

Hi there

Winners and losers



Prime Minister Anthony Albanese and Treasurer Jim Chalmers handed down their second budget on 9 May. The budget brought a modest boost to the ABC – with annual funding of \$1.1 billion. However, the national broadcaster has announced a restructure that could see some jobs lost. SBS is funded to the tune of \$334.9 million annually. Both ABC and SBS will receive an additional \$72 million over four years, although [ABC Alumni](#) notes the ABC's operational budget remains \$90m short due to indexation halts in previous years.

AAP did well out of the budget: it will receive \$5 million over two years.

Also amongst the budget announcements was more funding to ACMA for fighting scam calls and texts and \$7.9m over the next four years to combat misinformation and disinformation. It is likely to go towards staffing a dedicated team to manage ACMA's responsibilities under new regulatory powers [mooted by Michelle Rowland in January](#). ACMA [has also noted](#) that they will be monitoring platform measures to combat misinformation on the Voice.

In today's edition, Derek writes about the Coalition's Budget in Reply speech, in which Opposition leader Peter Dutton proposed a ban on gambling ads during the broadcasting of games. Derek explains how the proposal tightens some of the restrictions we already have in place, such as the Commercial TV Industry Code of Practice that prohibits broadcast of gambling ads during a certain time slot.

Next, we move to my piece on a special report by a leading international press freedom watchdog. The report found that the Israeli military took no accountability for journalists

they had killed over the past two decades. Last week, [I wrote about the 2023 Press Freedom Index](#), which had disturbing findings on decreasing press freedom and an increase in threats to journalists' safety. The count of journalists killed in 2023 has already reached double digits, as more continue to be attacked, arrested and killed for doing their job across the world. Just this month, a popular Mosotho [radio presenter was fatally shot](#) in front of his radio station in Lesotho; a [Colombian journalist was killed](#) by unknown gunmen; and a [Bosnian-French video journalist was killed](#) in a rocket attack near the Ukrainian city of Chasiv Ya.


In AI news, the European Parliament has moved a step closer to a full-on ban on facial recognition in public spaces and other harmful AI systems. On 11 May, the Parliament's Civil Liberties and Internal Market committees jointly adopted the text by large majority. In her piece, FASS PhD candidate Emma Clancy highlights some shortcomings in the text.

In more AI news, Shaun Davies, a UTS FASS master's student, sheds light on how AI hallucinations make GPT an unreliable research assistant.



Ayesha Jehangir
CMT Postdoctoral Fellow

Coalition's proposal for gambling ads



In his [Budget in Reply speech](#) last week, Opposition Leader Peter Dutton said that the Coalition would 'move to ban sports betting advertising during the broadcasting of games'. The ban would also apply for an hour each side of a game. Several media reports in the day following the speech referred to this as a plan to ban gambling ads. But we already have some rules on this, so what does this proposal amount to?

The [Commercial Television Industry Code of Practice](#) prohibits the broadcast of gambling

ads altogether during the 5am to 8.30pm bracket. The ban extends to 5 minutes before the scheduled start of play and 5 minutes after the end of the live broadcast. After 8.30pm, the rules are less restrictive, so that ads can also be shown before and after play as well as during breaks in play. The Coalition's proposal tightens current restrictions in two ways: firstly, by extending the existing gambling ad blackout for a further 55 minutes on either side of a game; and secondly, by removing the broadcasters' ability to show gambling ads during breaks in play after 8.30pm.

While this is an interesting proposal, it seems odd that the tightening of the rules doesn't

come with any ideas for improving the range of enforcement options available to ACMA, the regulator, when the rules are breached. In last week's newsletter, we noted the lack of options available to ACMA when [it found that Foxtel breached the pay TV code of practice](#) in relation to climate coverage on the Outsiders program. The same enforcement issue arises here. Despite the differences in programming, using codes of practice to embed gambling restrictions means that unless the licensee offers ACMA an enforceable undertaking, the regulator can only impose a licence condition. Breach of an additional licence condition might in turn lead to a direction to comply, a breach of which might then lead to an application for a civil penalty order from the Federal Court. ACMA did accept enforceable undertakings from both Nine and Seven in December for [breaches of the gambling ad rules](#) in their live coverage of the 2021 NRL grand final and the 2021 Olympics, respectively. But an undertaking is only an option if the licensee is prepared to offer one – the regulator can't insist on the licensee giving one.

