Can Online Activity Be Regulated?

Evidence from Adult Websites

Matthew Brown Emily Davis Devin G. Pope*

June 3, 2026

Abstract

The consequences of online regulations depend on the extent to which users can circumvent restrictions or substitute toward noncompliant platforms. Since 2023, 25 U.S. states have implemented age verification laws that caused prominent adult websites (including Pornhub) to restrict local access for all users. We study how these restrictions affected browsing activity using individual-level panel data. Access restrictions reduced overall time spent on adult sites by roughly 10%. Specifically, for every 100 hours spent on top adult sites before restrictions, about 50 hours remained accessible at noncompliant sites that never restricted access, 30 hours persisted through VPN-based circumvention, 10 hours were substituted from compliant sites to noncompliant sites, and 10 hours were no longer spent on adult sites.

JEL Classifications: L86, K42, L51, I18
Keywords: online regulation, age verification, platform substitution, pornography, adult websites

---

* Matthew Brown, Kenneth C. Griffin Department of Economics, mattbrownecon@uchicago.edu; Emily Davis, University of Notre Dame, Department of Economics, edavis25@nd.edu; Devin G. Pope, NBER and Booth School of Business, University of Chicago, devin.pope@ChicagoBooth.edu

We thank seminar participants at UC Los Angeles and UC Santa Barbara for their comments and suggestions. Kasey Buckles, Matthew Gentzkow, and Aaron Leonard provided helpful feedback. We thank Hannah Lybbert for outstanding research assistance. We gratefully acknowledge funding from the Kenneth C. Griffin Applied Economics Incubator. Davis is grateful for financial support from the Institute for Humane Studies Expense Support Grant and the American Institute for Boys and Men's Graduate Student Fellowship.

1 Introduction

Digital platforms have become an increasingly important target of regulation. Governments around the world have sought to restrict access to online content, often motivated by concerns about addiction, misinformation, or harms to minors. These regulatory efforts range from broad internet restrictions, such as China's censorship regime, to targeted platform bans, including TikTok in India, Twitter/X in Brazil, and social media restrictions for minors in Australia.

Regulating online activity presents unique challenges. Unlike in many offline markets, digital consumers can often circumvent restrictions through technologies such as virtual private networks (VPNs). It is also sometimes challenging to enforce regulations on digital content providers, in which case users may substitute toward noncompliant platforms. As a result, the consequences of online regulation are often uncertain.

In this paper, we study how regulations affect online activity in the context of internet pornography. Pornography is one of the most widely consumed forms of digital content, generating billions of visits each month globally.[1] Since 2023, 25 U.S. states have passed age verification laws regulating websites that host pornographic content. The laws require such websites to verify that users are at least 18 years old using government-issued identification or comparable verification methods.[2] These policies represent one of the most aggressive attempts to regulate online content in the United States and provide a natural setting for studying responses to regulation.

Adult websites have responded to these laws in different ways. Pornhub, the world's largest adult website, blocks access for users in states that enact age verification laws. The block applies to all users, including adults as well as minors. Other sites chose to comply by restricting access to content until a user had completed an age verification procedure. By contrast, some sites continued operating without implementing age verification requirements, relying

[1] In April of 2026, the top 5 adult websites received ∼8.9B visits, more than any individual website (Facebook, Instagram, ChatGPT, Reddit, etc.) with the exception of Google.com (93.8B visits) and YouTube.com (51.1B visits) (Semrush, 2026). In the U.S., 46% of men and 16% of women between the ages of 18 and 39 self-report watching pornography on a weekly basis (Regnerus et al., 2016).

[2] We discuss these methods in detail in the background section. But, crucially, the laws go beyond "check the box if you are 18 or older," and sites that do not comply with these requirements can face large fines.

on the practical difficulty of enforcement. Users who had been visiting these noncompliant sites before the law was passed were able to continue doing so without interruptions.

Alongside this noncompliance by websites, there is scope for user responses that diminish the impact of the regulation. In many cases, including that of Pornhub, users could circumvent blocks with VPNs that masked their location. Users also could substitute from compliant sites to noncompliant sites and continue to access pornography without using a VPN. The possibility remains that the laws could cause users to stop visiting adult websites. But because of site noncompliance and user circumvention and substitution, whether age verification laws reduce overall pornography consumption remains an open empirical question.

We answer this question using high-frequency, individual-level online activity data from a rotating panel of approximately 550,000 U.S. desktop and mobile internet users provided by Comscore. Using a stacked difference-in-differences design that exploits the staggered nature of the state law rollout, we show that access restrictions significantly alter online behavior. Specifically, for every 100 hours spent at the top 25 adult websites before the shutdowns, roughly 50 hours were spent on websites that did not comply with the law. Of the remaining 50 hours, 30 hours persist through VPN-based circumvention, 10 hours are substituted from compliant websites to noncompliant websites, and 10 hours are no longer spent at adult websites. We do not observe meaningful substitution to "fringe" adult websites outside of the top 25. These effects remain stable for several weeks after implementation, suggesting that the observed responses reflect more than short-run disruption following website closure.

Our individual-level data allow us to document heterogeneous responses across household types and across device types (desktop and mobile). We find evidence of circumvention and substitution for all subgroups, but in most cases we still observe significant reductions in overall pornography consumption. Importantly, our panel consists entirely of adults. Thus, our estimates capture the impact of age verification laws on adult users rather than minors. Although comparable data for minors are unavailable, we argue that the behavioral responses of young adults (we provide estimates for 18-24 year-olds) likely provide useful insight into how these laws affect teenagers.

Our paper contributes to a growing literature studying the effects of online regulation across a variety of domains. Recent work has examined the impact of access to uncensored internet

(Chen and Yang, 2019), bans on targeted advertising (Kircher and Foerderer, 2024), crackdowns on pirated content (Danaher et al., 2020), and data privacy regulation (Goldberg et al., 2024). We combine rich individual-level panel data with an unusually clean natural experiment involving the staggered rollout of regulations across jurisdictions. This combination allows us to quantify the importance of three responses (noncompliance, circumvention, and substitution) that undermine restrictions on internet activity. While our quantitative estimates are of course specific to adult websites, these channels are relevant in many settings of interest: for example, Australian teens self-report using VPNs to evade social media restrictions (Bursztyn et al., 2026).

More closely related to our work is research studying regulation of adult websites specifically. Madio et al. (2026) examine a natural experiment in which Pornhub removed all unverified videos from its platform following a widely discussed New York Times opinion article concerning child pornography. Using aggregate website traffic data, they find that this purge induced substitution from Pornhub to competing adult sites. Three papers (one published and two concurrent working papers) study the same age verification laws we analyze. Using Google Trends data, Spencer (2025) and Lang et al. (2026) use a difference-in-differences approach and document declines in search interest for compliant websites (e.g. Google searches for "Pornhub") alongside increases in interest for noncompliant competitors. Using aggregate website traffic data from five treated states and ten control states, Agarwal et al. (2026) similarly find evidence of substitution from compliant to noncompliant sites and an overall reduction in visits to 24 popular online adult websites.

These three papers provide compelling evidence that age verification laws are impacting website traffic. At the same time, our paper extends this literature by overcoming important data limitations present in existing work. For example, because Google Trends data are normalized, they are better suited for identifying directional changes in search intensity than for quantifying substitution patterns across platforms. In addition, Google Trends relies on IP-based geographic assignment and therefore cannot effectively capture activity routed through VPNs.

Our individual-level data allow us to extend the literature in several ways. First, we provide direct evidence of circumvention, since visits to adult websites are observable independent of whether a VPN is being used. We also quantify the relative importance of circumvention

compared to other possible user responses. Second, we measure the exact amount of time (in minutes) that individuals spend on adult websites. Analyses that use aggregate website traffic or Google search data can only study changes in searches or in visits. Third, we track substitution from compliant adult websites to the full set of alternative adult websites (200,000+) rather than only a subset of the most visited platforms. Finally, our individual-level data allow us to examine heterogeneity in behavioral responses across users and households.

Our analysis comes with caveats. Most importantly, our quantitative estimates are local to the Comscore panel, which is a selected sample of internet users. Within the Comscore panel, we do not observe website visits in private browsing sessions (such as Google Chrome's "Incognito Mode"), even though many adult website visits may take place in private browsing mode. We emphasize that private browsing does not allow users to circumvent state-level restrictions. We also do not observe users increasing their usage of other privacy-preserving browsing tools after the law, so we think of this limitation as affecting the external validity of our estimates rather than as a possible source of bias.

The rest of the paper proceeds as follows. Sections 2-5 present the institutional background, data, empirical strategy, and results respectively. Section 6 concludes, including a discussion of policy implications concerning the regulation of internet pornography and of online activity more broadly.

2 Institutional Background

2.1 The Market for Internet Pornography

The online pornography market is dominated by so-called "tube sites," which function similarly to YouTube by hosting large libraries of free, user-accessible video content. These sites make money through advertising, premium subscriptions, and affiliate programs. While other segments of the industry exist, including cam sites and subscription-based creator platforms such as OnlyFans, tube sites remain the primary source of online adult content traffic.

Although there are many adult websites online, the market itself is relatively concentrated. A small number of dominant platforms account for a large share of total traffic, while the remaining market share is distributed across numerous smaller competitors. The three focal websites for our analysis (Pornhub, XVideos, XNXX) are all tube sites and together account for 50% of time spent on adult websites.

This market structure is central to our analysis. Because Pornhub is one of the dominant firms in the market, its decision to block access in states adopting age verification laws represents a substantial shock to online pornography access. However, the availability of substitute platforms makes the overall effect on pornography consumption theoretically ambiguous.

2.2 Age Verification Laws

Although the welfare effects of pornography consumption remain debated,[3] concerns about youth exposure generate unusually broad consensus.[4] In the United States, more than half of children report having viewed pornography before age 13, and a substantial share of adolescents intentionally consume pornography on a regular basis (Robb and Mann, 2023).[5] To address these concerns, Louisiana passed House Bill 142 (HB 142) on January 1, 2023, which was the first age verification law for pornography websites in the United States. The law established the template for similar legislation later adopted in 24 additional states, most of which closely mirror Louisiana's original language. In Table A.1, we list these states, the date when the law went into effect for each state, and also our best estimate as to when Pornhub complied with the law. In Appendix A.1, we detail the provisions of these

[3]Meta-analyses of experimental, cross-sectional, and longitudinal studies show evidence for an association between consumption of pornography and acts of sexual aggression (Wright et al., 2016) and violence (Mestre-Bach et al., 2024). In economics, pornography consumption has been plausibly linked to increases in rape and other sex crimes (Gibbons and Rossi, 2021; Bhuller et al., 2013). Research also suggests an association between pornography consumption and marital instability and relationship dissatisfaction (Doran and Price, 2014). At the same time, other research has pushed back on the link between pornography and sexual violence (Kendall, 2006) and some research has suggested potential benefits of pornography in certain contexts (Kohut et al., 2021). Public opinion remains somewhat divided: roughly one third of Americans consider pornography morally acceptable (Newport, 2025; Bowman and Cox, 2022).

[4]A 2023 poll found that 83% of registered voters support a national age verification mandate for adult content sites (McCarthy, 2023) and a 2022 poll found that 86% of parents believe it is too easy for children to access pornography online (Toscano, 2024).

[5]According to this survey from Robb and Mann (2023), 40% of boys and 24% of girls between the ages of 13 and 17 reported they intentionally viewed pornography in the last week.

laws, how pornography websites have responded, and the legal framework surrounding their enforcement. In brief, the laws require pornography websites (not social media or other platforms that can have some pornographic material) to verify users' ages either directly or through third-party verification providers, else face fines and potentially other punishments. Common forms of verification include uploading a government-issued ID or providing credit card information.

Pornography websites have responded to these laws in different ways. Pornhub, the world's most visited pornography website, chose to block access entirely in most states that enacted age verification laws.[6] In contrast, the second and third most visited tube sites (XVideos and XNXX) chose not to comply with the laws and have continued operating without implementing age verification systems. Differences in decisions may have been driven by the legal uncertainty of whether states can sue foreign companies[7] and Supreme Court precedent for striking down laws that could restrict adult access to pornography.[8] We conjecture that Pornhub's visibility as the industry leader and prior scrutiny for hosting child pornography (Kristof, 2020) drove its decision to block access, both to lead the industry's legal battle against age verification laws and to avoid another high-profile public relations crisis. Since XNXX and XVideos are less well-known, they have relied on the legal difficulties in enforcing the laws to continue operating.[9]

Pornography consumers can circumvent these restrictions with limited fear of punishment. Specifically, because Pornhub uses IP addresses to determine whether to block access, consumers can use a Virtual Private Network (VPN) that masks their location. For example, the consumer can connect to a server located in a different state (one without age verification requirements) or a different country. When connected to this other server, the consumer can enter Pornhub even if they are physically located in a state with age verification laws.

