Responsibility of OpenAI for Defamation and Serious Invasions of Privacy by ChatGPT

Author: Dr Michael Douglas

ChatGPT is an amazing product, underpinned by amazing technology. It can be incredibly useful. But notoriously, the information the bot produces is not always reliable. As much is admitted by the organisation behind ChatGPT, OpenAI.

Inconsistent reliability is a necessary consequence of the large language model (**LLM**) underlying ChatGPT. It works by choosing words in a sentence on the basis of probabilities. There is a contingent relationship between the propositions produced by LLMs and the truth. For example, ChatGPT can produce some basic mathematical errors;[1] it does not always give effect to deductive logic. This is why some AI gurus, like Gary Marcus,[2] argue that AI research needs to go in another direction—one where truth and deference to facts is built into the technology from the ground up.

A truth-centred AI would be wonderful, but that's not the kind of technology that has taken the world by storm in recent months. ChatGPT vacuums the content from the web and pumps out the best content it can come up with in response to a prompt. Some of that content is great, and some of it is nonsense.

Where nonsense damages a person's reputation, it may be the subject of a defamation dispute. Recently, that prospect has been realised with respect to content produced by ChatGPT.

Brian Hood is Mayor of Hepburn Shire in Victoria, Australia.[3] Members of the public told Hood that ChatGPT was identifying Hood as party to a high-profile bribery scandal.[4] Prompts fed into ChatGPT that included Hood's name would return results that would identify him as a wrongdoer in the scandal, who was convicted of a crime and sentenced to prison.

Those results were utterly false. Hood was actually the whistle-blower.[5]

At the time of writing, Hood had engaged lawyers who sent OpenAI a concerns notice. Unless the company remedies the situation, a defamation case is likely. That case will break ground in confirming the capacity of ChatGPT content to be the subject of a defamation action, and the prospective liability of OpenAI for that content.

Prospective defamation liability for content produced by ChatGPT

It should be uncontroversial that content produced by ChatGPT may be amenable to defamation liability.

Throughout Australia, defamatory 'matter' is the cause of action: *Defamation Act 2005* (WA) s 8. 'Matter' is defined in a technology-neutral way to comprehend any potential means of communication: *Defamation Act 2005* (WA) s 4. Text generated by a bot is within its scope. So for example, auto-complete results of the Google search engine are capable of being defamatory matter, as the High Court recognised in *Trkulja v Google LLC* (2018) 263 CLR 149.

Does ChatGPT matter have defamatory capacity? It depends on what is pumped out, but Hood's case suggests that it can where the generated text disparages a person's reputation. OpenAI may argue that it has sufficiently disclaimed any such content. At the time of writing, the landing page of the interface of the chatbot (or at least the free version) says the following under the heading "limitations": 'may occasionally generate incorrect information'; 'may occasionally produce harmful instructions or biased content'; 'limited knowledge of world and events after 2021'. Those disclaimers are better than nothing but won't be enough to provide a bane and antidote defence. Further, the ordinary reasonable ChatGPT user engages with the bot to get useful information. Even if that hypothetical user accepts that the matter is contingently reliable, that could be enough to cause the user to think less of a person smeared by the generated content: see *Trkulja v Google LLC* (2018) 263 CLR 149, [60].

If defamatory matter is comprehended by a person other than the person defamed, there will be publication. Following the introduction of the statutory serious harm element in much of Australia (see, eg, *Defamation Act 2005* (NSW) s 10A), ordinarily, the content must be seen by a non-negligible number of people to be actionable. But even if the matter is seen by just a handful of people, if the meaning conveyed is serious enough, the serious harm threshold may be satisfied: see *Scott v Bodley (No 2)* [2022] NSWDC 651, [45].

Hood's case may be somewhat unique in that members of the public told Hood that ChatGPT was producing damaging content. But if a person were to simply plug their own name into the machine, not only would the publication element be a problem, so too would satisfaction of the serious harm element.

Provided the publication element is satisfied, who might be liable for that publication? Any person, including a company, that participates in the communication of defamatory matter to any degree may be on the hook as a publisher: *Fairfax Media Publications Pty Ltd v Voller* (2021) 95 ALJR 767, [30].

1 See, eg Josh Zumbrun, 'ChatGPT Needs Some Help With Math Assignments', *The Wall Street Journal*, 10 February 2023 <https://www.wsj.com/articles/ai-bot-chatgpt-needs-some-help-with-math-assignments-11675390552>.

