



Youth Internet Bans, Circumvention, and Substitution: Policy Briefing

Sam Glover, May 2026

Executive summary

- The best existing evidence suggests that youth internet bans are unreliable tools for ensuring that under-16s do not access harmful online content. The evidence that we have from a variety of countries and settings suggests that attempts to exclude children from online activities are frequently met with either circumvention or substitution into other online activities.
- The best direct evidence on circumvention and the weak effect of bans comes from South Korean and Chinese internet and/or gaming bans, evidence from the United States, and preliminary evidence from the Australian social media ban. Research consistently shows that circumvention of bans is common, with common methods of circumvention including VPN usage, borrowing unbanned accounts, adversarial methods to bypass age-assurance checks, and retaining use of accounts owned prior to the implementation of the ban.
- The preliminary evidence from Australia is especially relevant given that it is the only example of a youth social media ban for under-16s globally. While early evidence should not be taken to be definitive, the research that does exist indicates that circumvention is relatively straightforward, and that many young people have retained social media accounts from prior to the implementation of the ban, although there is also some evidence suggesting a modest reduction of usage of banned platforms.
- There is also a concern that exclusion from online spaces may result in young people substituting into other online activities that do not comply with the ban. Evidence suggests that this is highly likely to be the case; studies on online pornography bans show that searches for non-compliant rival sites increase dramatically after the implementation of pornography bans.
- The central policy take-away from this rapid research synthesis is that youth internet bans should not be taken to be reliable policy instruments for preventing children from accessing harmful content. This is true for two reasons: firstly, because circumvention is sufficiently common that many children simply will not be kept off the platforms that are trying to exclude them. Secondly, because many children will substitute their use of compliant platforms with similar platforms that do not comply with the ban.

Review question and scope

The policy briefing addresses three questions:

- | | |
|----|---|
| 1. | What happens when children and adolescents are banned or restricted from particular online services, content categories, or times of use? |
| 2. | Do young people circumvent these restrictions, and through what routes? |
| 3. | Among young people who do not circumvent restrictions, is there evidence of substitution onto other platforms? |

The term “ban” is used broadly. Included evidence covers social-media age restrictions, gaming curfews and playtime mandates, pornography age-verification laws, platform age gates, home or ISP filters, parental controls, school phone bans, and related age-assurance systems. The reasoning here is that essentially no completed peer-reviewed studies evaluate broad social-media bans for under-16s, so other evidence is needed to give the broadest overview of the likely effects and efficacy of youth social media bans.

This report is a rapid synthesis of available evidence on youth internet bans, age assurance, substitution, and circumvention. Sources were identified through iterative database and web searching, backward and forward source chasing from relevant policy briefings, review of regulator and parliamentary materials, and targeted searches for Australian implementation evidence. Although we use a mixture of peer-reviewed evidence and preliminary evidence from non-peer-reviewed sources (including working papers), we prioritised evidence that had strong causal identification and/or was peer-reviewed; we used non-peer-reviewed evidence where stronger evidence was unavailable. In the case of the Australian under-16 ban, no peer-reviewed evidence with strong causal identification existed. However, we believe that because this is the only example of an under-16 social media ban globally, it would have been ill-advised to ignore the evidence that does exist.

Sources were collected in a spreadsheet, and relevant information such as citation, year, policy area, design, and strength of causal identification was included. With regard to strength of causal identification, Randomised Controlled Trials and studies with strong quasi-experimental designs (such as difference-in-differences designs) were noted as being particularly high-quality. Weaker studies included those that relied primarily on observational and correlational data that made them more susceptible to omitted variable bias or other problems in establishing causal relationships.