[6]The only states with age verification laws in which Pornhub still remains accessible are Louisiana and Ohio. In Louisiana, Pornhub relies on a third-party age verification provider, while Pornhub has argued that Ohio's law was written in a manner that does not require the platform to comply.

[7]Aylo, the parent company of Pornhub, is headquartered in Canada, while WGCZ Holding, the parent company for XVideos and XNXX, is headquartered in the Czech Republic.

[8]For example, *Reno v. ACLU* (U.S. Supreme Court, 1997) ruled that the anti-indecency provisions of the 1996 Communications Decency Act were unconstitutional. These provisions were the first attempt to protect minors from obscenity on the internet.

[9]More recently, states such as Florida have sued XVideos and XNXX for noncompliance. However, during our sample period, these sites remained largely free from legal scrutiny.

3 Data

3.1 Comscore Data

We use internet browsing data from a large panel of U.S. households observed across both mobile and desktop devices, provided by Comscore, covering the period from January 1, 2022 through December 31, 2024. Comscore is a U.S.-based media measurement and analytics company that collects detailed information on online browsing behavior and user demographics from a panel of internet users.[10]

For each machine (i.e., an individual mobile or desktop device) registered in the panel, we observe the date and timestamp for each website visited, the number of pages viewed on each website, and the duration of the visit. We also observe a limited set of demographic characteristics associated with each machine, including the age of the household head, household size, household income, presence of children in the household, and geographic location.

One advantage of the Comscore data is that a machine's geographic location is stable over time.[11] In contrast, many alternative sources of website traffic data (including Google Trends) rely on IP-based geographic assignment, which can be altered through VPN usage. In our data, activity generated while a panelist is connected to a VPN remains observable because the machine's location is determined by the stable demographic information associated with the panelist rather than by the contemporaneous IP address. Although we do not observe a direct indicator for whether a VPN is being used during a particular browsing session, Comscore confirmed that VPN usage does not prevent the underlying browsing activity from being recorded in the panel.

[10] Panelists are recruited online through banner advertisements and related methods and are typically compensated through incentives such as free software downloads. Participants install metering software on their mobile devices or computers, which allows Comscore to track browsing activity on those devices.

[11] The geographic location corresponds to Comscore Markets, which are similar to Nielsen DMAs. These markets are more granular than state-level, and several markets cross state lines. Our treatment variation is at the state-level, so for machines in markets that straddle state boundaries, we assign the machines to states based on the share of the population in the market that comes from each surrounding state. Specifically, if the majority ($> 50\%$) of the market's population comes from one state, we assign the market to that majority state. We describe this process in more detail in the Appendix A.2.

One limitation of the data is that we do not observe websites visited in private browsers like Google's Incognito Mode. Private browsing is not itself a method for circumventing age verification laws, because it does not alter a user's IP address. Accordingly, we do not expect the use of incognito browsing to respond systematically to treatment status. However, incognito browsing will cause us to understate total pornography consumption both before and after the implementation of age verification laws. Similarly, to the extent that Comscore panelists are aware that they are being tracked, they may change their pornography consumption. Because panelists in both treated and control states are tracked, this concern affects only external validity.

Our data allow us to construct a measure of the weekly number of minutes each machine spends on pornography websites, which serves as the primary outcome variable in our analysis. Using the timestamp and duration information for each website visit, we aggregate browsing activity to the machine-week-site level. Prior to aggregation, we winsorize all observations to the 95th percentile of website visit length. We describe these cleaning methods in more detail in Appendix A.3.

Much of our analysis focuses on the three largest tube sites in our data: Pornhub, XVideos, and XNXX. We also observe visits to more than 200,000 other adult websites, which we aggregate into a single category labeled "Other XXX" sites. We identify these websites using Comscore's "XXX Adult" website classification.[12]

3.2 Summary Tables

Table 1 reports summary statistics on how pornography consumption varies with respect to demographic characteristics within our sample. The average machine in our sample spends approximately 4.5 minutes per week on adult websites. There are at least two reasons that this figure may understate the total amount of time the average American spends consuming

[12]We independently validated the accuracy of Comscore's "XXX Adult" classification using an external blocklist of pornographic domains obtained from https://github.com/Bon-Appetit/porn-domains GitHub repository. This exercise confirmed that the classification captures both major and fringe pornography websites. The classification is also relatively restrictive: it excludes platforms such as OnlyFans, which is technically categorized as a social media platform despite hosting substantial amounts of pornographic content. Similarly, websites such as Reddit are excluded even though they contain prominent adult-content communities, because they are not primarily pornography-focused websites.

online pornography. First, we do not observe browsing activity conducted in incognito or private browser modes. Second, our sample is disproportionately composed of desktop devices (87%) whereas pornography consumption occurs primarily on mobile devices.[13]

Table 1 also reports consumption for the three largest tube sites in our sample. Of the 4.5 weekly minutes spent on adult websites, approximately 1 minute is spent on Pornhub, 0.67 minutes on XVideos, and 0.56 minutes on XNXX. Together, these websites account for roughly half of observed time spent on adult websites in our data.

There is substantial heterogeneity in pornography consumption across users and devices. Men consume much more pornography than women (12.6 versus 3.9 minutes per week), and mobile devices exhibit much higher usage than desktop devices (18.8 versus 2.3 minutes per week). These patterns are consistent with existing evidence on self-reported pornography consumption (Regnerus et al., 2016) and with industry reports (Pornhub, 2024). Consumption is also highly skewed among users who visit adult websites. Conditional on any adult-site usage, machines in the highest tercile of consumption spend an average of 36.8 minutes per week on adult websites, whereas those in the lowest tercile spend only 0.13 minutes per week.[14]

4 Empirical Strategy: Stacked Difference-in-Differences

We implement a stacked difference-in-differences design that exploits both geographic and temporal variation in adult website shutdowns while addressing well-known sources of bias that can arise in staggered-adoption settings (Baker et al., 2022; Deshpande and Li, 2019). In our context, the key identifying assumption is that the timing of shutdowns across treated states is exogenous with respect to underlying trends in pornography consumption. To assess the plausibility of this assumption, we estimate an event-study specification that traces the evolution of pornography consumption throughout the analysis period.

[13] According to Pornhub's published traffic statistics, approximately 92% of visits in the U.S. originate from mobile devices (Pornhub, 2024).

[14] We emphasize that these statistics condition on having visited at least one adult site. 66% of devices in our sample never visit an adult website during the sample period.

Our event study compares trends in pornography consumption for machines located in states where major adult websites shut down to trends in states where those websites either never shut down or had not yet shut down during the relevant analysis window. Since Pornhub is by far the most prominent adult website that shut down in our sample, we use Pornhub's shutdown date to define the event.[15] Specifically, for each treated state k, we define the event time $\tau = 0$ as the week in which Pornhub access was blocked in that state, and estimate dynamic treatment effects during an event window $\tau \in [-16, +8]$. To do this, we construct a 25-week panel for each treatment state k at the machine-week level.[16] We refer to state k's panel as "cohort k." Each cohort includes machines located in: (i) the treated state k, (ii) states that never experience a Pornhub shutdown during our sample period, and (iii) states in which Pornhub has not yet shut down by event week $\tau = 8$. We then vertically stack the cohort-specific panels and estimate the following specification:

$$Y_{itk} = \alpha_{i,k} + \gamma_{t,k} + \sum_{\tau=-16}^{8} \beta_\tau * treated_{i,k} * \mathbb{I}\{t_{kt} = \tau\} + \varepsilon_{itk} \tag{1}$$

where Y_{itk} is minutes spent on a given website for machine i in cohort k in calendar week t, $\alpha_{i,k}$ are machine by cohort fixed effects, $\gamma_{t,k}$ are calendar week by cohort fixed effects, $treated_i$ is an indicator equal to one if a machine i is located in the treated state k, and $\mathbb{I}[t_{kt} = \tau]$ is an indicator function for whether the calendar week t for machines in cohort k is within the relative time period for the analysis. The reference period is $\tau = -1$, the week before Pornhub shuts down in state k.

We estimate this regression separately using various websites as the dependent variable. After accounting for the included fixed effects, identification comes from within-machine changes in website minutes over time relative to the timing of Pornhub shutdowns within each cohort. Accordingly, the coefficients β_τ capture differential changes in website minutes τ weeks after Pornhub has shut down, comparing machines in treated states to machines in the corresponding control states. We cluster standard errors at the state-level.

[15] See Table A.1 for the Pornhub shutdown dates. In practice, most shutdowns for compliant websites occurred in the weeks around which age verification laws went into effect. For Louisiana, the only treated state in which Pornhub did not shut down, we use the age verification law's effective date as the event date.

[16] The panel is balanced with respect to which states are included: we restrict to states that have at least two months of post-treatment data. The panel is unbalanced at the machine-week level due to panel churn.

In addition to the event study analysis, we present a pooled version of our estimates to increase power and to provide a concise, quantitative summary of our estimates. Following Bottan and Perez-Truglia (2015) and Miller (2023), we estimate the following regressions for each of our dependent variables:

$$\begin{aligned} Y_{itk} = \beta_{Pre} * treated_{ik} * \mathbb{I}\{t_{kt} = \tau \in [-16, -5]\} \\ + \beta_{ShortTerm} * treated_{ik} * \mathbb{I}\{t_{kt} = \tau \in [0, 3]\} \\ + \beta_{LongTerm} * treated_{ik} * \mathbb{I}\{t_{kt} = \tau \in [4, 8]\} + \varepsilon_{itk} \end{aligned} \tag{2}$$

For this pooled regression, the reference period is $\tau \in [-4, -1]$. The coefficients $\beta_{ShortTerm}$ and $\beta_{LongTerm}$ summarize the average effects of shutdowns on minutes for the website of interest in the first month and the second month after the shutdown, respectively. In the rest of this paper, when we say "long-term" effects, we are referring to $\beta_{LongTerm}$ estimates. The coefficient β_{Pre} summarizes our pre-trends test, that is, whether the website of interest was receiving more traffic than usual in treated states during the periods *before* the shutdowns occurred. In general, we expect β_{Pre} to be small or zero. We include the same fixed effects as in the event-study analysis, estimate the effects for each site separately, and cluster the standard errors at the state-level.

Finally, to explore heterogeneity in cessation by subgroup characteristics, we estimate the pooled regression in equation 2 separately for subgroups defined by the device and demographic characteristics summarized in Table 1. Because subgroups differ in their baseline consumption patterns, when we report results, we divide our coefficient estimates by the treatment group's pre-treatment pooled consumption for all XXX sites in period $\tau \in [-4, -1]$. With this normalization, the βs can be interpreted as percent changes relative to each group's baseline XXX consumption.

5 Results

5.1 Event study results

Panel A of Figure 1 illustrates trends in Pornhub consumption in the weeks after Pornhub shuts down in a state, using our event study specification from equation 1. Consumption is measured in units of minutes per machine per week. In the pre-treatment period, coefficients are close to zero, which shows that trends in Pornhub consumption were similar in treated and control states prior to shutdowns. Following the shutdown event, the estimates become negative, indicating that Pornhub consumption declined sharply in treated states. The estimated effects stabilize quickly, with relatively little change after the first week following the shutdown.

