

Centre for Media Transition



Hi there

Games of trouble



Welcome to the latest addition of the CMT newsletter!

This week, Derek investigates ACMA's latest decency breach finding against the Kyle and Jackie O show, and asks why we seem to be back at the beginning of the enforcement path.

Kieran examines the definitional hair-splitting underlying the public discussion about which platforms will fall under the social media ban, and queries whether the

regulatory approach is genuinely responsive to risk.

Finally, I look at Elon Musk's announcement of Grokipedia, his AI-generated alternative to Wikipedia, in the context of alarming drops in user traffic to both Wikipedia and news websites.



Michael Davis
CMT Research Fellow

Everybody's cross about KIIS



The media regulator, the Australian Communications and Media Authority (ACMA) announced at the start of this month that KIIS FM had breached the commercial TV code. ACMA [described](#) broadcasts in the second half of last year as “vulgar, sexually explicit and deeply offensive”.

There’s no arguing with that. If relaying the sound of staff urinating was not enough, the accompanying guessing game and explicit, sexualised language ensured the

segments did not conform with “generally accepted standards of decency”. As notoriously difficult as it is to nail these standards, it’s hard to see many people disagreeing with ACMA’s finding. Where people might disagree is on the forthcoming regulatory response.

When it published its investigation report and announced it was considering enforcement action, ACMA said it wanted to make sure the network, ARN, “takes full responsibility for the content broadcast on their stations”. This is, after all, the fifth breach finding this decade against this program for the same decency rule. And it comes after a [finding](#) in March that a broadcast of the same program segment, ‘Tradie vs Lady’, in the first half of last year “went beyond the bounds of decency expected by the community and was done so deliberately and provocatively”.

So we can expect a serious penalty, right? Maybe not.

There are serious limitations on ACMA’s enforcement powers when it comes to breaches of industry codes of practice, including the Commercial Radio Code of Practice. Apart from informal agreements about improved compliance and staff training, there are only two regulatory tools available to ACMA. Firstly, it can, as it has in the past, accept an enforceable undertaking in which a licensee promises, for example, to improve its compliance systems. An enforceable undertaking is a useful instrument because a breach of it empowers ACMA to seek a more meaningful penalty such as a civil penalty order from the Federal Court.

But ACMA can’t just make ARN or its licensee companies deliver an undertaking; an undertaking needs to be *offered*. ACMA’s own version of this regulatory action, also with serious consequences if it’s breached, is a remedial direction. But remedial directions are only available for breaches of licence conditions, program standards and other obligations imposed by the Broadcasting Services Act, not for breaches of industry codes.

This gets us to the second option: an additional licence condition which can, for example, require compliance with the existing code rules. That's not nothing, but it has its own limitation because a breach of an additional licence condition does not empower ACMA to seek a civil penalty – it can only lead to a remedial direction!

The problem with this convoluted path to significant enforcement action has been known for decades. It's not ACMA's fault that its powers are weak – the responsibility for that lies with successive federal governments. But there is an aspect that further complicates this particular enforcement dilemma, and one that might cause some reflection on ACMA's previous response to breaches by KIIS FM. All of the broadcasts that are the subject of the latest two investigations took place in 2024 and at that time, there was actually an enforceable undertaking in place for the Kyle and Jackie O Show. This undertaking, which ran for a period of two years, had been accepted in response to other breaches of the decency rule, including comments about the Paralympics and about gay men and monkeypox. A breach of this undertaking would mean ACMA could seek an order from the Federal Court. Except ... it looks like the undertaking itself has not been breached, even if the principal rule – the obligation to observe standards of decency – has been flouted on multiple occasions. This is because the undertaking required training and processes to promote compliance but didn't actually require compliance with any rule. One thing it did require was a "second, back-up censor, to assist the primary censor to monitor compliance". And we know from ACMA's media release that ARN did actually have a second censor in place, but still the system failed.

The approach taken in this undertaking stands in contrast to most licence conditions, which do impose an enforceable obligation to comply with a requirement such as the decency rule in the radio code. Back in 2010, when the Kyle and Jacki O program was on 2Day FM, ACMA even used a licence condition to require the licensee to have regard to the welfare and wellbeing of children who are interviewed on air. That followed the notorious episode where 14-year-old girl was attached to a lie detector and asked questions about her sexual conduct.

