Addendum: technology companies, burden-shifting and safety at source

Core finding

The government's claim is that tech companies failed to keep children safe, so the state is stepping in. That is partly accurate. The Online Safety Act gives platforms duties to protect users from illegal content and children from harmful content, and Ofcom can enforce those duties. GOV.UK says the Act gives social media companies and search services legal duties to protect users from illegal content and content harmful to children. (GOV.UK)

But the proposed under-16 ban changes the centre of gravity. Instead of only forcing platforms to make their products safer, it also makes children's access the object being controlled. That means the policy can look like a punishment for tech companies while being experienced as a restriction on children and a verification demand on users.

The plain reading is:

Government claim	Practical reading
tech companies failed	largely true, especially on age enforcement, harmful content, algorithms and risky contact
government is taking power from tech giants	partly true legally, but operational power also moves to Ofcom, age-verification systems and platform compliance teams
power is going back to parents	parents get a shared rule, but they do not get direct control over platform design
children will be protected	some access and functions may be reduced, but harm can move if design incentives are unchanged
companies will be held responsible	companies face duties and penalties, but users may carry the day-to-day burden of proving age and losing access

What the government has done to rein in technology companies

The government and regulators have already created several layers of control over technology companies.

Measure	What it does	How strong it is	Main limitation
Online Safety Act 2023	makes in-scope services implement systems and processes to reduce illegal-content risks and protect children from harmful or age-inappropriate content	strong legal framework	it relies on Ofcom enforcement, platform compliance and risk assessments rather than removing all harmful design incentives
illegal-content duties	platforms must assess illegal harms and take steps to remove illegal material quickly and reduce priority criminal content appearing	enforceable since March 2025	illegal content is only part of the harm children experience
child-safety duties	services likely to be accessed by children must assess child risks, put protections in place and keep records updated	legally direct on providers	depends on how well platforms assess their own risks and how aggressively Ofcom enforces
age-assurance duties	services must use effective age checks in some contexts, especially pornography and harmful content	central to the regime	shifts privacy and proof burdens onto users
Ofcom codes and guidance	Ofcom has set out practical measures for companies, including safer feeds, strong age checks and protection from stranger contact	important implementation layer	codes do not automatically change the profit model behind attention-maximising design
Ofcom enforcement powers	companies can face enforcement action and, in serious cases, fines up to 10% of qualifying worldwide revenue or £18 million, whichever is greater	strong on paper	large platforms may absorb fines unless enforcement is fast, repeated and reputationally costly
enforcement actions	Ofcom has fined and investigated providers, including age-check failures, a suicide forum, Telegram and other services	shows the regime is active	enforcement is reactive and can lag behind harms

Children's Wellbeing and Schools Act 2026	requires government to introduce age or functionality restrictions for children using social media	adds a specific route to restrict access and features	turns safety into access control, not only product redesign
proposed under-16 social media ban	blocks specified social media platforms from offering services to under-16s	highly visible and politically clear	may be bypassed and may push children elsewhere
proposed feature restrictions	restricts livestreaming and stranger communication for under-16s across wider services, including gaming; switches these off by default for 16- and 17-year-olds	closer to safety by design than a blanket ban	still depends on definitions, enforcement and age checks

Ofcom's compliance guide says the Online Safety Act makes providers legally responsible for keeping UK users safe, especially children, even if the provider is based outside the UK. It also says serious failures can lead to enforcement and fines of up to 10% of qualifying worldwide revenue or £18 million, whichever is greater. (www.ofcom.org.uk)

Ofcom has also finalised more than 40 practical child-safety measures for sites and apps used by children, including measures for safer feeds, strong age checks, more help and control for children, and protection from harmful content. (www.ofcom.org.uk) Its June 2026 bulletin lists recent enforcement activity, including fines for failures to implement highly effective age assurance, a fine against 4chan, a fine against an online suicide forum provider, and an investigation into Telegram over child sexual abuse material concerns. (www.ofcom.org.uk)

There are also parallel controls outside the social media ban. The ICO's Children's Code contains 15 standards for online services likely to be accessed by children, including privacy by default, data minimisation, limits on detrimental use of children's data and risk-based age recognition. (Information Commissioner's Office) The CMA's digital markets regime can regulate the very largest technology firms with strategic market status, including conduct requirements and pro-competition interventions, although that regime is about market power rather than child safety specifically. (GOV.UK)

What the government has not fully done

The government has strengthened regulation, but it has not yet fully forced technology companies to change the core business model that creates much of the risk.