Importantly, the magnitude of this consumption reduction implies that a meaningful share of pre-shutdown Pornhub consumption persists even after access is blocked. To see this, consider the summary coefficient estimates from our pooled specification 2, which are reported in Column (1) of Table 2. Our point estimate for impacts four to eight weeks after the shutdown indicates that Pornhub consumption fell by approximately 0.76 minutes per machine ($p < 0.001$), which is a 54% decline relative to the baseline mean. Therefore, 46% of time spent on Pornhub remained observable after the shutdown. Our preferred interpretation is that this residual is driven by VPN-based circumvention.[17]

We next study impacts on the largest noncompliant substitute websites, XVideos and XNXX. Panels B and C of Figure 1 show that time spent on both of these competing platforms increases following Pornhub shutdowns. An analyst who failed to account for this substitution and only considered the impacts on Pornhub consumption (from Panel A) would overstate the impact of the shutdowns on total porn consumption. That said, substitution to noncompliant sites did not completely undo the impacts of the shutdowns. In Panel D of Figure 1, we plot

[17]Some of the observed time spent on Pornhub after the shutdowns is surely due to factors other than VPN-based circumvention. For example, a user could visit the Pornhub "landing page" (which displayed a message announcing the shutdown) without viewing any content or there could be measurement error in our machine-level geographical variable. We do not think that such traffic accounts for a quantitatively meaningful amount of Pornhub time. Some patterns in our data are consistent with our interpretation. We do not observe an effect that grows over time, even though mistaken visits to the landing page would presumably diminish as people become familiar with the law. And our estimates are stable when we use a more restrictive sample definition where state is measured with less error (Table A.2).

the effect of the shutdowns on total time spent across all adult websites in our sample. On net, time spent on adult websites in treated states falls relative to control states following the shutdown events. Using our pooled regression, we estimate that long-term total adult website consumption falls by roughly 0.5 minutes per machine, which is 8% of the baseline mean ($p < 0.001$).

One natural question concerns the ultimate destination of those 0.5 minutes. In Figure A.2, we summarize our estimates of how restrictions on adult websites affect time spent on some non-adult websites of interest. Most estimates are imprecise, because these other websites are exposed to other local demand shocks beyond the age verification laws. We estimate positive effects on some social media platforms that hold adult content (TikTok and Reddit) but not others (Twitter/X).[18] Of particular interest is the effect on use of DuckDuckGo, a search engine that is designed and marketed as a privacy-preserving tool. We do not observe an increase in usage of DuckDuckGo after restrictions are implemented (our point estimate is negative and insignificant). This fact alleviates concerns about differential trends in the use of other privacy-preserving tools, like Google Chrome's Incognito Mode, in treated versus control states.

Overall, these results indicate that age verification policies did reduce total time spent at adult websites, even among a population of adults who were not directly targeted by the laws. We also show, though, that the magnitude of the policy-induced consumption reduction was muted by consumers' and firms' behavioral responses. We now turn to an exercise that quantifies the relative importance of distinct kinds of behavioral responses for the overall impact of the policy.

5.2 Cessation, noncompliance, circumvention and substitution

The goal of our exercise is to decompose the baseline time spent on adult websites into components representing the impacts of distinct behavioral responses to the policy. Conceptually, there are four mutually exclusive possibilities for how baseline consumption can respond to

[18] One hypothesis might have been that any time that was eliminated from adult websites would have flowed to these social media based alternatives. This was not the case: the sum of our point estimates for effects on time spent across these three sites is less than half as large as our estimate of the reduction in time spent on adult websites.

the policy. The first is that people may stop visiting adult websites; we call this response *cessation*. Then, there are three ways that consumption may persist. Some baseline consumption was occurring on websites that did not implement access restrictions, like XVideos. There is no obvious channel by which the regulation would reduce consumption at these websites. We call this source of persistent consumption *noncompliance*. There is also persistence by some users who continued to visit compliant websites like Pornhub through VPN-based *circumvention*. Finally, we observe *substitution* from compliant sites to noncompliant sites.

To analyze the relative importance of these channels, we need to categorize how adult websites responded to the policy. Therefore, for this subsection only, we restrict attention to the top 25 adult websites in our sample (in terms of minutes per week) and manually code each website as either compliant or noncompliant. We coded websites as compliant if they implemented any access restrictions in states with age verification laws.[19] The top 25 websites account for roughly three quarters of time spent on adult websites in the baseline period, and we do not detect evidence of substitution from top 25 websites to other websites or vice versa.[20]

We illustrate our decomposition in Figure 2. The bar on the left side of the figure divides the 4.68 minutes spent visiting the top 25 adult websites in the baseline period into minutes at five website categories: Pornhub (1.40 minutes), other compliant websites (0.99 minutes), XVideos (0.97 minutes), XNXX (0.78 minutes), and other noncompliant websites (0.54 minutes). This bar alone immediately demonstrates the importance of noncompliance: 49% of baseline minutes were spent at websites that did not implement access restrictions. It is not surprising that porn consumption at noncompliant websites would persist.

The rest of Figure 2 divides the remaining 2.39 minutes into cessation, circumvention, and substitution. In our pooled specification with total minutes at top 25 sites as the dependent variable, we estimate that access restrictions reduce total time on top 25 sites by 0.49 minutes per machine-week. This measure of cessation corresponds to 10% of total baseline minutes and 21% of baseline minutes spent on compliant sites. Since we have partitioned the top 25 sites

[19] Access restrictions include either state-level shutdowns (like Pornhub) or website changes that do not allow users to access adult content without completing an age verification process. We coded the websites by first looking for evidence from contemporary news reports, and then by checking whether content was accessible in states with laws in May of 2026.

[20] We summarize estimates of long-term effects on time spent at each of the top 25 websites separately and for time spent at adult websites outside of the top 25 in Figure A.1

into five mutually exclusive categories, the overall reduction in time spent at these sites can be expressed as a sum of analogous effects on time spent at each of the five website categories. We illustrate these category-specific estimates in Figure 2.[21] We measure circumvention by comparing the estimated reduction in consumption at compliant sites to a "no circumvention" benchmark where users spend zero minutes at compliant sites.[22] Circumvention accounts for 1.41 minutes, which is 31% of total baseline minutes and 59% of baseline minutes spent on compliant sites. We measure substitution as the sum of the effects on the three kinds of noncompliant sites. Substitution accounts for 0.49 minutes, which is 10% of total baseline minutes and 21% of baseline minutes spent on compliant sites. Almost all substitution flows to the two dominant noncompliant sites, XVideos and XNXX.

5.3 Heterogeneity Analysis

To understand how effects vary across user characteristics and device types, we estimate equation 2 separately for relevant subgroups. Figure 3 summarizes heterogeneity in cessation, reporting long-term effects on total time at adult websites for each subgroup. To facilitate comparison across subgroups with different baseline adult website consumption levels, we report effects in units of percent changes relative to baseline time at adult websites. Recall that in our full sample, we estimated that 8% of pre-shutdown adult website consumption ceased. We report analogous group-specific estimates of substitution and circumvention in Figure A.3 and Figure A.4.

One especially important subgroup from a policy perspective is minors. After all, the age verification laws were explicitly designed to target pornography consumption by children under age 18. It is not obvious whether the restrictions' effects would be smaller or larger for minors. Teenagers may be more technologically sophisticated than adults, and therefore more aware of VPN-based circumvention. On the other hand, younger users, having less experience,

[21] To be precise, we illustrate estimates of $\beta_{LongTerm}$ from equation (2), using minutes spent at websites in the category as the dependent variable.

[22] This definition classifies time spent by users who successfully (and legally) verify their age to access content at a compliant website as circumvention. In practice, we expect that such activity is quantitatively unimportant. Users could not verify age to access Pornhub's content in states with restrictions, because Pornhub shut down completely. And Pornhub accounts for the vast majority of time on compliant sites in the baseline data.

may be less familiar with the available substitutes. Unfortunately, Comscore does not collect data on devices owned by minors, so we do not have direct evidence regarding effects on users under the age of 18.

In the absence of direct evidence, we discuss two sources of indirect evidence on how minors responded to the law. The first concerns how young adults (18-24 year olds) responded – presumably, teenagers' internet habits are not too different from this age group's. We find less cessation among young adults than among 25-44 year olds ($p = 0.045$). The difference is apparently driven by more VPN-based circumvention from young adults (Figure A.3). Our second source of indirect evidence concerns the difference in effects between desktop machines in households with children present versus similar machines in households without children present. The idea is that some of the adult website consumption on computers in households with children likely arises from children browsing on family computers.[23] The evidence here is consistent with more cessation among minors: cessation was larger among households with children present than without ($p = 0.026$). We highlight that this second test points in the opposite direction from our first test, which indicated less cessation among minors. Overall, the evidence is inconclusive as to whether minors were more or less impacted by age verification laws.

While some of the differences across subgroups in Figure 3 are hard to interpret, a key high-level takeaway is that behavioral responses meaningfully attenuate the effects of age verification laws on porn consumption across a broad variety of users and device types. For all subgroups that we study, we estimate that total time spent on adult websites fell by 15% or less after restrictions were implemented.

6 Discussion

This paper studies the staggered rollout of age verification laws across U.S. states beginning in January 2023. Using detailed browsing data from a large panel of internet users, we examine how consumers respond when Pornhub blocks access in states adopting these laws. About

[23] We focus on desktop computers for this comparison because children may have a greater ability to access a family computer, compared to their parent's phone.

half of the time spent at adult websites was spent on websites that did not comply with the regulation. Among users who had been spending time at compliant sites, we document three central behavioral responses: circumvention, substitution, and cessation. Following the shutdown, a substantial share of consumption at compliant websites persists through VPN-based circumvention, some consumption shifts toward competing noncompliant websites, and some online pornography consumption ceases completely.

These results come with important limitations. First, our sample is not perfectly representative of adult website browsing in the U.S., both because the Comscore panel is not perfectly representative of the U.S. population along demographic and device-type dimensions and because we do not capture browsing activity conducted in private browser modes. We do note that within our sample, circumvention and substitution are relevant margins of adjustment across demographic groups and across device types. Second, our analysis focuses on adult users rather than minors, even though age verification laws are primarily intended to affect children's access to pornography. Because comparable browsing data for minors are not available, our results should be interpreted as evidence on how adults respond to these laws and as an indirect proxy for the mechanisms through which minors may also adapt.

Despite these limitations, our results have implications both for the regulation of online pornography and for digital regulation more broadly. We tackle these implications in order from narrowest to broadest.

Policymakers in the U.S. and worldwide are currently deciding whether and how to implement age verification for adult websites. While one could support or oppose age verification regulations for many reasons, surely one important consideration is how the regulations impact pornography consumption. We show that in the U.S., the impact of these regulations on consumption has been limited by multiple forms of leakage, including circumvention through VPNs and substitution toward noncompliant alternatives. Some have proposed alternative paradigms for age verification. For example, with "device-based" age verification, instead of the website conducting an age check, devices like smartphones collect verified age information at setup and transmit age information whenever the device tries to access an adult site

(whether through a VPN or not).[24] Such an approach could reduce VPN-based circumvention, but it would not block users from substituting to noncompliant substitutes.

Many empirical questions that are important for pornography regulation were beyond the scope of our analysis. We document that a prominent regulation reduced pornography consumption, but we have nothing to say about the welfare effects of these consumption reductions, either for pornography consumers or for other members of society. Our analysis also does not speak to how regulations affect the production of adult content, even though the supply side of the industry is rapidly evolving and regulations could affect the welfare of performers and others involved in the production process. The applied microeconomics toolkit seems well-suited to answer such questions, but to date economics has paid little attention to pornography.[25] We hope that future work will provide useful insights.

Beyond pornography, policymakers are increasingly considering age verification requirements and related access restrictions in settings such as online gambling, social media, and other forms of regulated digital content. Our study provides some of the strongest evidence to date on how age verification laws affect online behavior and how consumers adjust to new forms of digital regulation. The effectiveness of such policies may depend critically on the ease of technological circumvention and on the availability of substitutes.

Most broadly, our findings contribute to a growing literature on online content regulation and consumer adaptation to digital restrictions. We analyze a case study where regulations targeted a dominant online platform in a market with close substitutes and relatively low circumvention costs. In our context, technological circumvention and substitution toward alternative providers muted the impact of access restrictions on overall consumption. These results highlight that the design and enforceability of online regulations matter for their ultimate impact. In any given context, understanding the relevant channels by which consumers can continue to engage in the targeted online behavior will help policymakers design policies that align better with their own goals.

[24] Aylo, the parent company for Pornhub, advocates for device-based age verification. See for example https://www.aylo.com/assets/files/age_verification_fact_sheet.pdf

[25] Notable exceptions include Kendall (2006), Edelman (2009), Bhuller et al. (2013), Doran and Price (2014), Malcolm and Naufal (2016), Gibbons and Rossi (2021), Mubasshera (2024), and recent working papers from Madio et al. (2026) and Agarwal et al. (2026).

References

Agarwal, S., U. M. Ananthakrishnan, L. Madio, M. F. Mitchell, M. Quinn, and C. Reggiani (2026, April). Are You 18? Age Verification and Adult-only Consumption in the US.

Baker, A. C., D. F. Larcker, and C. C. Y. Wang (2022, May). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics 144*(2), 370–395.