2 See The Ezra Klein Show, 'Transcript: Ezra Klein Interviews Gary Marcus', *The New York Times*, 6 January 2023 <https://www.nytimes.com/2023/01/06/podcasts/transcript-ezra-klein-interviews-gary-marcus.html>.

3 See Hepburn Shire Council, 'Mayor and Councillors', 2023 <https://www.hepburn.vic.gov.au/Council/About-Council/Councillors>.

4 Byron Kaye, 'Australian mayor readies world's first defamation lawsuit over ChatGPT content', *Reuters*, 6 April 2023 <https://www.reuters.com/technology/australian-mayor-readies-worlds-first-defamation-lawsuit-over-chatgpt-content-2023-04-05/>.

5 See Commonwealth Department of Public Prosecutions, 'Securency and Note Printing Australia foreign bribery prosecutions finalised', 2023 <https://www.cdpp.gov.au/case-reports/securency-and-note-printing-australia-foreign-bribery-prosecutions-finalised>.

When defamatory matter is produced by a machine, and where that machine is produced by a company, the attribution of responsibility for that matter to the company would ordinarily be straightforward, applying orthodox principles on the attribution of corporate responsibility.

The situation is trickier where 'the machine' provides links to content, and the content underlying the link is defamatory, but the presentation of the link is not defamatory on its face. In *Google LLC v Defteros* (2022) 96 ALJR 766, the majority of the High Court held that Google did not publish webpages linked by the results page of its search engine, although it published the page of links. The majority view still leaves room for the possibility that a person could be responsible for content underlying the link, particularly where the link is not presented in a neutral way; for example, where its presentation is prioritised as a sponsored link.

The content produced by the ChatGPT bot is quite different to the content under consideration in the *Defteros* case. The machine has produced content which is defamatory on its face; the user need not dig any further. For the purposes of Australian defamation law, the company behind the tech is a publisher of any content produced by the bot, whatever the prompt inputted by the user. If an analogy were to be drawn with a Google case, ChatGPT content is closer to the defamatory snippets considered in *Duffy v Google Inc* (2015) 125 SASR 437.

Defensibility of defamatory ChatGPT content

Would OpenAI have a defence? I previously thought that it would with respect to defamatory content it does not know about.

Relevant defences include those available under the *Online Safety Act 2021* (Cth) (**Online Safety Act**) s 235, and also the innocent dissemination defence provided by the Defamation Acts of the States and Territories (see, eg, *Defamation Act 2005* (WA) s 32). The innocent dissemination defence may be defeated if the defendant 'ought reasonably to have known' that the matter was defamatory, and its lack of knowledge was not due to its own negligence. Neither of those defences would be available in Mayor Hood's case, as the publisher has been made aware of the content via a concerns notice.

Since I initially expressed that view, Associate Professor Jason Bosland and others pointed out the flaws in that view.[6] He was right to do so: I wrote my initial thoughts very quickly and had not thought it all through well enough.

Bosland argues that OpenAI could not be considered to be an innocent disseminator of content created through its own algorithm. By contrast, my previous view would entail that OpenAI be characterised as a 'subordinate distributor' under *Defamation Act 2005* (WA) s 32(2) and equivalents:

(2) For the purposes of subsection (1), a person is a subordinate distributor of defamatory matter if the person

(a) was not the first or primary distributor of the matter;

(b) was not the author or originator of the matter; and

(c) did not have any capacity to exercise editorial control over the content of the matter (or over the publication of the matter) before it was first published.

The provision requires that a publisher satisfy all three criteria in order to be characterised as a 'subordinate distributor'. Each limb presents difficulties for OpenAI. For example, by developing the underlying tech and providing it to the public, OpenAI might be characterised as the 'author or originator' for the purposes of s 32(2)(b).

In response, OpenAI may argue that it is merely acting as a conduit: that any content produced depends on what is available (and scraped) on the broader internet, and whatever the user plugs into the platform; accordingly, OpenAI is not the author or originator. Google has advanced similar arguments with respect to its search engine: see eg *Google Inc v Duffy* (2017) 129 SASR 304, [123], [598].

But even if OpenAI succeeds in that argument, it will struggle to establish that it was not the 'first or primary distributor' under s 32(2)(a). That is particularly the case where the content produced by its chatbot departs from the meaning of content available on the internet, like in Hood's case.