Findings

Implemented restrictions show weak and incomplete exclusion effects, and may lead to substitution

Some of the strongest evidence on internet bans of various kinds comes from South Korea and China, as both countries have implemented bans to prevent young people from accessing certain online content. In South Korea, a “Cinderella Law”, designed to prevent under-16s from accessing gaming services between midnight and 6am (with the aims of reducing disruption to sleep and preventing gaming addiction), worked by requiring gaming companies to verify users’ names and ages when they signed up for online gaming services. Choi et al. (2018) use a difference-in-differences design with more than 240,000 observations from a web-based survey to determine how effective the restriction was. They find that the law resulted in a reduction of daily internet usage of only 3.6 minutes, and no detectable impact on the policy’s aims of reducing sleep disruption and/or internet addiction. Lee et al. (2017) use a regression-based study to examine the same law, finding that the ban made extremely small and insignificant differences to the outcomes it targeted: it showed only a modest decrease in the probability of internet addiction (a 0.7pp fall), an increase in the amount of time spent sleeping of 1.5 minutes, and a statistically significant increase in the probability of being in a high-use internet group (by 1.6 percentage points). Neither study was able to track individual behaviour, only looking at behaviour in the aggregate, and neither study was able to determine whether young people substituted online gaming activity with other online activities.

China introduced a similar law targeting online gaming in 2021, wherein under-18s could only access online gaming services for around 3 hours per week at certain times. Gaming companies were expected to enforce this through verifying the identities of people signing up for and using online gaming services. Zhou et al. (2024) surveyed a subset of the population who identified as heavy-gamers, and found that while many of them had complied with the law, a minority were able to circumvent the ban through acquiring alternative identities or renting accounts that were not restricted from accessing the services, and more notably, the majority of respondents said that they had substituted their gaming behaviour with short-form video.

Zendle et al. (2023) look at a 2019 version of the same Chinese gaming ban, and conduct before-and-after and within-participant analyses, finding that there is no evidence that the gaming restrictions actually reduced heavy play, and in fact their point estimate suggests that it may have increased heavy play, although this finding was below their pre-specified threshold for practical significance. Sensitivity analyses found a similar absence of reduction after the stricter 2021 adjustment. It should be noted that the study examines the gaming profiles of accounts on online gaming services generally, rather than accounts confirmed to belong to those targeted by the restrictions, so this evidence should be taken to mean that the law did not reduce the use of online gaming services among users as a whole, rather than young users specifically.

Other evidence on the efficacy of age verification (and the extent to which users engage in “substitution”, replacing their usage of a compliant site with a non-compliant site) comes from the United States. The US has introduced legislation mandating that sites providing adult content require age-verification before they allow access. One study, Lang et al. (2026), uses a pre-registered quasi-experimental design exploiting the fact that these laws underwent a staggered roll-out. They found that several months after implementation, searches for compliant adult sites fell significantly (a 51% reduction), but searches for a non-compliant rival adult-content site increased by 48.1%, and searches for VPN providers (as a means of circumvention) increased by 23.6%. It should be noted that absolute changes cannot be inferred from the available data, so the study authors are only able to report relative changes. Even given

this caveat, this study should be taken as meaningful evidence that age-verification laws result in a meaningful increase in users accessing non-compliant alternatives to restricted sites, and a meaningful increase in VPN usage to bypass restrictions.

Other studies add to the concern that substitution is a real issue regarding these bans: a working paper by Agarwal et al. (2026) uses a similar staggered design, employing a difference-in-differences approach to study the extent to which substitution occurs after the implementation of adult-content laws that require age-verification. They find evidence that users substitute into non-compliant sites: specifically, they find that while compliant sites lose significant traffic, the traffic for non-compliant sites increases by between 35 and 40%, suggesting that there is a meaningful degree of substitution rather than it simply being a case of overall demand reduction. It should be noted, however, that the substitution that occurs primarily seems to be to the benefit of large non-compliant sites, and that there was no similar large-scale substitution into fringe websites that did not adopt age-verification requirements. The authors do note that overall demand did fall somewhat overall, meaning that the age-verification laws may not have resulted in total substitution: the study finds that the overall reduction in demand for adult-content sites was 23%.