Bhuller, M., T. Havnes, E. Leuven, and M. Mogstad (2013, October). Broadband Internet: An Information Superhighway to Sex Crime? *The Review of Economic Studies 80*(4), 1237–1266.

Bottan, N. L. and R. Perez-Truglia (2015, September). Losing my religion: The effects of religious scandals on religious participation and charitable giving. *Journal of Public Economics 129*, 106–119.

Bowman, K. and D. A. Cox (2022, October). Gender, Generation and Abortion: Shifting Politics and Perspectives After Roe - The Survey Center on American Life.

Bursztyn, L., A. L. Duckworth, R. Jiménez-Durán, A. Leonard, F. Milojević, C. Roth, and C. R. Sunstein (2026, April). Why Bans Fail: Tipping Points and Australia's Social Media Ban.

Chen, Y. and D. Y. Yang (2019, June). The Impact of Media Censorship: 1984 or Brave New World? *American Economic Review 109*(6), 2294–2332.

Danaher, B., J. Hersh, M. D. Smith, and R. Telang (2020, June). The Effect of Piracy Website Blocking on Consumer Behavior. *MIS Quarterly 44*(2), 631–660.

Deshpande, M. and Y. Li (2019, November). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy 11*(4), 213–248.

Doran, K. and J. Price (2014, December). Pornography and Marriage. *Journal of Family and Economic Issues 35*(4), 489–498.

Edelman, B. (2009, March). Markets: Red Light States: Who Buys Online Adult Entertainment? *Journal of Economic Perspectives 23*(1), 209–220.

Gibbons, M. A. and M. A. Rossi (2021). When You Can't Tube... Impact of a Major Youtube Outage on Rapes. *Economic Inquiry 59*(2), 762–775.

Goldberg, S. G., G. A. Johnson, and S. K. Shriver (2024, February). Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy 16*(1), 325–358.

Kendall, T. D. (2006, September). Pornography, Rape, and the Internet.

Kircher, T. and J. Foerderer (2024, February). Ban Targeted Advertising? An Empirical Investigation of the Consequences for App Development. *Management Science 70*(2), 1070–1092.

Kohut, T., K. A. Dobson, R. N. Balzarini, R. D. Rogge, A. M. Shaw, J. K. McNulty, V. M. Russell, W. A. Fisher, and L. Campbell (2021, July). But What's Your Partner Up to? Associations Between Relationship Quality and Pornography Use Depend on Contextual Patterns of Use Within the Couple. *Frontiers in Psychology 12*.

Kristof, N. (2020, December). Opinion | The Children of Pornhub. *The New York Times*.

Lang, D., B. Listyg, B. V. Ross, A. V. Musquera, and Z. Sanderson (2026, January). Age Verification and Public Adaptation: A Pre-Registered Synthetic Control Multiverse. *Journal of Law & Empirical Analysis*, 2755323X251415424.

Madio, L., M. F. Mitchell, M. Quinn, and C. Reggiani (2026, March). Asymmetric Content Moderation in Search Markets: The Case of Adult Websites.

Malcolm, M. and G. Naufal (2016, June). Are Pornography and Marriage Substitutes for Young Men? *Eastern Economic Journal 42*(3), 317–334.

McCarthy, S. A. (2023, August). Poll Shows Vast Majority of Americans Want Age Verification Laws for Porn Sites. https://washingtonstand.com/news/poll-shows-vast-majority-of-americans-want-age-verification-laws-for-porn-sites.

Mestre-Bach, G., A. Villena-Moya, and C. Chiclana-Actis (2024, April). Pornography Use and Violence: A Systematic Review of the Last 20 Years. *Trauma, Violence & Abuse 25*(2), 1088–1112.

Miller, D. L. (2023, May). An Introductory Guide to Event Study Models. *Journal of Economic Perspectives 37*(2), 203–230.

Mubasshera, H. (2024). Pornography usage during adolescence: Does it lead to risky sexual behavior? *Health Economics 33*(8), 1682–1704.

Newport, F. (2025, September). Trends in U.S. Adults' Acceptance of Moral and Values Behaviors. https://news.gallup.com/opinion/polling-matters/694550/trends-adults-acceptance-moral-values-behaviors.aspx.

Pornhub (2024, December). 2024 Year in Review - Pornhub Insights.

Regnerus, M., D. Gordon, and J. Price (2016, September). Documenting Pornography Use in America: A Comparative Analysis of Methodological Approaches. *Journal of Sex Research 53*(7), 873–881.

Robb, M. B. and S. Mann (2023, January). Teens and Pornography. Technical report, Common Sense Media.

Semrush (2026, March). Most Visited Websites in the World, Updated March 2026.

Spencer, N. (2025, March). Age Verification Laws, Firm Responses, and Porn Consumption.

Toscano, M. (2024, March). Online Age Verification Laws Are a Bet Worth Making | Opinion - Newsweek. https://www.newsweek.com/online-age-verification-laws-are-bet-worth-making-opinion-1874764.

U.S. Supreme Court (1997, June). Reno v. American Civil Liberties Union, 521 U.S. 844. 521 U.S. 844. No. 96-511.

Wright, P. J., R. S. Tokunaga, and A. Kraus (2016). A Meta-Analysis of Pornography Consumption and Actual Acts of Sexual Aggression in General Population Studies. *Journal of Communication 66*(1), 183–205.

Figure 1: Event Study Estimates of the Effects of Pornhub Shutdowns

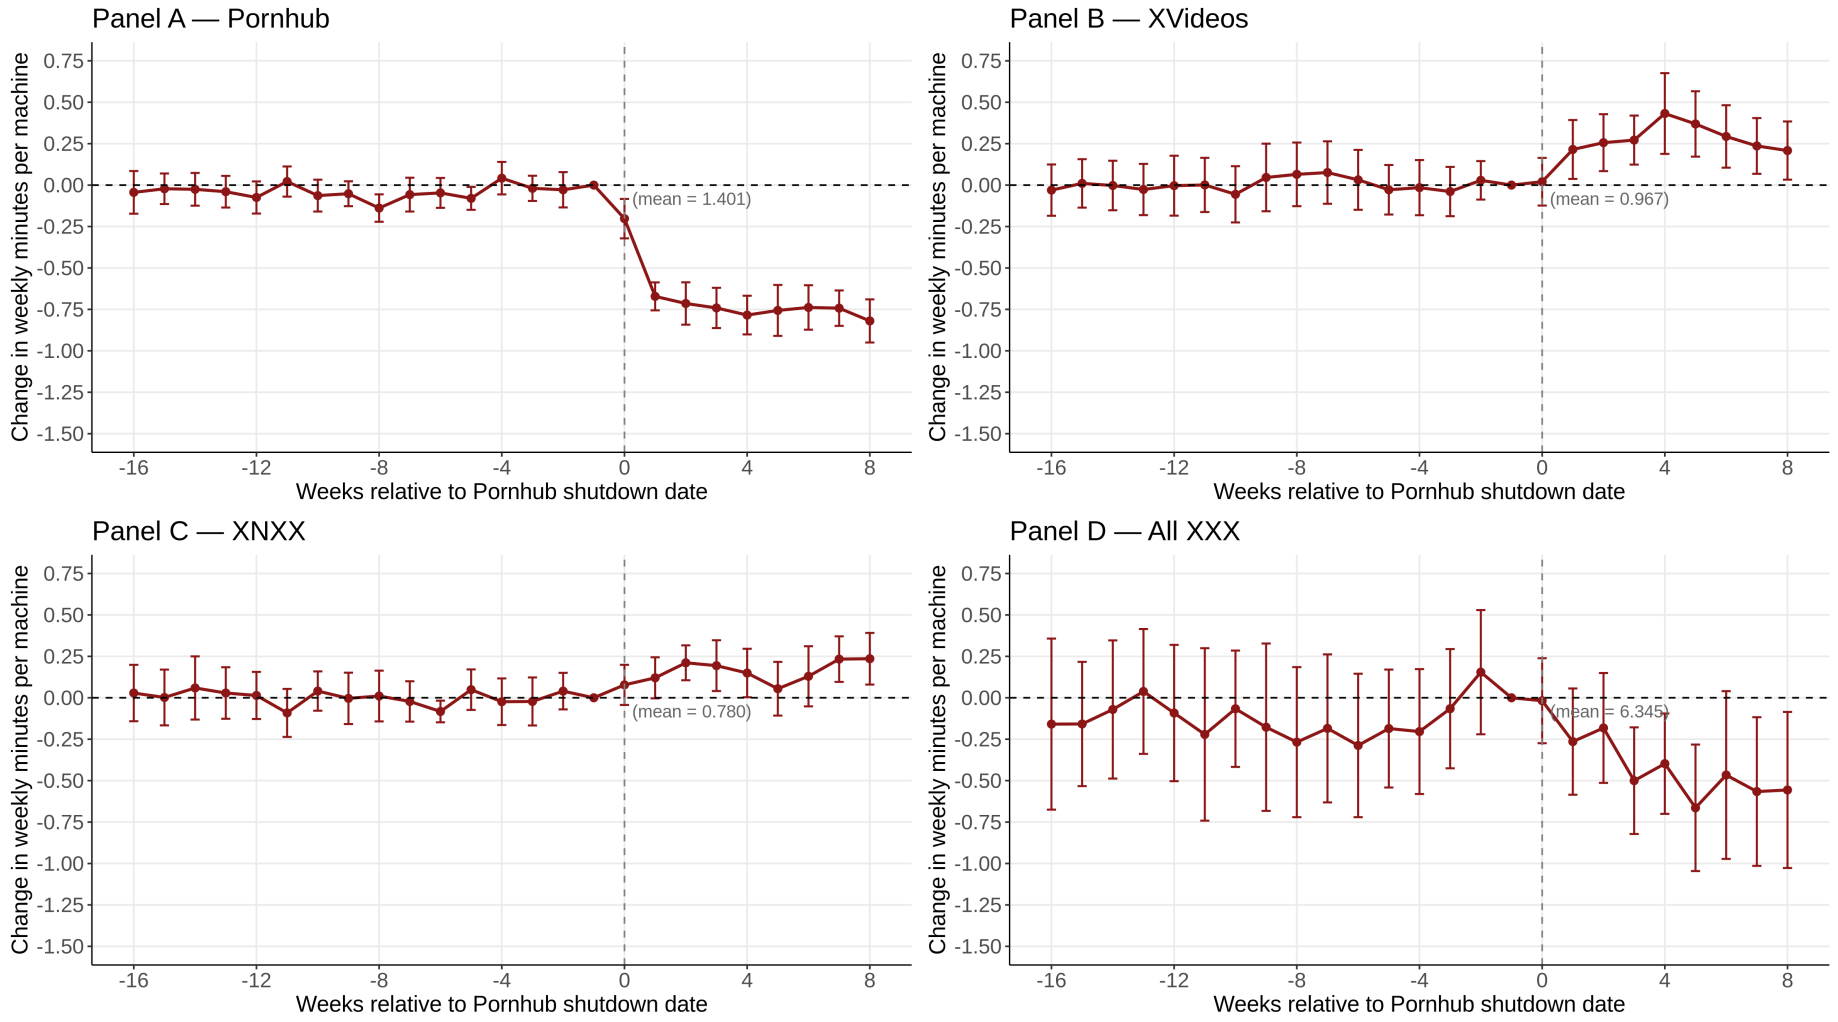

Notes: We report estimates from our event study specification 1 with dependent variables for time spent on three major adult websites (Panels A-C) and for time spent on all adult websites (Panel D). The dependent variable is winzorized minutes per machine-week. The reference period is $t = -1$ (one week before Pornhub shut down). The baseline mean is the average of the dependent variable in that reference period for treated machines. All four panels share a common y-axis for ease of comparison. Error bars show 95% confidence intervals. Standard errors are clustered at the state level.