Further, OpenAI may similarly struggle to establish that it lacks the capacity to exercise editorial control over the content for the purposes of s 32(2)(c). OpenAI's success on such an argument would likely require evidence of how it makes its sausages. Would the company want to produce that evidence, or would it settle to avoid the hassle? Does it even understand the ghosts in its machine?

Bosland is right: OpenAI's characterisation as a 'subordinate distributor' of ChatGPT content is unlikely. Still, I feel for them. Section 32(3)(b) provides that 'a person is not the first or primary distributor of matter merely because the person was involved in the publication of the matter in the capacity of … a librarian'. You might argue that ChatGPT is a really fast research librarian—one who won't just find you books but will summarise them for you; but one who will sometimes get their summary wrong. Arguably, technology has advanced to a point where the policy of the statutory innocent dissemination defence is no longer achieving its purpose.

Bosland may be right about the Online Safety Act too.

The operation of the Online Safety Act s 235 defence may turn on whether OpenAI may be characterised as an 'Australian hosting service provider' under s 235(1)(b). 'Hosting service' means (per ss 5, 17) a social media service, a relevant electronic service, or a 'designated internet service'. 'Relevant electronic service' means technologies that facilitate communication, like email and instant messaging (see s 13A), and would not apply here. 'Designated internet service' is defined under s 14 to include, inter alia, a service that allows end-users to access material using the internet (s 14(1)(a)); whether ChatGPT could be shoe-horned into this definition remains to be seen.

There is also a mechanism for the Minister to designate certain material as a 'designated internet service' under Online Safety Act ss 14(1)(g), (2). ChatGPT has not been so

6 See LinkedIn chat of April 2023 available at: <https://www.linkedin.com/feed/update/n:li:activity:7050390391392505856?commentUrn=urn%3Ali%3Acomment%3A%28activity%3A7050390391392505856%2C7050394533695520768%29&replyUrn=urn%3Ali%3Acomment%3A%28activity%3A7050390391392505856%2C7051768800723038209%29>.

designated, but if it were, that could provide a defamation shield to OpenAI: it would mean that OpenAI operates a 'hosting service' that is captured by the s 235 defence insofar that ChatGPT involves 'hosting material *in Australia*' (emphasis added) under s 5. Anyway, I can't see an Aussie Minister making a relevant designation any time soon.

On the other hand, in recent years, there has been appetite from Australian legislators, cheered on by big-tech lobbyists and well-meaning bureaucrats seduced by the free-speech vibes of the lobbying, to give mostly American tech companies new shields from defamation liability: see the proposals of the NSW-led Stage 2 review into the Model Defamation Provisions.[7] Perhaps AI-friendly defo defences will be the subject of future lobbying.

While defamation defences for tech companies will soon be strengthened by Australian legislators,[8] even if those defences were in force now, they would not prevent Hood from succeeding. The incoming law requires tech companies to have a system for receiving complaints of defamatory content, and to remove the content within a certain period of time. ChatGPT has no such system.

The transnational dimension

A person defamed by ChatGPT cannot sue ChatGPT itself. They would need to sue the company, or companies, responsible for it.

These days, 'OpenAI' is a corporate group. The non-profit in which Elon Musk and others invested is OpenAI Inc.[9] OpenAI Inc has a for-profit subsidiary, Open AI LP.

OpenAI LP describes itself as a 'capped-profit' company; it is a 'Delaware Limited Partnership' controlled by a single-member Delaware company.[10] That Delaware company is controlled by OpenAI Inc. Under Delaware law, a Limited Partnership seems to be a distinct entity from its members. OpenAI LP has its place of business in California, but is relevantly located in Delaware, USA.[11]

While the corporate machinations of OpenAI are a little hard to digest, it seems that OpenAI LP is the appropriate entity to go after if you find yourself defamed by ChatGPT content. From our perspective in Australia, OpenAI LP is a foreigner located outside of our jurisdiction.

Australian courts will readily entertain defamation suits against foreigners outside Australia where the foreigner publishes defamatory content in Australia. John Barilaro did just that when he sued Google LLC over content the company published on YouTube, which was originally created by commentator / activist / pest FriendlyJordies: see *Barilaro v Shanks-Markovina (No 1)* [2021] FCA 789.