In a study on social media usage and wellbeing using a randomised controlled trial, Collis and Eggers (2022) find that university students who were restricted to ten minutes per day on both Facebook and other social media platforms do not end up spending any less time overall on their devices than participants in the control group, and instead mostly replace their usage of social media platforms with instant messaging platforms such as WhatsApp. These studies focus on adult university students rather than people under the age of 18, so they cannot be used to make broad inferences about the impact of youth internet bans in particular, but they do suggest that restricting platforms can simply result in users of those platforms replacing their usage with similar services that have not been restricted.

A study by Madio et al. (2026) provides further quasi-experimental evidence on substitution. In December 2020, MindGeek (the owner of Pornhub) deleted approximately 80% of its video library overnight following pressure from both the media and payment processors. Using a difference-in-differences design across eight countries, the authors find that MindGeek lost 36 to 42% of visits within a month. The evidence suggests that users substituted into other platforms that had not deleted content: mainstream rivals like XVideos and XNXX gained 9 to 23% over six months, and fringe sites saw traffic increase by 56 to 58%. Notably, the substitution from mainstream sites to fringe alternatives was not mostly driven by users directly navigating to fringe sites. Instead, the uptick in traffic was typically a result of using search engines and other aggregation services. As with the previous study, it should be noted that this was not specifically a youth restriction, and is instead intended to provide evidence generally that we see significant evidence of substitution when restrictions are implemented more broadly.

Another useful source of evidence is the preliminary data from the Australian youth social media ban. However, as the ban was only implemented relatively recently (in December 2025), this evidence should be taken as quite preliminary, and there is not much evidence that has strong causal identification techniques. Nonetheless, because this is the only evidence we have of the direct impacts of a youth social media ban, it seems particularly pertinent to the possible impacts of a UK youth social media ban. In Australia, before the implementation of the ban, eSafety reported widespread underage use of social and messaging services: 84% of 8 to 12-year-olds had used at least one social-media or messaging service since the start of 2024 despite a limit to over 13s; 40% of users had their own account; 80% of account-holders had help setting up accounts; and only 13% had an account shut down for being underage (eSafety Commissioner, 2025). Qualitative research commissioned by eSafety found that children and parents anticipated workarounds, including false ages, use of a parent's face, and migration to similar platforms (eSafety Commissioner, 2026a). A large ABC Behind the News non-probability

survey of 17,144 respondents aged 9 to 15 found that almost 80% used social media in some capacity, 75% of social-media users had their own account, 75% of social-media users did not plan to stop after the ban, and only 6% thought the ban would work (ABC Behind the News, 2025).

The initial evidence from Australia suggests that the ban may have had a real but incomplete effect on the number of children who had access to social media accounts. A survey commissioned by the eSafety Commissioner in late January and early February 2026 asked around 900 parents and carers about their children's (aged 8 to 15) use of restricted platforms. The proportion of children who parents and carers reported having their own social media account fell by 18.4 percentage points, from 49.7% to 31.3%. On most platforms, the majority of children who already had access to a social media account were able to retain access, with the proportion able to retain access ranging from 48.5% on YouTube to 69.4% on Snapchat. However, it should be noted that this evidence comes from data reported by surveys of parents and carers, meaning they may understate the true extent of the circumvention by under-16s if children were not willing to share with their parents that they were still able to access these platforms (or if the parents/carers were otherwise unaware of their children's online activity).

Of the parents who had a child who was able to keep access to their account, 66.8% said the platform had not yet asked for age verification, 23.7% said age verification got the child's age wrong, 23.3% reported another workaround such as makeup or someone else's ID, 10.3% said the child opened a new account after deactivation, 7.6% reported parent/carer help, and 6.7% reported VPN use (eSafety Commissioner, 2026c). The official compliance report update has evidence that points in the same direction: large numbers of children were able to keep access to their old accounts, create new accounts, or bypass the systems designed to keep them off the platforms. Platforms were also found to keep prompting children to attempt age-verification, which allowed them to repeat their attempts until allowed onto the platform (eSafety Commissioner, 2026b).