Figure 2: Noncompliance, Circumvention, Substitution, and Cessation

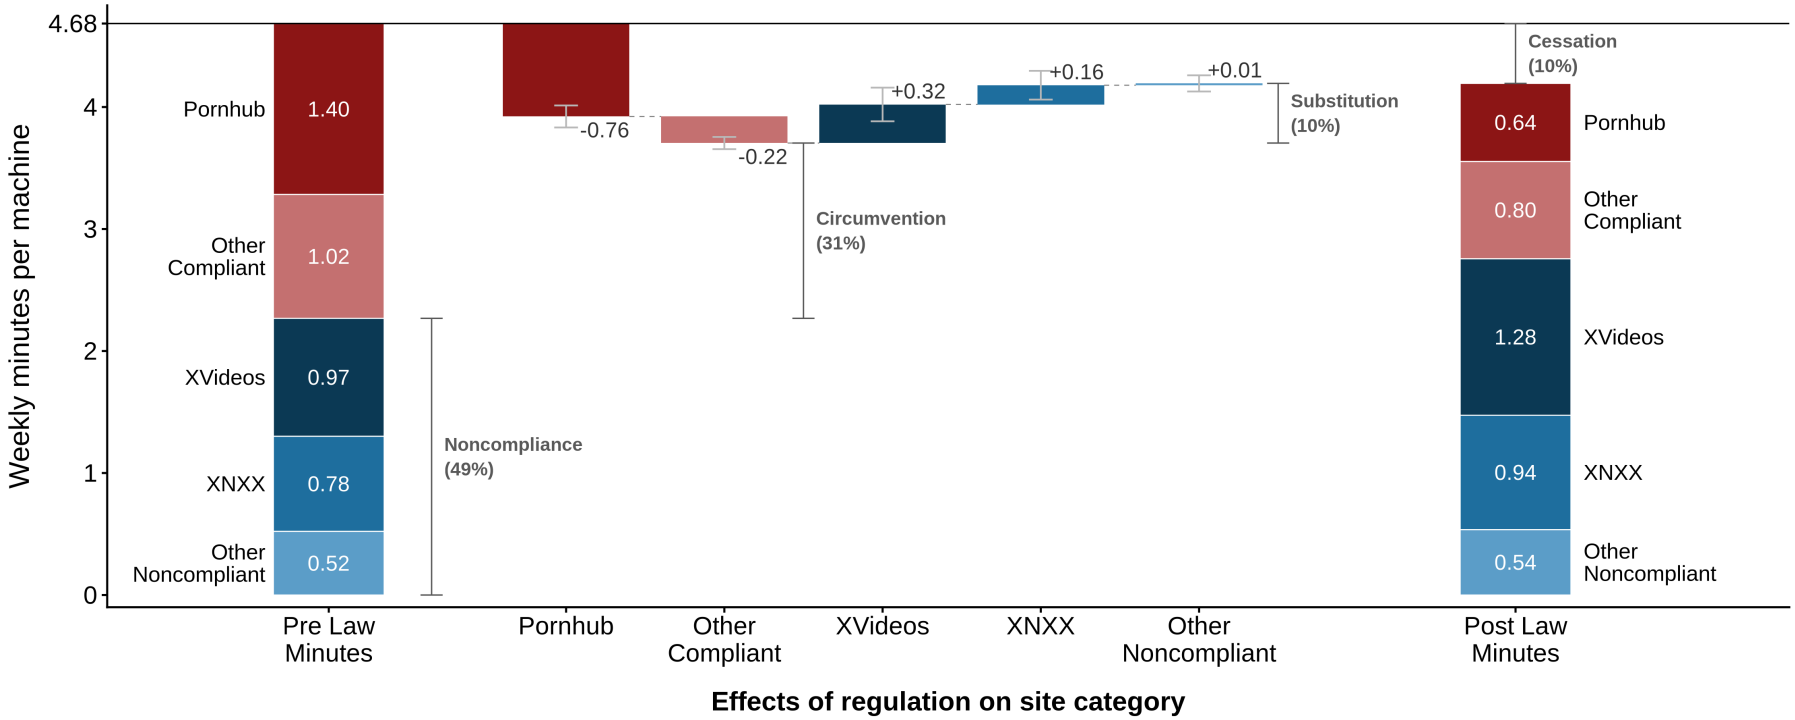

Notes: This figure illustrates how the regulation affected adult website browsing time across five different categories of adult website within the top 25. Compliant sites (Pornhub and Other Compliant) are shown in shades of red; noncompliant sites (XVideos, XNXX, and Other Noncompliant) are shown in shades of blue. For the full list of websites, see Figure A.1. The leftmost bar shows minutes at each site category for treated machines in the baseline period($t = -1$). Total baseline time spent on top 25 adult websites (4.68 minutes) is the sum of these category means. The five bars in the middle of the plot illustrate long-term estimates of the effects of regulations on browsing time at each category ($\beta_{LongTerm}$ from equation 2). These bars are laid out as a "waterfall plot," so that the height where each bar "ends" is the cumulative sum of that effect plus the effects to the left of it. Since the five categories are mutually exclusive, the height of the final bar indicates the effect of regulation on total browsing time at top 25 sites. We illustrate the composition of post-law browsing time, given these effects, in the right-most bar. We label four mechanisms of interest (noncompliance, circumvention, substitution, and cessation) as described in the main text. Briefly, noncompliance is the share of pre-law minutes already spent on noncompliant sites before the law took effect, circumvention is the share of pre-law minutes that remain on compliant sites after the shutdown, substitution is the share of pre-law minutes shifted to noncompliant sites, and cessation is the share of pre-law consumption that does not persist on adult websites.

Figure 3: Heterogeneity in Cessation

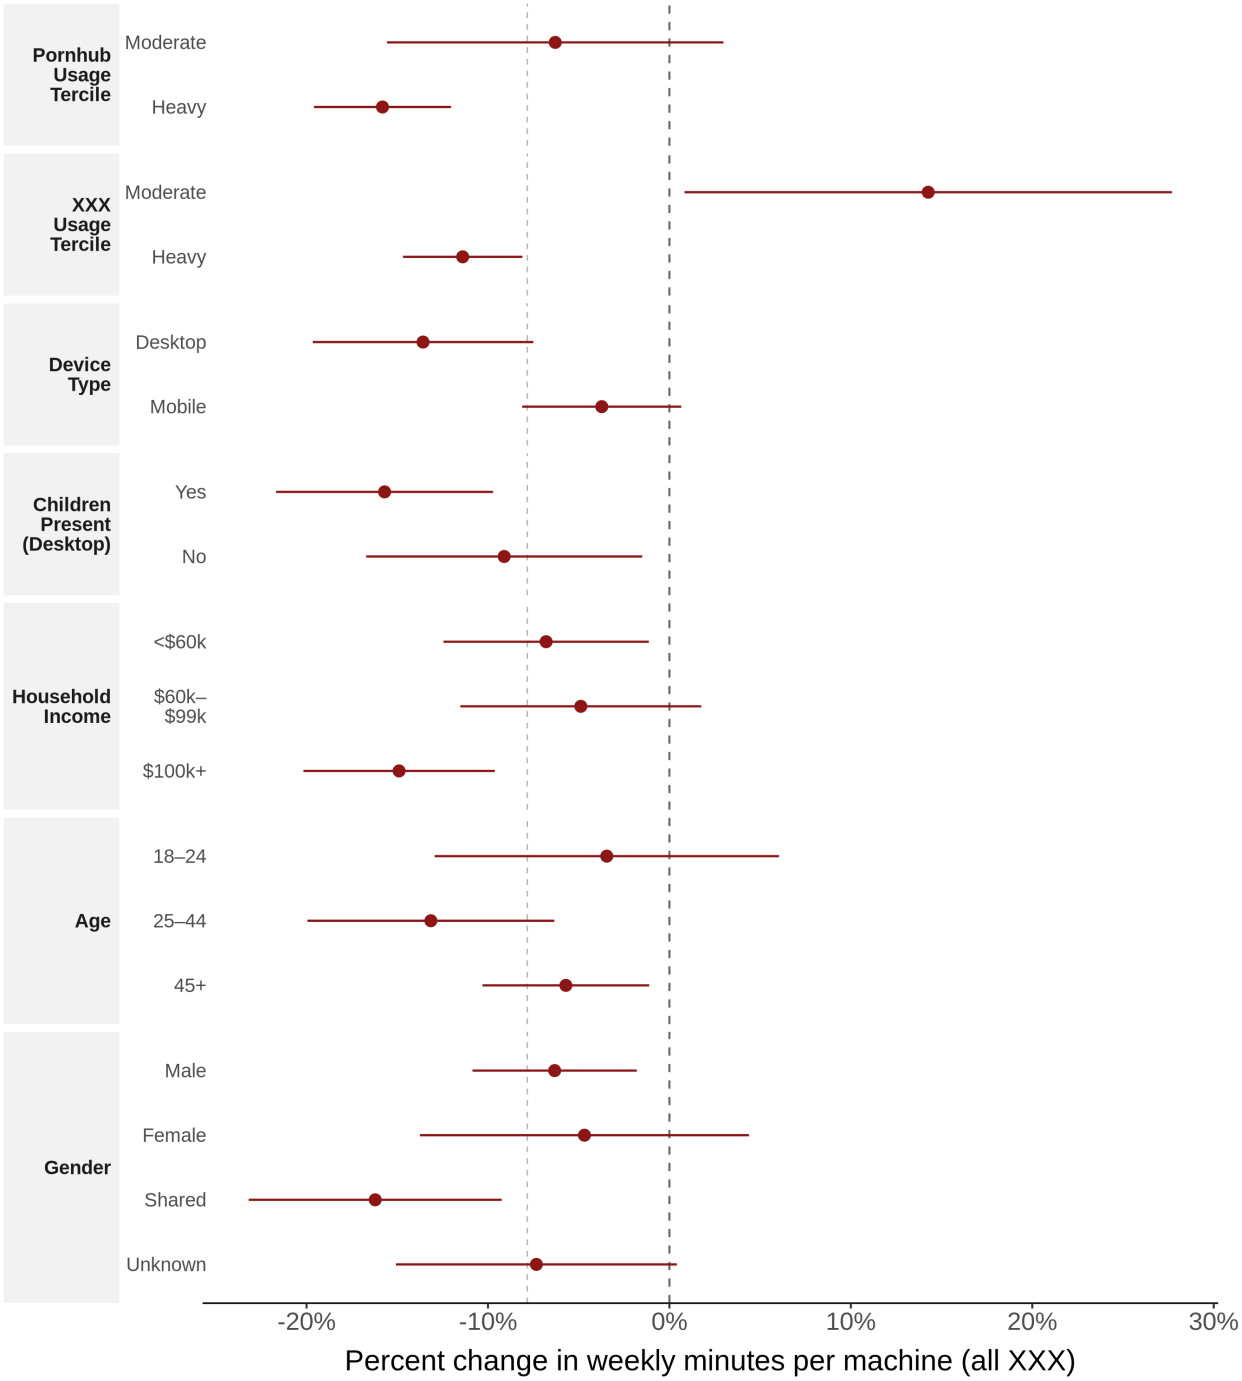

Notes: Each point represents a subgroup-level estimate of how the regulation affected total time spent on adult websites. Specifically, each point shows a long-term estimate from our pooled specification (equation 2) with All XXX minutes as the dependent variable, estimated for the indicated subgroup. We report all estimates in units of "percent of baseline consumption;" to do this we divide each subgroup's $\beta_{LongTerm}$ estimate by the baseline ($t \in [-4, -1]$) dependent variable. Error bars show 95% confidence intervals. Standard errors are clustered at the state level. The lighter dashed line marks the analogous estimate for the full analysis sample. Pornhub and XXX usage terciles are defined among machines with positive average weekly minutes on Pornhub or any XXX site, respectively, during the pre-period (weeks −16 through −1). Non Visitor and Light tercile subgroups are excluded from this figure; when baseline consumption is zero or very small reporting effects in percent units is uninformative. When reporting on the "children present" split, we restrict both subsamples to desktop machines only. See the notes for Table 1 for other details on demographic variables.

