is especially true because the biggest traffic receivers according to our analysis are non-adopting mainstream players. These platforms are likely to be subject to relatively stricter regulatory provisions and monitoring more broadly due to their size, suggesting that users diverting to these outlets remain within relatively well-supervised spaces, rather than migrating to ungoverned fringe platforms. Yet, the heterogeneous nature of the adoption of measures to meet the age verification regulation requirements has led to some traffic reallocation, which should be carefully tracked by the policy maker. Third, the strong role of pre-existing business models and authentication requirements implies that the incidence of age verification laws falls unevenly across adopting platforms: unauthenticated-access, ad-funded sites bear most of the effective demand reduction, whereas registration- and payment-based platforms are largely pre-compliant and therefore insulated. Fourth, our analysis shows that age verification laws achieve limited demand reduction when large platforms remain non-adopting. Since these platforms concentrate a large share of traffic, such gaps in adoption can considerably undermine aggregate effectiveness.

Finally, we emphasize that our analysis primarily captures within-industry reallocation across adult platforms. In principle, substitution in the broader attention economy can extend across services and content channels (Aridor, 2025), including toward social networks and messaging apps where pornographic material may remain accessible.

[32] Related survey evidence from the UK context is broadly encouraging in this respect. Research by Childnet found that the increase in VPN use following the introduction of age assurance measures was not attributable to children, suggesting that circumvention behavior was concentrated among adult users rather than the targeted group of minors. See: https://www.childnet.com/blog/new-research-from-childnet-shows-that-the-surge-in-vpn-use-following-the-introduction-of-age-verification-in-the-summer-is-not-attributable-to-children/.

References

Agarwal, S., Ananthakrishnan, U.M., Tucker, C.E., 2022. Deplatforming and the control of misinformation: Evidence from Parler. SSRN Working Paper 4232871 .

Allcott, H., Braghieri, L., Eichmeyer, S., Gentzkow, M., 2020. The welfare effects of social media. American Economic Review 110, 629–676.

Allcott, H., Gentzkow, M., Song, L., 2022. Digital addiction. American Economic Review 112, 2424–2463.

Aridor, G., 2025. Measuring substitution patterns in the attention economy: An experimental approach. RAND Journal of Economics 56, 302–324.

Bar-Isaac, H., Deb, R., Mitchell, M., 2025. Content moderation for sale: Pricing attention through steering and certification. arXiv preprint arXiv:2506.12604 .

Beknazar-Yuzbashev, G., Capozza, F., Rutherford, K., Stalinski, M., 2026. Online toxicity, emotions, and mental health. Mimeo.

Beknazar-Yuzbashev, G., Jiménez Durán, R., McCrosky, J., Stalinski, M., 2022. Toxic content and user engagement on social media: Evidence from a field experiment. SSRN Working Paper 4307346 .

Beknazar-Yuzbashev, G., Jiménez-Durán, R., Stalinski, M., 2024. A model of harmful yet engaging content on social media. AEA Papers and Proceedings 114, 678–683.

Brynjolfsson, E., Collis, A., Eggers, F., 2019. Using massive online choice experiments to measure changes in well-being. Proceedings of the National Academy of Sciences of the United States of America 116, 7250–7255. doi:10.1073/pnas.1815663116.

Callaway, B., Sant'Anna, P.H., 2021. Difference-in-differences with multiple time periods. Journal of Econometrics 225, 200–230.

Chan, J., Mojumder, P., Ghose, A., 2019. The digital Sin City: An empirical study of Craigslist's impact on prostitution trends. Information Systems Research 30, 219–238.

Johnson, G.A., Lin, T., Zhong, L., Cooper, J., 2026. COPPAcalypse? the Youtube settlement's impact on kids content. Management Science Forthcoming.

Kircher, T., Foerderer, J., 2024. Ban targeted advertising? An empirical investigation of the consequences for app development. Management Science 70.

Kominers, S.D., Shapiro, J.M., 2024. Content moderation with opaque policies. NBER Working Paper 32156 .

Kumarswamy, N., Singhal, M., Nilizadeh, S., 2023. Impact of stricter content moderation on Parler's users' discourse. arXiv preprint arXiv:2310.08844 .

Lang, D., Listyg, B., Ross, B.V., Musquera, A.V., Sanderson, Z., 2025. Age verification and public adaptation: A pre-registered synthetic control multiverse. Journal of Law & Empirical Analysis , 2755323X251415424.

Liu, Y., Yildirim, P., Zhang, Z.J., 2022. Implications of revenue models and technology for content moderation strategies. Marketing Science 41, 831–847.

Madio, L., Mitchell, M., Quinn, M., Reggiani, C., 2025. Asymmetric content moderation in search markets: The case of adult websites. CESifo Working Paper 11842.

Madio, L., Quinn, M., 2025. Content moderation and advertising in social media platforms. Journal of Economics & Management Strategy 34, 342–369.

Manolios, S., Sala, A., Sundorph, E., Chaudron, S., Gomez, E., Beullens, K., Bozzola, E., Cataldo, I., Hale, L., Kent, M., Montag, C., Nivins, S., O'Reilly, M., Rubaek, L., Schiotz Thorud, H.M., 2025. Minors' health and social media: an interdisciplinary scientific perspective. JRC Policy Report 141090.

Rendo, I., 2024. Content moderation in presence of fringe platforms. Mimeo .

Spencer, N., 2025. Age verification laws, firm responses, and porn consumption. SSRN Working Paper 5165280 .

Zeng, H.S., Danaher, B., Smith, M.D., 2022. Internet governance through site shutdowns: The impact of shutting down two major commercial sex advertising sites. Management Science 68.