The key unresolved issue is not only harmful content. It is harmful design.

That includes:

Design issue	Why it matters
algorithmic feeds	they can intensify exposure to harmful, distressing or compulsive content
engagement incentives	platforms profit when users stay longer, scroll more and return often
data profiling	children can be shaped, targeted and nudged through data-driven systems
stranger contact	grooming, coercion, bullying and unwanted sexual contact often depend on contact pathways
livestreaming	real-time content is harder to moderate and can expose children quickly
infinite scroll and autoplay	these features increase time spent and reduce stopping cues
default settings	children may be exposed before they understand privacy, contact or recommendation controls

The government recognises some of this. Its announcement says algorithmic feeds can intensify exposure to dangerous, distressing or overly engaging material, and it is looking at overnight curfews and breaks in infinite scrolling. It also plans restrictions on livestreaming, stranger communication and sexualised AI chatbot functions. (GOV.UK)

The weakness is that the headline measure is still an access ban. The government is not yet saying: platforms must prove their recommender systems, advertising model, profiling practices and engagement features are safe before children can encounter them. It is saying: under-16s should be kept away from some platforms, and some functions should be restricted.

That is a different kind of accountability.

How the burden of safety is being shifted

The government says it is shifting responsibility to technology companies. In law, that is partly true. In practice, the burden is spread across the system.

Burden	Who carries it	How it shows up
legal compliance	platforms	risk assessments, age checks, record-keeping, safer systems, enforcement exposure

enforcement	Ofcom	defining effective age assurance, monitoring companies, investigating failures, issuing penalties
proof of age	users	adults and young people may have to prove they are old enough
privacy risk	users	facial checks, ID checks, age estimation or third-party verification may expose sensitive data
household policing	parents and carers	managing workarounds, arguments, shared accounts, device access and peer pressure
access loss	children	losing mainstream spaces for learning, creativity, identity, support and friendship
displacement risk	children, parents, Ofcom	harm may move to VPNs, adult accounts, private groups or less regulated services
offline replacement	schools, youth services, families	children need real alternatives for friendship, activity, support and belonging
market adaptation	smaller platforms	compliance costs may be easier for large platforms than small or non-profit alternatives
commercial opportunity	age-assurance providers	verification becomes a new market created by safety regulation

This is the central burden-shift:

The law tells platforms to keep children out or restrict features. But the practical checkpoint is the user's body, account, device, identity, face, parent or data trail.

The government fact sheet says many adults may not need checks if their account has been open more than 16 years, has a credit card attached, or is linked to an age-verified email address. But it also says some adults may have to do age verification, potentially as simply as a facial recognition check for over-18s. (GOV.UK) It also acknowledges that children may use VPNs or parents' accounts to get around checks, and says it has commissioned research into children's VPN use. (GOV.UK)

That means the system does not only regulate companies. It also asks users to become legible to the system.

Does the policy rely on consumers?

Formally, no. The legal duties fall mainly on providers and Ofcom. Companies must assess risks, comply with duties, enforce minimum ages, restrict functions and face penalties if they fail.

Ofcom says services likely to be accessed by children must carry out children's access assessments, children's risk assessments and put protections in place. (www.ofcom.org.uk)

Practically, yes. The policy relies heavily on consumers because the enforcement point is access.

It relies on:

Consumer role	What the user is expected to do
adults	prove they are over the relevant age if the platform cannot infer it
children	accept exclusion, move to permitted services, or be blocked
parents	stop children using adult accounts, shared devices or VPN workarounds
young people aged 16 and 17	accept default restrictions on livestreaming and stranger communication
families	replace lost online connection with offline alternatives
users with errors	challenge mistaken blocks or age-assurance failures

So the more precise conclusion is:

The policy places legal responsibility on companies, but it places practical friction on consumers.

A stronger safety-at-source policy would make the risky feature or business practice change before the child reaches it. This policy does some of that through restrictions on livestreaming and stranger contact. But the ban itself works mainly by stopping a category of users from entering.

Safety at source versus safety at the gate

This distinction is important for the report.