Table 1: Summary Statistics

	Share of Sample	N	Pornhub		XVideos		XNXX		All XXX	
			Visits	Minutes	Visits	Minutes	Visits	Minutes	Visits	Minutes
State										
Treated	52.8%	292,985	19.1%	0.98	12.3%	0.69	9.2%	0.58	33.2%	4.30
Non-Treated	47.2%	261,881	20.4%	1.05	12.6%	0.64	9.2%	0.52	35.8%	4.63
Gender										
Male	18.9%	104,795	31.5%	2.56	22.7%	2.29	18.2%	2.04	50.0%	12.54
Female	13.5%	74,727	21.9%	1.14	13.1%	0.59	12.3%	0.64	39.4%	3.87
Shared	18.4%	102,189	22.4%	0.70	13.5%	0.29	8.9%	0.16	39.5%	2.94
Unknown	49.2%	273,155	13.6%	0.50	7.9%	0.20	5.0%	0.11	25.2%	2.08
Age										
18–24	14.6%	81,225	27.5%	1.35	15.7%	0.48	9.6%	0.23	44.9%	4.06
25–44	29.2%	162,267	23.8%	1.27	14.8%	0.77	11.3%	0.60	40.6%	4.93
45+	56.1%	311,371	15.6%	0.79	10.4%	0.66	8.0%	0.62	28.5%	4.31
Household Income										
<$60k	55.3%	306,652	18.8%	1.12	12.6%	0.82	10.0%	0.73	33.2%	5.16
$60k–$99k	12.3%	68,518	22.3%	1.29	15.0%	0.93	11.9%	0.80	39.2%	6.24
$100k+	32.4%	179,696	20.3%	0.72	11.2%	0.31	6.8%	0.16	34.8%	2.58
Children Present										
Yes	50.3%	279,252	23.3%	1.01	13.9%	0.49	9.6%	0.36	39.8%	3.83
No	49.7%	275,614	16.1%	1.01	11.0%	0.84	8.8%	0.75	29.0%	5.09
Device Type										
Desktop	87.1%	483,291	16.3%	0.56	9.6%	0.23	6.2%	0.13	29.7%	2.34
Mobile	12.9%	71,575	42.9%	4.10	31.9%	3.59	29.5%	3.45	66.4%	18.75
XXX Usage Tercile										
Non Visitor	65.6%	363,762	-	-	-	-	-	-	-	-
Light	11.5%	63,702	26.7%	0.03	10.0%	0.01	5.9%	0.01	100.0%	0.13
Moderate	11.5%	63,701	64.1%	0.68	33.9%	0.18	22.4%	0.13	100.0%	1.90
Heavy	11.5%	63,701	80.9%	8.11	64.5%	5.61	51.9%	4.70	100.0%	36.79
Total	100.0%	554,866	19.7%	1.01	12.4%	0.67	9.2%	0.56	34.4%	4.46

Notes: This table reports summary statistics at the machine level. For each of three websites (Pornhub, XVideos, and XNXX) we report the average of winsorized minutes per week across all machines in our analysis sample ("Minutes"). We also report the share of machines with any time on each site, over the whole analysis period ("Visits"). The "All XXX" column reports these statistics with reference to all adult websites, not just our top three. The bottom row reports on these characteristics for the whole sample; other rows split the sample according to Comscore demographics. "Treated" states are states where Pornhub shutdown at some point during the analysis period. For Gender, Male and Female correspond to machines with a single linked user; Shared corresponds to desktop machines with multiple linked users; Unknown corresponds to machines for which no user gender information was available in Comscore demographic files. For Age, subgroups reflect the household head's age bracket for shared desktop machines. Children Present refers to whether children are present in the household, not necessarily whether children use the device in the panel. Desktop includes laptops and other personal computers; mobile includes tablets. For XXX usage tercile, machines with zero visits to adult websites are classified as Non Visitor; the three terciles divide machines with positive usage into equal thirds by average weekly winsorized All XXX minutes.

Table 2: Pooled Regression Results

	Pornhub (1)	XVideos (2)	XNXX (3)	Other XXX (4)	All XXX (5)
β_{Pre}	-0.0506**	0.0139	0.0036	-0.0921	-0.1252
	(0.0221)	(0.0364)	(0.0412)	(0.0928)	(0.1251)
$\beta_{ShortTerm}$	-0.5794***	0.1960***	0.1520***	0.0184	-0.2130***
	(0.0405)	(0.0452)	(0.0360)	(0.0867)	(0.0760)
$\beta_{LongTerm}$	-0.7619***	0.3163***	0.1578***	-0.2093**	-0.4971***
	(0.0460)	(0.0704)	(0.0601)	(0.0834)	(0.1025)
Baseline mean	1.4015	0.9669	0.7796	3.1969	6.3449
N (obs)	33,558,770	33,558,770	33,558,770	33,558,770	33,558,770
N (machines)	554,866	554,866	554,866	554,866	554,866

Notes: We report estimates from the pooled regression in equation 2. All regressions include cohort × machine and cohort × calendar week fixed effects. The dependent variable is (winsorized) weekly minutes per machine; the baseline mean row reports these averages for treated machines in week −1. N (obs) is the number of machine × cohort × calendar-week observations in the stacked panel. N (machines) is the number of unique machines in the stacked panel. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A.1 Age Verification Laws

Under Louisiana's HB 142, "commercial entities" that distribute "material harmful to minors" on websites containing a "substantial proportion" of such material are required to implement "reasonable age verification methods." Websites that fail to do so can be held liable for damages. A follow-up law, HB 77, enacted in August 2023, strengthened enforcement by authorizing civil penalties and allowing the state attorney general—not just private individuals—to bring legal action against noncompliant websites.[26]. Subsequent state laws largely adopted these same provisions, though penalties and enforcement mechanisms vary across states.

In practice, the laws are designed to target pornography websites specifically rather than platforms where adult content appears only incidentally.[27]. "Commercial entities" refers to the operators of the websites themselves, while internet service providers, search engines, and cloud hosting services are generally exempt from liability. "Material harmful to minors" is typically defined using language derived from Ginsberg v. New York (1968), referring to material appealing to the prurient interests of minors and lacking serious literary, artistic, political, or scientific value. Most states define a "substantial proportion" as websites in which at least one-third of the content is pornographic.

To comply with these laws, websites must verify users' ages either directly or through third-party verification providers. Common methods include uploading a government-issued ID, providing credit card information, or using biometric tools such as facial recognition software to estimate age. Many websites rely on specialized third-party verification services, in part because state laws also impose restrictions and potential penalties related to the handling of users' personal information.

[26]The civil penalties described in HB 77 can reach up to $5,000 for each day of violation. The attorney general may also request that the court impose an additional $10,000 for each section of the law the commercial entity violates. Besides these civil penalties, commercial entities are also responsible for all costs, expenses, and fees related to the investigation, including attorney fees.

[27]Social media is the second most popular source for viewing pornography among teens. Robb and Mann (2023) report that among teens who intentionally consume pornography (44% of all teens), 38% consume pornography weekly on social media, while 44% consume it weekly on pornography sites.

Enforcement occurs primarily through the court system. Individuals or state attorneys general may sue websites that fail to implement adequate age verification procedures, and penalties can be substantial. In states such as Texas and Louisiana, fines can reach thousands of dollars per day of noncompliance, in addition to legal costs and investigative expenses.

There have been legal responses to age verification laws. The most prominent response has been from the Free Speech Coalition (FSC). The FSC is a trade organization that represents the adult industry. They have filed a number of lawsuits against state legislatures, claiming that age verification laws violate First Amendment rights. The most prominent one of these cases is *Free Speech Coalition, Inc. v. Paxton (Texas)*. This case made it to the Supreme Court, and in June, 2025, the Court upheld the law, affirming the government's interest in protecting minors from accessing adult content. This court case established important precedent for the legality of age verification laws across the country.

A.2 Mapping Comscore Markets to States

We were unable to obtain a crosswalk that mapped Comscore Markets onto counties or states, as these are Comscore's proprietary geographic market definitions. However, we were able to crosswalk Comscore's 210 Markets to Nielsen's 210 Designated Market Areas (DMAs), and the Kilts Center for Marketing at the University of Chicago Booth School of Business provided us with a crosswalk of Nielsen's DMAs to counties. The primary difference between Comscore Markets and DMAs that is relevant for our analysis is that the DMAs precisely follow county boundaries, while Comscore Markets follow a ZIP code-based framework.

We first map DMAs into Comscore Markets. Then, using county population data from the Census's estimates for 2024, we calculate the total population for each DMA and the share of that population coming from each surrounding state. We assign each DMA a majority state and record the majority state's share of the DMAs' total population.

A.3 Winsorization

We winsorize the raw session-level data from both the desktop and mobile samples at the 95th percentile of session length for each month in our sample. Winsorization reduces a significant amount of noise and ensures that results are not driven by a few machines with unreasonable session durations in a given week. Then, we aggregate the winsorized session data to the machine-week-site level. We create a balanced panel of mobile and desktop machines for every week from January 1, 2022 to December 31, 2024. If a machine is not observed with any activity in a given week, we record the duration on every focal site as missing. If a machine is observed with any activity in a given week, but does not visit one of the focal sites, the duration on this site is recorded as zero. In the end, each observation reflects the weekly amount of time spent on a focal website or category of websites for each machine in the sample when that machine is active.

Table A.1: Pornhub Shutdown Dates Relative to Age Verification Law Effective Dates by State

	State	Shutdown Date	Law Effective Date
Treated States used in Analysis	Louisiana	N/A	01-01-2023
	Utah	05-01-2023	05-03-2023
	Mississippi	06-29-2023	07-01-2023
	Virginia	06-29-2023	07-01-2023
	Arkansas	07-28-2023	07-31-2023
	Montana	12-28-2023	01-01-2024
	North Carolina	12-28-2023	01-01-2024
	Texas	03-14-2024	09-19-2023
	Kansas	06-28-2024	07-01-2024
	Indiana	06-28-2024	08-16-2024
	Idaho	06-28-2024	07-01-2024
	Kentucky	07-10-2024	07-15-2024
	Nebraska	07-15-2024	07-18-2024
	Alabama	09-26-2024	10-01-2024
	Oklahoma	11-01-2024	11-01-2024
States Treated after Analysis Period	Florida	01-01-2025	01-01-2025
	South Carolina	01-01-2025	01-01-2025
	Tennessee	01-14-2025	01-14-2025
	Wyoming	06-30-2025	07-01-2025
	South Dakota	06-30-2025	07-01-2025
	North Dakota	06-30-2025	08-01-2025
	Georgia	07-01-2025	07-01-2025
	Arizona	09-26-2025	09-26-2025
	Ohio	N/A	09-30-2025
	Missouri	12-01-2025	11-30-2025

Notes: We list all U.S. states that have enacted age verification laws as of May 2026, along with Pornhub's shutdown date and each law's effective date. Our analysis includes all states above the dashed line: those with shutdown dates on or before November 1, 2024. Law effective dates are sourced from state legislature websites and LegiScan. The Age Verification Providers Association maintains a comprehensive list of sources for each law: https://avpassociation.com/us-state-age-verification-laws-for-adult-content/. Pornhub shutdown dates were identified through contemporaneous news coverage at the time of each state's effective date. Louisiana and Ohio have "N/A" for the shutdown dates because in Louisiana, Pornhub relies on a third-party age verification provider, and Pornhub has argued that Ohio's law was written in a manner that does not require the platform to comply. Texas and Indiana have larger gaps between law effective dates and shutdowns due to legal challenges. In Texas, the Free Speech Coalition challenged the law, and on September 1, 2023, a federal judge in Austin granted a temporary injunction to block enforcement. The injunction was lifted in November 2023, and on February 26, 2024, Texas Attorney General Ken Paxton sued Pornhub's parent company Aylo for noncompliance with the law. Pornhub blocked access to its site and affiliated sites in Texas on March 14, 2024. Indiana's law was originally scheduled to take effect on July 1, 2024. On June 28, 2024 a federal district court issued a preliminary injunction, blocking enforcement. Pornhub had already blocked access in Indiana one day before the injunction on June 27, 2024, and we found no evidence that access was reinstated during the injunction period.