Suing a foreign company requires service outside the jurisdiction, which in the case of a company would ordinarily require delivery of a court document to the company's corporate headquarters. In circumstances where you know where the foreigner is located, and you can present evidence to a court showing that you have been defamed by that foreigner's publication, service outside of the jurisdiction won't be too difficult.

Provided that the defamatory ChatGPT content is published to persons in Australia, proceedings against a foreign publisher will not be stayed: see *Dow Jones & Co Inc v Gutnick* (2002) 210 CLR 575. And even if the publications sued upon are made to persons outside of Australia, I am not convinced that the court should stay the proceeding where it is brought by an Australian resident. With great respect to her Honour McCallum J (as her Honour was), the decision in *Bleyer v Google Inc* (2014) 44 NSWLR 670 runs contrary to the 'clearly inappropriate forum' test for *forum non conveniens* mentioned in *Dow Jones & Co Inc v Gutnick* (2002) 210 CLR 575, [3]. As abuse of process stays (like in *Bleyer*) depend on the same inherent (or in the case of the Federal Court, implied) power, the test applied by the High Court in cases like *Voth v Manildra Flour Mills Pty Ltd* (1990) 171 CLR 538 must take precedence, with the effect that an Australian court may conceivably entertain a claim for *foreign* publication of defamation: see further Michael Douglas, 'Forum Shopping in Australian Defamation Litigation' (PhD Thesis, University of Sydney, 2022). A defamation plaintiff who has regularly invoked the authority of a superior court of Australia has a prima facie right to insist upon its exercise: *Oceanic Sun Line Special Shipping Co v Fay* (1988) 165 CLR 197, 243 (Deane J).

Once served, however, OpenAI LP might just ignore service, pretending they are not subject to the Australian court's jurisdiction. The companies behind Twitter have behaved like that when faced with Aussie lawsuits: see eg *X v Twitter* (2017) 95 NSWLR 301. But if OpenAI LP were to do that, it would risk default judgment.

Enforcing a defamation judgment against OpenAI

If an Australian court then determines that OpenAI LP is liable for defamation, the judgment creditor may face enforcement issues.

Enforcing a judgment against a judgment debtor within the court's territorial jurisdiction, or within Australia, is not too difficult—provided you can locate them. Courts have powers to hold people to account. In WA, for example, you can get an enforcement order to force a recalcitrant judgment creditor to pay up: see *Civil Judgments Enforcement Act 2004*

7 See NSW Department of Communities & Justice, *Review of Model Defamation Provisions*, 'Standing Council of Attorneys-General, Stage 2 Review of the Model Defamation Provisions (MDPs)' (2023) <https://www.justice.nsw.gov.au/justicepolicy/Pages/lpclrd/lpclrd_consultation/review-model-defamation-provisions.aspx>.

8 'On 9 December 2022, the Standing Council of Attorneys-General approved in principle final amendments for Part A of the Stage 2 Review of the Model Defamation Provisions. This is subject to final agreement in the first half of 2023': see Standing Council of Attorneys-General communiqué, December 2022 <https://www.ag.gov.au/about-us/publications/standing-council-attorneys-general-communiques>.

9 See Sawdah Bhaimiya, 'OpenAI cofounder Elon Musk said the non-profit he helped create is now focused on "maximum-profit," which is "not what I intended at all"', *Business Insider*, 17 February 2023 <https://www.businessinsider.com/elon-musk-defends-role-in-openai-chat-gpt-microsoft-2023-2>.

10 See OpenAI, *OpenAI LP* (2023) < https://openai.com/blog/openai-lp>.

11 See the US Securities and Exchange Commission record (2023) available at: <https://www.sec.gov/Archives/edgar/data/1877240/000187724021000001/xslFormDX01/primary_doc.xml>.

(WA). That order might be physically enforced by a sheriff (assisted by police), who can do things like rock up to a person's house and possess property: see *Civil Judgments Enforcement Act 2004* (WA) s 74(2)(b).

Courts can also enforce their orders by holding persons who disobey orders guilty of contempt of court. Contempt of court is serious. Australian courts have various powers to punish people for contempt. A contemnor might be smacked with a big fine, or even sent to prison: see eg *Doe v Dowling* [2017] NSWSC 202.

What happens, though, if a foreigner refuses to comply with an Australian court's orders?