Approach	What it controls	Who bears the main burden	Example
safety at source	platform design	company	recommender systems cannot push self-harm or eating-disorder content to children

safety at source	contact architecture	company	unknown adults cannot message children by default
safety at source	data model	company	children are not profiled or nudged using harmful data practices
safety at the gate	access	user and platform	user must prove age before entering
safety at the gate	eligibility	Ofcom and platform	regulator decides what counts as adequate age assurance
safety at the gate	household compliance	parents and children	families manage workarounds and conflict

The government has chosen a mixed model, but its public story centres the gate. It says children under 16 should not be offered certain services. That is easier to announce than product-safety reform, but it is also easier for harm to route around.

The Commons Library briefing records the criticism clearly. When an MP asked about legally enforceable safety-by-design duties for children who “slip through the net”, Liz Kendall answered that it was not an either-or, because the Online Safety Act already requires companies to remove illegal content and harmful content for children. But the same briefing also records Ian Russell’s criticism that the ban was the politically easy route, and that government should have compelled tech companies to change their business models and tackle product safety at source.

That is the unresolved accountability gap.

The hidden inversion

The government says the ban takes power away from tech giants and puts it back in parents’ hands. Its own announcement uses that framing directly. (GOV.UK)

The inversion is that parents do not get control over the feed, the recommender system, the advertising model, the profiling logic, the contact architecture or the design decisions that made the platform harmful. They get a rule. Platforms still own the design. Ofcom owns the compliance judgement. Age-assurance vendors may own or process the checkpoint. Children and adults experience the restriction.

So power does not simply move from tech companies to families.

It moves from one corporate system to another compliance system, with families positioned as the moral justification.

Where technology companies may still benefit

The ban may hurt large platforms by reducing under-16 accounts. But large platforms may also be better placed than small competitors to absorb compliance costs, negotiate technical standards, integrate age assurance and manage legal teams.

That is why the competition angle matters. If age verification becomes expensive, identity-heavy and legally risky, small or non-profit services may struggle more than the major platforms. One uploaded Guardian opinion makes this argument directly, warning that identity and age verification could consolidate power among large platforms and third-party verification providers, while making it harder for smaller, privacy-centred competitors to operate.

This does not mean all age assurance is bad. It means age assurance is not neutral. It creates a market, a data flow and a new layer of gatekeeping.

The consumer-safety paradox

The policy is sold as protection from technology companies.

But if it relies on age verification, users may have to give technology companies or verification providers more information to be protected from technology companies.

That is the paradox.

The government says not every adult will need to verify. But the fact sheet still normalises the possibility that access to mainstream platforms may depend on account history, credit-card connection, age-verified email or facial recognition. (GOV.UK)

That creates a civic question:

Are we reducing platform power, or are we making more of everyday life pass through identity checks controlled by platforms and their vendors?

What a stronger company-accountability model would look like

A stronger model would not rely mainly on users proving age or children being excluded. It would require platforms to prove safety before risky features reach children.

That would include:

Company duty	What it would require

recommender safety	prove that feeds do not amplify self-harm, eating disorder, misogynistic, violent or extremist material to children
contact safety	block unknown adult contact with children by default
livestream safety	restrict child livestreaming and real-time contact unless safety can be proved
data minimisation	collect less data from children and reduce profiling
harmful-design audits	independently test addictive and compulsive design features
product pre-clearance	test high-risk features before launch, not after harms appear
redress	provide clear routes for wrongly blocked users or harmed children
displacement monitoring	track where children go when blocked and whether risk moves
transparency	publish evidence on risk, moderation, age checks, algorithmic exposure and appeals
competition protection	make compliance possible for small, safer and non-profit services

The Online Safety Act and Ofcom codes already contain parts of this direction. The problem is that the ban may become the visible substitute for deeper design reform.

Unignorable reading

This is a “burden of proof” problem.

The Unignorable principles warn about systems that substitute process for support, shifting proof and labour onto vulnerable people while institutions present the process as care. Applied here, the risk is that child safety becomes a series of checks, blocks, errors, appeals and household conflicts, while the most profitable parts of platform design remain intact.

The user-needs guidance is also relevant: a need should describe the real problem, not smuggle in a preferred solution. The real need is not “children need to be banned”. The real need is “children need safe, age-appropriate digital spaces that do not exploit their attention, data, vulnerability or social dependence”.

A ban may be one tool. It is not the need itself.

Colophon

This analysis was generated by *Unignorable*, a civic clarity assistant designed to surface hidden needs, assumptions, and consequences beneath civic and political language, clearly, plainly, and without jargon. Made by Happygeneralist, Ryan Haney.