It's tempting in this situation to think in terms of 'ifs': if ACMA had used a licence condition rather than accepted an undertaking, and the same breaches in the first half of 2024 had been addressed by a remedial direction, then breaches in the second half of 2024 could have been addressed by an application to the Federal Court – or even a licence suspension.

In the end, we can't accurately say how it might have turned out if a different approach had been taken. But we can be unhappy about being back at the beginning of the enforcement path, like a game of Trouble or Snakes and Ladders. We can also say that having to choose between a licence condition and an undertaking is an unnecessary restraint on the regulator. Our broadcasting laws are not serving the public interest, and it's time they were changed.



Derek Wilding
CMT Co-director

Trust issues



Australia is currently in the midst of figuring out which social media platforms the country's youth will be banned from accessing come December 10, fuelling criticism that the entire approach is a knee-jerk, rushed intervention.

Case in point: recent comments from the eSafety Commissioner [indicated](#) that notorious cesspit of the internet – 4chan – will likely not fall under the ban because it is an “image board”. There are a few points in this comment to unpack.

First, 4chan doesn't require a user to sign up before posting or interacting. In that sense, 4chan isn't even captured by the legislation, which requires age-restricted social media platforms to prevent under16s from creating or maintaining accounts. This account-based limitation undermines the logic of protecting kids from harmful content: if the risk sits in the feed, removing accounts just shifts kids into feeds without protections.

While accounts may underpin other risks, such as hyper-targeted algorithms, which have given rise to concerns about addiction and echo chambers, banning accounts does not protect against harmful content. It may, in fact, make it worse. A recent Guardian Australia [test](#) found that simply accessing YouTube Shorts or TikTok in a logged-out state on a fresh device quickly surfaced gambling promos, violent incidents and far-right content. eSafety has confirmed the logged-out experience isn't subject to the ban, even while urging platforms not to undermine the intent of the law. That's the core loophole.

Google, the owner of YouTube, [reiterated](#) this point when it told a Senate inquiry that not only will the ban be “extremely difficult to enforce”, but that forcing teens into a logged-out experience will strip away the safety features parents rely on – no autoplay, break reminders, teen modes, or personalised filters. Meta is already [testing](#) AI to identify suspected under-18 Instagram users and automatically shift them into teen settings. Under this law, those protections will vanish if the user is simply logged out.

Meanwhile, definitional hair-splitting is doing a lot of work. Ignoring the account issue, the eSafety Commissioner's framing of 4chan as an "image board" is being used to justify an exemption from the ban, while YouTube's arguments that it's a "video-streaming platform" rather than social media haven't [carried any weight](#).

If you're looking for any coherence on what social media is, the list of companies approached by eSafety to prompt them to self-assess whether their platforms fall under the law doesn't help. Alongside the expected Facebook, Instagram, TikTok and X, eSafety [reached out](#) to GitHub – a developer collaboration platform – which speaks to how elastic the definition of "social media" has become. Contrast this with LinkedIn, which, despite being a full-blown social network, has argued its way out of the ban by being [too boring](#) for children. Does the government consider it unlikely that kids are on LinkedIn, therefore sparing Microsoft (LinkedIn's owner) the onerous burden of implementing age verification, or does it think it is valuable for kids to start professional networking and development early? Probably the former, considering recent comments that eSafety would [focus](#) on platforms with a large number of young Australian users. If this reflects eSafety's motivation for contacting GitHub, the future of open-source code development in Australia looks bright, if, of course, GitHub manages to escape the ban.

Back to 4chan. Officials said that 4chan would instead be captured by other Online Safety Act codes (designated internet services) and could face penalties up to \$49.5m, but also acknowledged eSafety hasn't formally assessed it and is taking a [risk-based approach](#) focused on "mainstream" platforms. This is the heart of the trust issue: leaving a well-documented dark-web-adjacent platform outside the flagship initiative to address youth online harm sends precisely the wrong signal to families about where the risk actually lives. It also ignores the practical reality that 4chan [reportedly](#) refused to pay a proposed UK fine – so even if captured by codes later, compliance is hardly guaranteed.