A working paper by Bursztyn et al. (2026) also examines the Australian youth social media ban, looking into the extent to which teenage users complied with the ban. Using surveys of Australian 14- and 15-year-olds four months after the implementation of the ban, they find that around 27% of banned users actually complied with the restrictions, and note that most banned teenagers regarded bypassing the restrictions as easy. They also survey teenagers asking them how many other young users would have to comply with the ban before they would stop themselves, finding that it would require around two-thirds of their peers to comply before they also decided to do so; given that the majority of users were not complying with the ban, it seems unlikely that this threshold will be passed in the near future. Further polling commissioned by the Molly Rose Foundation found similar results, with 61% of respondents between 12 and 15 years old who previously had social media accounts reporting that they retained access to at least one social media account that they had had prior to the implementation of the ban (Molly Rose Foundation and YouthInsight, 2026).

Data on the extent to which younger children are able to circumvent platform-wide bans exists in Europe, with EU Kids Online finding that around 38% of 9 to 12-year-olds had access to at least one social media platform despite the platforms limiting usage to those who are thirteen or older, and approximately 20% of younger children had access to a Facebook profile (Livingstone et al., 2011, 2013). In a similar study by Net Children Go Mobile, they found that around 42% of children below the age threshold of thirteen had a Facebook profile. Regression models suggested that the main variables that predicted underage usage of social media sites were social motivations, communication skills, peer mediation, age, and online disinhibition rather than a clear problematic-use profile (Barbovschi et al., 2015).

There is also further evidence from the United States on youth internet restrictions specifically, and one clear finding is that circumvention of age restrictions is often assisted by family members. Boyd et al. (2011) find that 36% of parents were aware that their children had accounts even though they were below the age limit to create those accounts, and that of those who had accounts, the majority (68%) had been assisted by their parents in creating those accounts. Pew found that a sizeable minority of children (39%) reported lying about their age in order to gain access to a restricted site or account, while also finding that around a quarter of children lied about other information (such as their name or location), although it should be noted that this was primarily in the context of online privacy rather than necessarily to circumvent restrictions (Madden et al., 2013). Evidence from the UK from the Advertising Standards Authority and Ofcom commentary similarly suggests that children frequently use false ages and are often treated as adults by online platforms, although the sample size in the ASA study was relatively small (n = 24) (Advertising Standards Authority, 2013; Ofcom, 2022).

Discussion and Policy Implications

The evidence generally supports the conclusion that youth internet bans and age-verification methods generally are not particularly effective at keeping children (or other users) off restricted platforms; circumvention is quite widespread and generally relatively straightforward, substitution into non-compliant sites also appears to be relatively common, and workarounds such as accessing VPNs, keeping access through accounts owned prior to the restriction, and/or gaining access through assistance from family, friends, or other older account owners, all seem to be reasonably effective ways of circumventing restrictions.

The strongest lesson for policy-makers is that bans should be assumed to come with a large degree of both circumvention and substitution; many young people will simply get around the bans using the aforementioned workarounds, and many of those young people who are either unable or uninterested in circumventing the bans are likely to substitute to other online activity (at least some of which is likely to be enabled by non-compliant sites, rather than through moving to “safer” online activity). While we don’t have particularly strong evidence on the question of who exactly is more likely to be able to engage in circumvention, it seems reasonable to assume that technically sophisticated and/or determined young people are both more likely to have the prowess to circumvent restrictions, and are possibly also more likely to be at risk of harms from the internet and social media.

For a youth social media ban specifically, policymakers should assume that technologically able and capable young people will almost certainly be able to circumvent bans, but even a large proportion of ordinary young people are likely not to be wholly restricted from online social media sites. Much of the evidence we have presented in this briefing should also be taken as being likely to understate the true extent to which young people have been able to access restricted sites after the implementation of such bans; data from parents and from young people are both liable to result in understatement for separate reasons: the former because parents/carers may not have accurate information about whether their children continue to be able to access restricted sites, and the latter because children may not be perfectly honest about whether they are continuing to access restricted sites. Given this, it seems plausible that even the observational studies that suggest a modest reduction in usage are overstating the case, and in fact the reduction may be minimal or even entirely nonexistent.