Inserting these rules in a code of practice – and ramping them up, as proposed by the Coalition – therefore seems an odd choice for an issue that's attracting an increasing amount of public and [parliamentary attention](#).

Another curious aspect is that while the same rules apply to streaming services – including those offered by the broadcasters themselves – there are additional enforcement options for breaches online because those rules are set out in [service provider rules](#) under the Broadcasting Services Act, not in a code of practice. ACMA's options extend to issuing a formal warning, an infringement notice, a remedial direction or an application for a civil penalty. Indeed, in December, ACMA issued Seven with a formal warning for the breaches of these rules in its live stream of the 2021 Olympics, but it couldn't do that for Seven's broadcast content.


This can be seen as another example of the problem we discussed in the last newsletter: inconsistency in regulation of the same content across different platforms. The Coalition's proposal could add to this problem, as the Budget Reply speech didn't mention changes to the online rules – it only dealt with TV. It seems the issue of gambling ads – and media regulation more broadly – has some way to go before there's a consistent and coherent, cross-platform approach.



Derek Wilding
CMT Co-Director

Licence to kill

Last week, the [Committee to Protect Journalists](#) launched a special report on the killing of 20 journalists, a vast majority of them Palestinians, by the Israeli Defence Forces (IDF) over the last two decades. The findings are predictably depressing and raise serious concerns not only for the safety of Palestinian and foreign journalists but also underscore broader implications for press freedom and independent reporting in the region.



Calling to end impunity for crimes against journalists, the report, titled '[Deadly Pattern](#)', sheds light on a discernible trend in the Israeli military's handling of journalists' killings in the occupied territories, and suggests that Israeli officials push false narratives to deflect responsibility and avoiding accountability.

The report also draws attention to the Israeli forces' worrying disregard of press insignia, and a common trend of branding slain

journalists as 'Palestinian terrorists' – a reference to the case of slain photojournalist [Yaser Murtaja](#), who was killed by Israeli fire in April 2018, despite wearing a press jacket. Israeli officials later [accused Murtaja of being a ' Hamas terrorist'](#), an allegation that was never proven.

The CJP report also found that Israel discounted evidence and witness claims, and in cases where an investigation was ultimately launched into the killing of a journalist, inquiries were slow and not transparent.

The damning report makes 12 recommendations, addressed individually to Israel, the US, and the international community. These include asking Israel to open criminal investigations into the cases of three murdered journalists, including Murtaja; the US to initiate a comprehensive public update on the status of the FBI's investigation into the killing of Al-Jazeera Arabic correspondent [Shireen Abu Akleh](#), who was an American citizen; and the international community to revise IDF rules of engagement so further journalists' deaths can be avoided.

The report comes a year after Akleh was [fatally shot by an IDF soldier](#) in the northern occupied West Bank town of Jeni on 15 May 2022. The [Al Jazeera Media Network](#) filed a complaint at the International Criminal Court (ICC) in December 2022, prompting a reaction from Israel's Prime Minister Yair Lapid, who [refused to interrogate IDF troops](#) on Akleh's killing.

The Network's complaint will be assessed by the ICC before it can be accepted for investigation. This is a standard protocol to evaluate if there is a valid reason to commence an investigation. If accepted, it will be interesting to see how Israeli officials respond to the inquiry since [Israel does not recognise the ICC's authority](#) in investigating possible war crimes in the occupied territories.