Table A.2: Pooled Regression Results From Sample With Less Error in Geography Measurement

	Pornhub (1)	xVideos (2)	XNXX (3)	Other XXX (4)	All XXX (5)
β_{pre}	-0.0662**	-0.0124	0.0386	-0.0458	-0.0859
	(0.0306)	(0.0344)	(0.0390)	(0.0967)	(0.1287)
β_{ST}	-0.6293***	0.1942***	0.1918***	0.0745	-0.1688*
	(0.0469)	(0.0418)	(0.0387)	(0.0940)	(0.0994)
β_{LT}	-0.8161***	0.3176***	0.2014***	-0.1881**	-0.4852***
	(0.0592)	(0.0825)	(0.0551)	(0.0749)	(0.1207)
Baseline mean	1.3522	0.9965	0.7811	3.1865	6.3163
N (obs)	30,004,950	30,004,950	30,004,950	30,004,950	30,004,950
N (machines)	489,023	489,023	489,023	489,023	489,023

Notes: For this table, we report estimates from the pooled regression in equation 2, restricting to the sample to machines in Comscore Markets where the majority state in that market comprises more than 80% of the population. We lose about 12% of the original analysis sample after making this restriction. All regressions include cohort × machine and cohort × calendar week fixed effects. The dependent variable is (winsorized) weekly minutes per machine; the baseline mean row reports these averages for treated machines in week −1. N (obs) is the number of machine × cohort × calendar-week observations in the stacked panel. N (machines) is the number of unique machines in the stacked panel. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: Long-term Effects of Restrictions on Top 25 Adult Websites

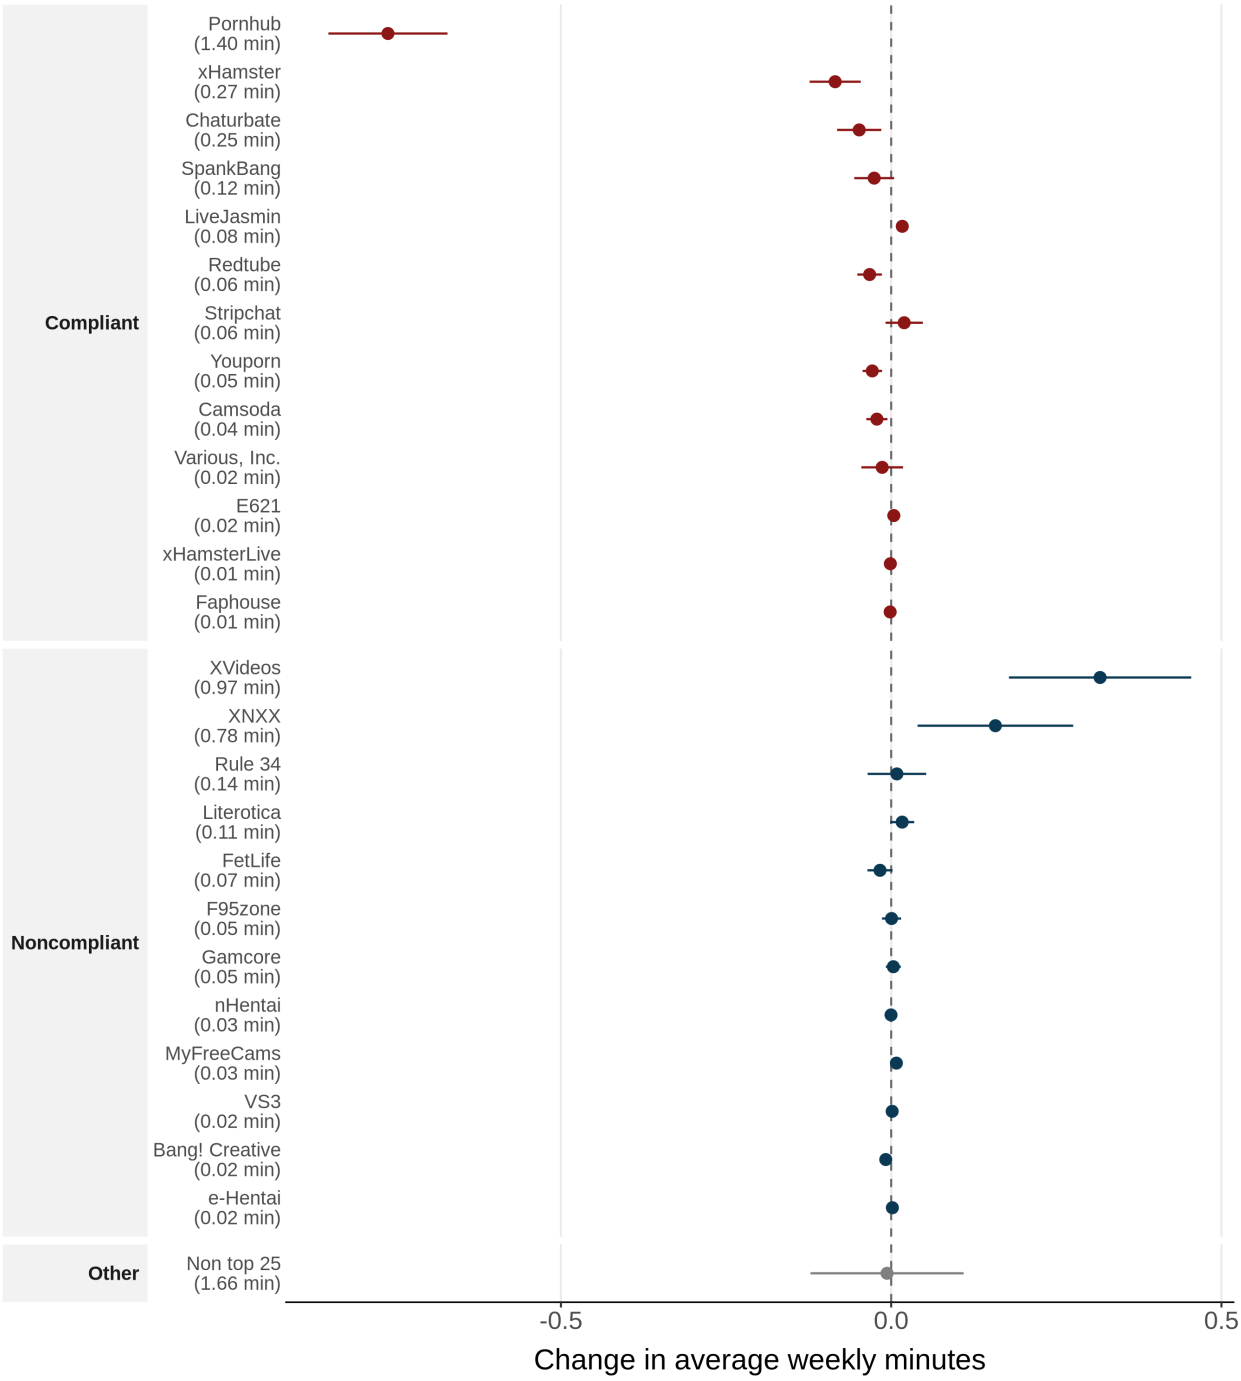

Notes: Each point represents the long-term estimate of how the regulation affected time spent on a top 25 adult site: $\beta_{LongTerm}$ from equation 2 estimated separately for each website. Estimates for compliant sites are grouped together and shown in red, noncompliant sites are group together and shown in blue. At the bottom, we show the pooled estimate for all other sites Comscore categorizes as "XXX Adult" outside of the top 25 ("Non top 25"). Error bars show 95% confidence intervals. Standard errors are clustered at the state level. Baseline consumption for each site is reported in parentheses: average weekly minutes spent on that site in $t \in -4, -1$ for treated machines. We coded compliance status by first looking for evidence from contemporary news reports, and then by manually checking whether content was accessible in states with laws in May of 2026.

Figure A.2: Long-term Effects of Restrictions on Selected Non-adult Websites

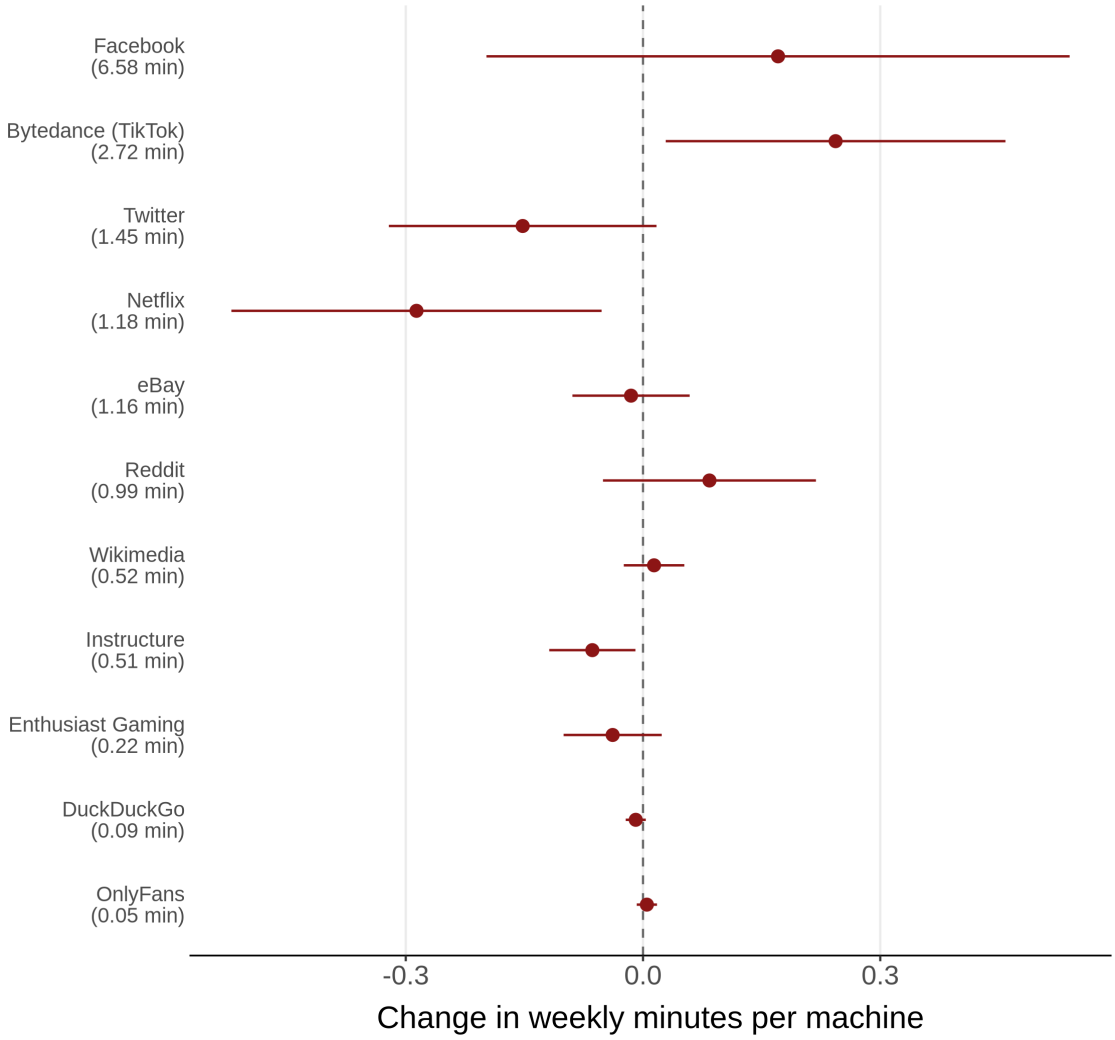

Notes: Each point represents the long-term estimate of how the regulation affected time spent on the selected non-adult sites or group of sites: $\beta_{LongTerm}$ from equation 2 estimated separately for each site. Error bars show 95% confidence intervals. Standard errors are clustered at the state level. We group sites by their parent company, so "Facebook" includes time spent on Instagram too (since both are part of Meta). "Twitter" also includes X.com after Twitter's name change in July 2023. Wikimedia is the parent company for Wikipedia and other Wiki sites. Instructure is the parent company for Canvas and Educational Tools. Enthusiast Gaming is a major digital media and entertainment company that operates networks of gaming communities like addictinggames.com. DuckDuckGo is a popular privacy-focused search engine. Baseline consumption of each site or group of sites is reported in parentheses: average weekly minutes in weeks $t \in [-4, -1]$ for treated machines.