Some of the rhetoric around the proposed social media ban in the United Kingdom has implied that the ban is likely to be relatively effective. For instance, the Commons Education Committee presented a cross-party report that argued that an outright ban would “arrest the current harm”, and that a legal boundary would give children “greater protection”. These are assumptions that need to be tested rather than asserted, and the best evidence seems to suggest that in fact many

children would not be prevented from accessing restricted websites, while other evidence suggests that a meaningful proportion of young people would simply begin to use non-compliant platforms. Apparent falls in traffic to social media sites cannot be said to reliably inform policy-makers about the actual impact of the ban: whether through using VPNs, or through substitution into other harmful content, the evidence needs to be taken in its totality rather than reviewed selectively. Ultimately, these questions are hard to answer without strong experimental and quasi-experimental evidence, and policymakers should bear this in mind when evaluating the efficacy of such bans, rather than relying solely on self-reports, from observational data, or from reported changes in traffic from certain localities (or from users of certain ages).

Substitution also risks increasing the harms to children rather than reducing them. The quasi-experimental evidence presented by Lang et al. (2026), using a pre-registered design, found that age-verification laws decreased the proportion of users accessing compliant adult-content sites, but significantly and meaningfully increased the proportion of users accessing non-compliant adult-content sites. In the case of social media sites, it may be the case that such unregulated and non-compliant sites pose more risk to children (especially if they are not adhering to other regulations) than are compliant rivals.

Overall, the best current evidence suggests that youth internet bans are weak exclusion tools: young people (and other excluded users in studies that look at internet bans more generally) are largely able to circumvent such bans, the time spent on restricted platforms seems to only fall by very small amounts (and sometimes does not fall at all). Early Australian evidence indicates that the majority of young people were able to retain access to social accounts that they had prior to the ban, and much of this evidence relies on self-reported data that may understate the true extent of circumvention. Expectations around a proposed social media ban should be modest: most children should be expected to retain access to restricted sites, some children will substitute into non-compliant rivals, and the exclusion effect may even end up being non-existent as knowledge of workarounds proliferates among young people. Future evaluations should not measure directly whether accounts are merely removed or whether parents/carers/young people report that access has reduced, but also include quasi-experimental methods to determine the extent to which usage has actually fallen, and examine the extent to which substitution has occurred.

References

- ABC Behind the News. (2025). What 17,000-plus young people think about the social media ban. <https://www.abc.net.au/news/2025-12-02/what-young-people-think-about-social-media-ban-btn-survey/106058976>
- Advertising Standards Authority. (2013). ASA research shows children are registering on social media under false ages. <https://www.asa.org.uk/news/asa-research-shows-children-are-registering-on-social-media-under-false-ages.html>
- Agarwal, S., Ananthakrishnan, U. M., Madio, L., Mitchell, M. F., Quinn, M., & Reggiani, C. (2026). Are you 18? Age verification and adult-only consumption in the US [SSRN working paper]. <https://doi.org/10.2139/ssrn.6619018>
- Barbovschi, M., Machackova, H., & Olafsson, K. (2015). Underage use of social network sites: It's about friends. *Cyberpsychology, Behavior, and Social Networking*, 18(6), 328–332. <https://doi.org/10.1089/cyber.2014.0553>
- Boyd, D., Hargittai, E., Schultz, J., & Palfrey, J. (2011). Why parents help their children lie to Facebook about age: Unintended consequences of the Children's Online Privacy Protection Act. *First Monday*, 16(11). <https://doi.org/10.5210/fm.v16i11.3850>
- Burszty, L., Duckworth, A., Jiménez-Durán, R., Leonard, A., Milojević, F., Roth, C., & Sunstein, C. R. (2026). Why bans fail: Tipping points and Australia's social media ban (Becker Friedman Institute for Economics Working Paper No. 2026-57). University of Chicago.
- Choi, J., Cho, H., Lee, S., Kim, J., & Park, E.-C. (2018). Effect of the online game shutdown policy on internet use, internet addiction, and sleeping hours in Korean adolescents. *Journal of Adolescent Health*, 62(5), 548–555.