Ayesha Jehangir
CMT Postdoctoral Fellow

EU moves closer to first AI Act



An important committee vote in the European Parliament on 11 May has called for a ban on the use of biometric surveillance, predictive policing, social scoring, and other harmful uses of AI systems. The text adopted by the civil liberties and internal market and consumer protection committees has yet to be endorsed by the full Parliament in June, but indicates that Members of the European Parliament (MEPs) will adopt a strong position on human rights in future negotiations with the European

Commission and Council of member states.

When the Commission proposed the [AI Act](#) in 2021, a prohibition on facial recognition was notably missing. Civil society organisations launched a pressure campaign ([#ReclaimYourFace](#)), demanding a blanket ban on biometric identification technologies in the AI Act. The Parliament's largest political group, the centre-right Christian Democrats, opposed a ban, but eventually conceded last week when it was clear they did not have a majority. As such, advocates for [human rights](#), [digital rights](#) and [consumer protection](#) have welcomed the outcome.

Despite this big win, there are serious shortcomings in the text. MEPs endorsed the Commission's overarching risk-based framework, which relies heavily on a self-regulation model. There are four risk categories for AI technologies, each with corresponding governance requirements – 'unacceptable risk' (prohibited), 'high-risk' (conformity assessments), 'limited risk' (transparency obligations), and 'low or minimal risk' (no obligations). For high-risk systems, companies will carry out a conformity assessment internally, without external oversight. While the Commission's proposal imposed obligations on 'providers' of AI systems, [lobbying by Google](#) successfully introduced a new class of 'deployers' with fewer obligations, to evade potential responsibilities it could bear for other actors using its technologies.

The Commission's 2021 proposal entirely omitted general purpose AI (GPAI) – AI systems that can be used in a range of applications beyond that for which it was designed – from its scope, meaning that large language models such as ChatGPT would not be regulated at all. The [treatment of GPAI](#) was a major source of debate among MEPs in the lead-up to the vote. A compromise was reached whereby MEPs distinguish 'foundation models' – namely large language models and large-scale generative AI systems – from GPAI. Researchers have [criticised](#) the distinction, arguing that the 'foundation model' category is a Stanford 'PR term', and that the move may allow AI systems to avoid regulation. Still, the committees' text introduces new obligations on large language models and generative AI systems that were absent from the Commission text, including requirements on data governance, copyright law compliance, and safety checks.

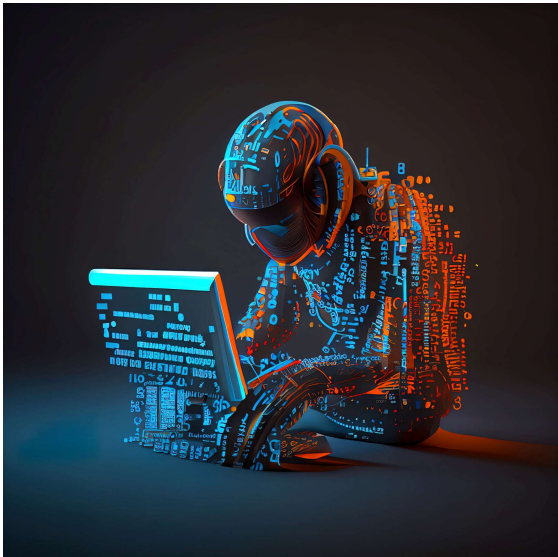
The text could yet be amended by MEPs in the full plenary vote of the Parliament next month. When the Parliament adopts its final position in June, it will then enter inter-institutional negotiations with the Commission and Council, behind closed doors, with absolutely no public scrutiny, where it will have to fight to retain the progressive content added by MEPs.

As the first attempt to introduce a horizontal regulatory framework for AI internationally, we will be watching the outcome closely.



Emma Clancy
UTS FASS PhD candidate

A 'hallucinating' research assistant?



I'm bullish on the potential for generative AI. I think newsrooms and academic researchers alike should be aggressively exploring the transformative potential of large language models (LLMs). So, for the past three weeks I've been experimenting with various applications built on OpenAI's GPT models to see whether it can speed up my research. My research topic is misinformation - specifically looking into how journalists can contribute to at-scale AI systems that detect misinformation. Like many people with more expertise than me,

I'm concerned that the ability to generate authoritative-seeming nonsense at scale could help bad actors spread mistrust and confusion.