The ban both over-reaches and under-reaches: it dampens incentives to build age-appropriate safeguards, yet limits restrictions to the logged-in experience only. If Australia were serious about protecting children, it would either rethink its approach or ban access to harmful sites outright, account or not.



Kieran Lindsay
CMT Researcher

Beware the Ires of Musk!



A little over two years ago, I [reported on a trip](#) to Singapore I and my colleague Heather Ford took to visit Wikimania, a conference attended by volunteer [Wikimedians](#) from around the world. We were there to investigate what they [thought about](#) the potential impact of generative AI on Wikipedia as a trusted source of knowledge.

Many were worried, naturally, that the then already-infamous weaknesses of generative AI would lead to Wikipedia

being infected by error and bias, or by the [puffery](#) to which many [chatbots are prone](#). But the overriding concern was a deeper one, centred on the potential for generative AI to undermine the long-term viability of Wikipedia. If you can get all the information you need from a chatbot, why visit Wikipedia at all?

This worry was magnified once generative AI was incorporated into search engines, such as Google's AI Overviews. Early data was inconclusive, but [recent figures](#) show that traffic to all Wikimedia sites is now taking a hit, with human pageviews declining 8% over the past year. Similarweb reported a [drop of 23 percent](#) in visitors specifically to Wikipedia over the 3 years to March 2025. Unique visitors have also been falling.

These declines in traffic mirror those hitting news websites, a matter of [strong concern](#) amongst Australian news companies, with Similarweb showing [declines of over 30%](#) at some major news sites. At the same time, traffic from AI scrapers has [increased markedly](#).

With the cost of bot traffic higher than for human users, this is straining Wikipedia's digital infrastructure and increasing running costs. Some scrapers are programmed [to act like human users](#) to avoid detection. This leaves a bad taste in the mouths of many Wikimedians, who freely contribute their time to build and maintain the encyclopaedia, only for the contributions to be exploited for profit by AI companies.

Elon Musk's recent [announcement](#) of Grokipedia, an AI-driven competitor to Wikipedia, will likely have made this bad taste a whole lot worse. Musk has had it in for Wikipedia for a while, calling it 'Wokepedia' for what he perceives as a leftist bias. To the bad taste you can add a large dose of irony. Musk is building Grokipedia in part by piggybacking on Wikipedia content: using Grok to analyse a page to distinguish what is true from what is 'partially true, or false, or missing', and then rewriting it. But Grok, like other consumer AI tools, draws on Wikipedia's vast archives for much of its data.

Grokipedia v0.1 was expected to launch this week, but it has [been postponed](#) as, Musk says, more work needs to be done 'to purge out the propaganda'. When propaganda is a view you disagree with, this may take longer than Musk expects. The project is also unlikely to garner much long-term interest, as various other ideologically driven [online](#)

'[encyclopedias](#)' attest. Musk says Groklopedia is a 'necessary step towards the xAI goal of understanding the Universe.' But a true encyclopedia has a purpose that is at once more modest and more ambitious. It does not purport to know the truth in advance, as an ideologist does, but seeks to record what is currently known by relying on the expertise of others. Wikipedia's open, communal editing process relies on a common commitment to a [set of principles](#) that [lead over time](#) to greater accuracy and reduced bias. AI can offer – at best – only a simulacrum of that process.

For more, check out our [DoubleTake](#) podcast with Wikipedia Co-Founder Jimmy Wales.



Michael Davis

CMT Postdoctoral Fellow




We hope you have enjoyed reading this edition of the *Centre for Media Transition newsletter* - Games of trouble - Issue 19/2025

ISSN 2981-989X

This serial can be accessed online [here](#) and through the National Library of Australia. Please feel free to share our fortnightly newsletter with colleagues and friends! And if this was forwarded to you, please subscribe by clicking the button below:

Subscribe

Please visit our [website](#) for more information about the Centre.



The Centre for Media Transition and UTS acknowledge the Gadigal and Guring-gai people of the Eora Nation upon whose ancestral lands our university now stands.

We pay respect to the Elders both past and present, acknowledging them as the traditional custodians of knowledge for this land.



[Privacy Statement](#) | [Disclaimer](#) | [Unsubscribe](#)

UTS CRICOS Provider Code: 00099F

This email was sent by University of Technology Sydney, PO Box 123 Broadway NSW 2007, Australia