Collis, A., & Eggers, F. (2022). Effects of restricting social media usage on wellbeing and performance: A randomized control trial among students. *PLOS ONE*, 17(8), e0272416. <https://doi.org/10.1371/journal.pone.0272416>

eSafety Commissioner. (2025). Children and social media. Australian Government. <https://www.esafety.gov.au/research/children-and-social-media>

eSafety Commissioner. (2026a). Attitudes of children and parents to social media age restrictions ahead of implementation. Australian Government. <https://www.esafety.gov.au/sites/default/files/2026-03/Attitudes-of-children-and-parents-social-media-age-restrictions-Mar-2026.pdf>

eSafety Commissioner. (2026b). Social media minimum age: Compliance update. Australian Government. <https://www.esafety.gov.au/sites/default/files/2026-03/SocialMediaMinimumAgeComplianceUpdateMarch2026.pdf?v=1774905032806>

eSafety Commissioner. (2026c). Under the new age restrictions: Early insights from Australian parents. Australian Government. <https://www.esafety.gov.au/sites/default/files/2026-03/Under-the-new-age-restrictions-Early-insights-from-Australian-parents-March2026.pdf?v=1775024586194>

Lang, D., Listyg, B., Ross, B. V., Musquera, A. V., & Sanderson, Z. (2026). Age verification and public adaptation: A pre-registered synthetic control multiverse. *Journal of Law & Empirical Analysis*. <https://doi.org/10.1177/2755323X251415424>

Lee, C., Kim, H., & Hong, A. (2017). Ex-post evaluation of illegalizing juvenile online game after midnight: A case of shutdown policy in South Korea. *Telematics and Informatics*, 34(8), 1597–1606. <https://doi.org/10.1016/j.tele.2017.07.006>

Livingstone, S., Olafsson, K., & Staksrud, E. (2011). Social networking, age and privacy. *EU Kids Online*. <https://eprints.lse.ac.uk/35849/>

Livingstone, S., Olafsson, K., & Staksrud, E. (2013). Risky social networking practices among underage users: Lessons for evidence-based policy. *Journal of Computer-Mediated Communication*, 18(3), 303–320. <https://doi.org/10.1111/jcc4.12012>

Madden, M., Lenhart, A., Cortesi, S., Gasser, U., Duggan, M., Smith, A., & Beaton, M. (2013). Teens, social media, and privacy. Pew Research Center. <https://www.pewresearch.org/internet/2013/05/21/teens-social-media-and-privacy/>

Madio, L., Mitchell, M., Quinn, M., & Reggiani, C. (2026). Asymmetric content moderation in search markets: The case of adult websites. SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5106235

Molly Rose Foundation & YouthInsight. (2026). Australia's social media ban: Is it working? Molly Rose Foundation.

Ofcom. (2022). How old is your child online? The risks of children using false ages on social media. <https://www.ofcom.org.uk/online-safety/protecting-children/risks-of-children-using-false-ages-on-social-media>

Zendle, D., Flick, C., Gordon-Petrovskaya, E., Ballou, N., Xiao, L. Y., & Drachen, A. (2023). No evidence that Chinese playtime mandates reduced heavy gaming in one segment of the video games industry. *Nature Human Behaviour*, 7, 1753–1766. <https://doi.org/10.1038/s41562-023-01669-8>

Zhou, X., Liao, M., Gorowska, M., Chen, X., & Li, Y. (2024). Compliance and alternative behaviors of heavy gamers in adolescents to Chinese online gaming restriction policy. *Journal of Behavioral Addictions*, 13(2), 687–692. <https://doi.org/10.1556/2006.2024.00021>