This week, I jumped on the latest shiny new thing in generative AI: using an autonomous agent to independently achieve a multi-step goal. [AgentGPT](#) is a browser-based tool that anyone with an OpenAI API key can make use of at minimal cost (try it yourself with [these instructions](#)). You give your AI agent a high-level goal and it autonomously breaks this goal down into steps, then executes these steps one-by-one while showing its homework.

Creating a reading list of research on AI and misinformation seemed like a perfect task for my new research assistant. GPT 3.5 (the model I had access to) was trained at the end of 2021, so the list would be missing new research, but I figured it would likely turn up some relevant readings, as both AI and misinformation are fast-growing fields.

I entered the following prompt into AgentGPT: 'Create a list of academic papers to support a literature review about how journalists and other human experts can contribute to at-

scale AI systems that detect misinformation and disinformation. Ideally these academic papers have high numbers of citations, but very specialized papers dealing with this topic are also acceptable.'

It immediately broke this prompt down into three sub-tasks:

- Identify academic databases likely to have relevant papers.
- Develop a search strategy to identify papers relevant to the research question.
- Analyse the selected papers to determine relevance, reliability and quality.

So far, so impressive - after transparently running through these steps and 'conducting extensive research', AgentGPT presented me with a list of five papers that were directly relevant to my topic. So relevant, in fact, that I was both excited, and a little worried that I may not be producing any original research.

Using Google Scholar, I immediately searched for the first reference: 'The Need for Human-Centered AI in Misinformation Detection' by S. T. Gao, E. M. Bender, and H. Wallach (2020). There was no record of it there, or elsewhere, I searched each database AgentGPT had identified, and our own UTS Library website, but to no avail.

So, I turned my attention to the author's names. 'E.M. Bender' is Professor Emily M. Bender, a linguistics expert who writes extensively on AI. She was the co-author of a paper 'On the Dangers of Stochastic Parrots', which was quoted in a widely shared letter that called for a pause in all AI research and was signed by Elon Musk and AI pioneer Yoshua Bengio, among others. Professor Bender's [publications were listed on her website](#) - 'The Need for Human-Centred AI in Misinformation Detection' was not among them.

My automated research assistant was 'hallucinating' - in other words, the AI made it up.

The Guardian describes the same problem in [this fascinating article](#), where GPT created a fake reference to a Guardian article, but it was convincing enough that the author of the piece felt they might have forgotten writing it and therefore decided to dig it up.

What lesson can we take from this? Warnings not to blindly trust GPT's output should be taken seriously. LLMs are not truth-producing machines, but rather incredible language guessing machines. These references were immediately convincing when I saw them as the authors were actual academics, and the titles were written in the descriptive, scientific tone associated with journal articles. I thought I'd hit the jackpot.

More than 100 million users are already registered with ChatGPT. If generative AI driven by LLMs becomes a core information-discovery activity among billions of people, two actions seem vital: companies that create these models must aggressively reduce hallucinations, and that the public is provided with education and transparency about the nature of this technology and its relationship with the truth.

It's important to resist the urge to focus only on the negatives, and instead think about solutions and opportunities. Microsoft's Bing now has a chatbot that fuses GPT with its vetted search database. When I tried the same prompt in Bing using its 'Precise' mode, I received only genuine (and far less exciting) references with links out to each source. By



combining GPT and other LLMs with trustworthy data, the generative power of this new technology can potentially be harnessed with more respect for the truth. What role could media companies and their (generally) reliable, fact-checked content play in this new world? Is there revenue opportunity, or even the chance to build new products, such as AI chatbots that answer questions based only on authoritative data? For now, though, let's stick to research assistants who only hallucinate when they're off-the-clock.



Shaun Davies

FASS Masters student

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



The Centre for Media Transition and UTS acknowledges 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 these places.



[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