Figure A.3: Heterogeneity in Circumvention

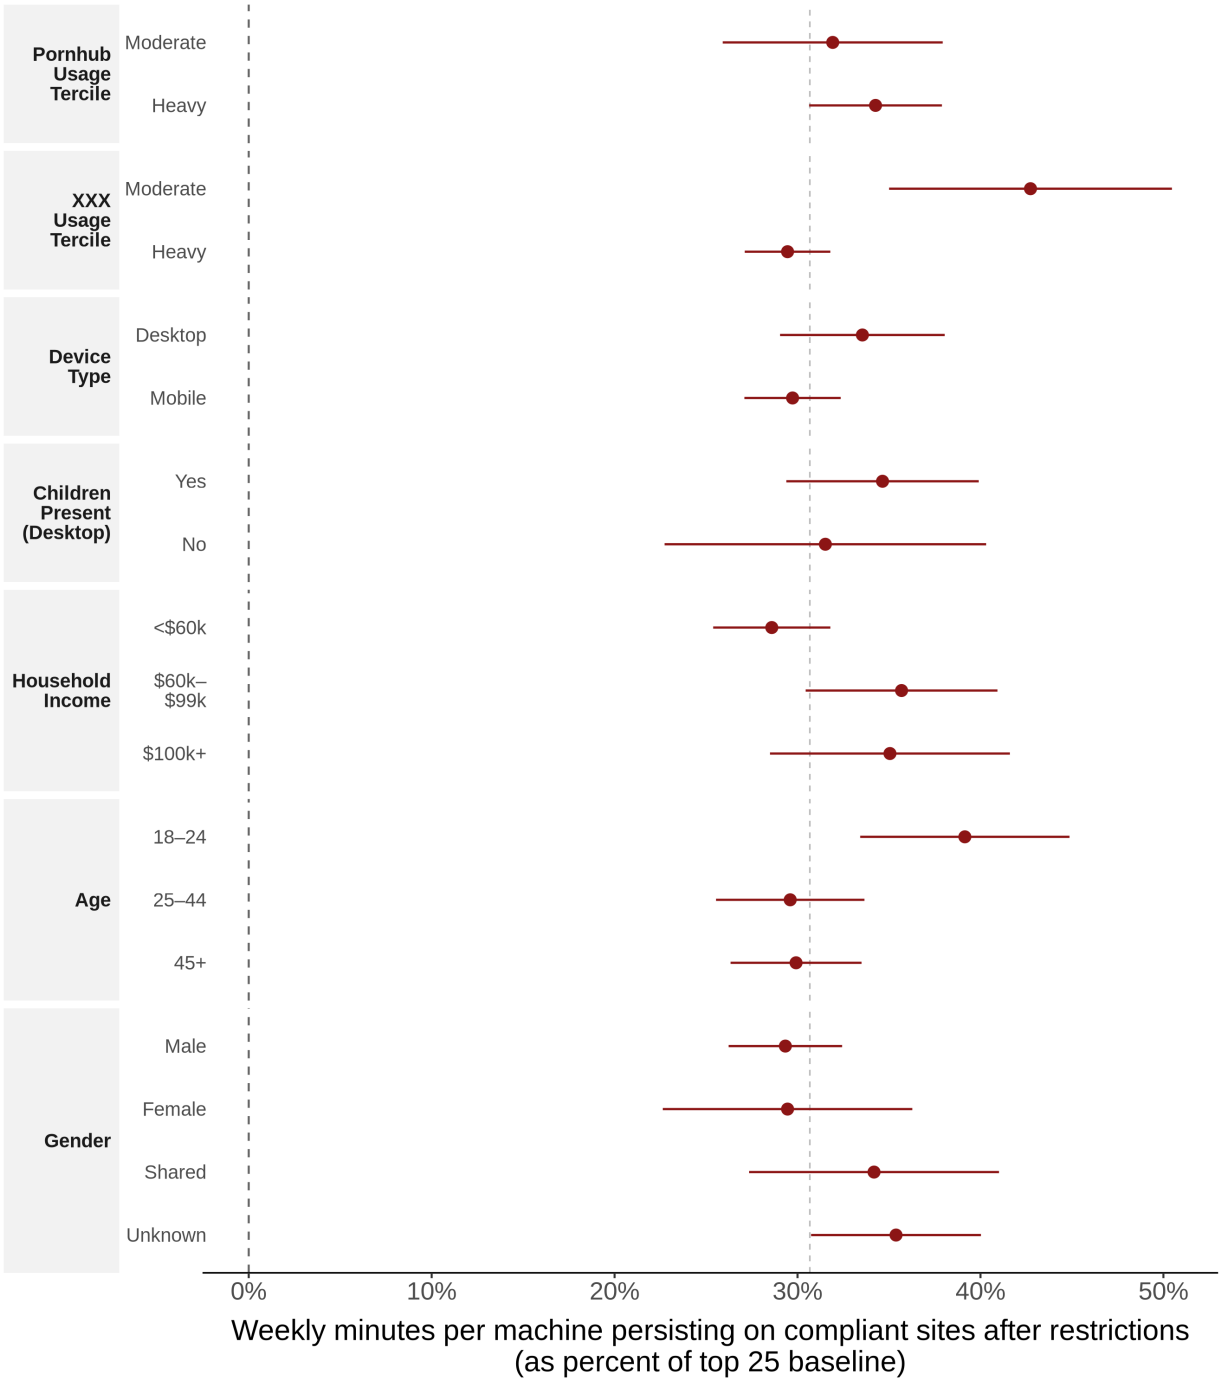

Notes: Each point represents a subgroup-level estimate of how the regulation affected remaining time spent on compliant sites within the top 25. We report all estimates in units of "percent of baseline consumption" on all top-25 sites. Specifically, the numerator is the estimated post-treatment consumption level on compliant top-25 sites: for each compliant site (Pornhub and other compliant top-25 sites), we add its long-term treatment effect ($\beta_{LongTerm}$ from equation 2) to its pre-period baseline ($t \in [-4, -1]$) and sum across sites. The denominator is the subgroup's baseline mean: average weekly minutes on all top-25 sites in $t \in [-4, -1]$ for treated machines. Error bars show 95% confidence intervals, constructed by taking the square root of the sum of squared standard errors from the two estimates that make up the numerator and dividing by the baseline consumption on top-25 sites. Standard errors are clustered at the state level. The lighter dashed line marks the analogous circumvention estimate for the full analysis sample. Pornhub and XXX usage terciles are defined among machines with positive average weekly minutes on Pornhub or any XXX site, respectively, during the pre-period (weeks −16 through −1). Non Visitor and Light tercile subgroups are excluded from this figure; when baseline consumption is zero or very small reporting effects in percent units is uninformative. When reporting on the "children present" split, we restrict both subsamples to desktop machines only. See the notes for Table 1 for other details on demographic variables.

Figure A.4: Heterogeneity in Substitution to Noncompliant Sites

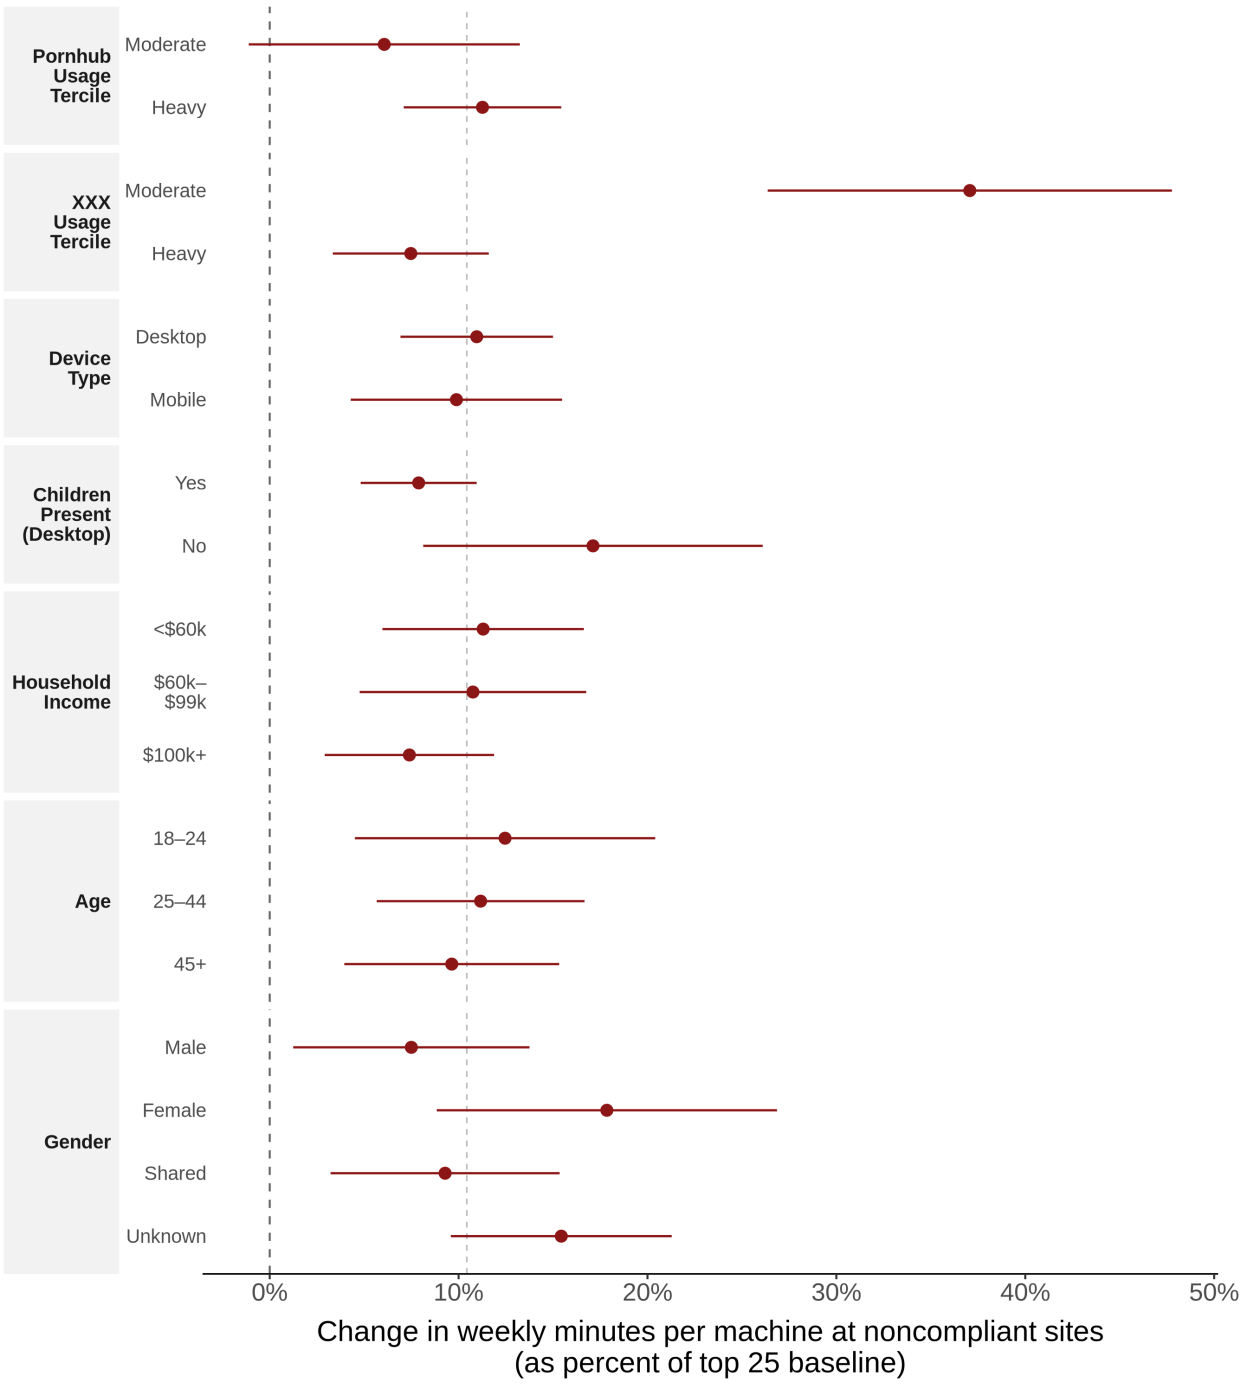

Notes: Each point represents a subgroup-level estimate of how the regulation shifted consumption toward noncompliant sites within the top 25. We report all estimates in units of "percent of baseline consumption" on all top-25 sites. Specifically, the numerator for each estimate is the sum of the long-term treatment effects ($\beta_{LongTerm}$ from equation 2) across XVideos, XNXX, and other noncompliant top-25 sites. The denominator is the subgroup's baseline mean: average weekly minutes on all top-25 sites in $t \in [-4, -1]$ for treated machines. Error bars show 95% confidence intervals, constructed by taking the square root of the sum of squared standard errors from the three estimates that make up the numerator and dividing by the baseline consumption on top-25 sites. Standard errors are clustered at the state level. The lighter dashed line marks the analogous substitution estimate for the full analysis sample. Pornhub and XXX usage terciles are defined among machines with positive average weekly minutes on Pornhub or any XXX site, respectively, during the pre-period (weeks −16 through −1). Non Visitor and Light tercile subgroups are excluded from this figure; when baseline consumption is zero or very small reporting effects in percent units is uninformative. When reporting on the "children present" split, we restrict both subsamples to desktop machines only. See the notes for Table 1 for other details on demographic variables.