With respect to defamation cases, that question is significant. The law of the USA is quite different with respect to defamation liability. Among other things, the US has a constitutional right to freedom of expression protected by their First Amendment. Australia lacks an equivalent. Compared to Australian law, the US takes a different approach to the balance to be struck between allowing freedom of expression on the one hand, and protection of reputation on the other.

In 2010, the US Congress passed the 'SPEECH Act'. It provides that a US 'domestic court shall not recognize or enforce a foreign judgment for defamation unless the domestic court determines that' the relevant foreign law would provide the same protections for freedom of speech as would be afforded by the United States Constitution (SPEECH Act s 3; United States Code, title 28, Part VI, § 4102).

A defamation judgment against OpenAI over ChatGPT content would be very unlikely under US law. This means that the SPEECH Act would likely prevent enforcement of an Australian defamation judgment in the US.

The practical consequence of all this is that even if Hood wins his defamation case, OpenAI may refuse to pay, hiding behind the shield of US law and the fact the company and its assets are not in Australia.

Does this mean the case is not worth bringing?

Not at all.

OpenAI may choose to pay a debt owed by a judgment of an Australian court in order to 'keep up appearances' in the Australian market. Although our country is not as powerful as those who would provide the bulk of ChatGPT's user base, it is a developed nation of 25 million people, and a US ally. It would be good politics, and arguably good business, for OpenAI to pay up.

If OpenAI refuses to pay, it risks being held in contempt of court. And if OpenAI is in contempt, then its directors face personal risk for refusing to cause the company to obey the court if they know that the company should have obeyed. An English court once explained:

> 'In our view where a company is ordered not to do certain acts or gives an undertaking to like effect and a director of that company is aware of the order or undertaking he is under a duty to take reasonable steps to ensure that the order or undertaking is obeyed, and if he wilfully fails to take those steps and the order or undertaking is breached he can be punished for contempt. We use the word 'wilful' to distinguish the situation where the director can reasonably believe some other director or officer is taking those steps': *Tuvalu v Philatelic Distribution Corp Ltd* [1990] 1 WLR 926, 936E–F (Woolf LJ).

Australian courts will follow the policy of that decision: see *Australian Competition and Consumer Commission v Goldstar Corporation Pty Ltd* [1999] FCA 585, [41] (Kiefel J).

English courts have also been willing to exert their contempt powers against persons outside of the court's territorial jurisdiction; see *Dar Al Arkan v Al Refai* [2015] 1 WLR 135. Australian courts would likely follow suit.

With respect to the situation of OpenAI, it is notable that its CEO and founder is Sam Altman. A little googling reveals that Altman is a very wealthy American. His current partner is Oliver Mulherin, an Australian software engineer.[12] If Altman wishes to continue visiting Australia, as he did earlier this month, then causing OpenAI LP to *not* be in contempt of an Australian court would be a sound decision.

Even in a worst-case scenario, if OpenAI never pays a defamation judgment pronounced by an Australian court, then the judgment creditors to such judgments—people like Mayor Hood—are not left without nothing. They will be left with vindication of their reputation; with proof that they have been defamed. For many defamation plaintiffs, that is the real point of suing.

Looking forward to liability for privacy wrongs

The liability risk for OpenAI under Australian law is magnified by proposed privacy law reforms, whose enactment may be on the horizon. The long-awaited *Privacy Act Review Report*, published by the Commonwealth Attorney-General's Department in early 2023, proposed a number of legislative changes that could be a huge deal for media lawyers and their clients.[13] Relevantly to this article, the many proposed reforms include the following changes with respect to the *Privacy Act 1988* (Cth):

- Creation of a statutory tort for serious invasions of privacy, covering misuse of information and intrusion upon seclusion (proposal 27). This would *finally* give effect to the recommendations of the 2014 ALRC Report 123, the report of an inquiry led by Professor Barbara McDonald;
- Creation of a new right to erasure, which would apply to APP entities (proposal 18, especially 18.3);
- Creation of a direct right of action, allowing persons to sue over various breaches of the Privacy Act—seemingly, including a breach of the right to erasure (proposal 26); and
- Allowing the direct right of action to be litigated via representative proceedings (see p 273ff).

Each of these could expose OpenAI to risk.

12 See Cade Metz, 'The ChatGPT King Isn't Worried, but He Knows You Might Be', *The New York Times*, 31 March 2023 <https://www.nytimes.com/2023/03/31/technology/sam-altman-open-ai-chatgpt.html?smid=nytcore-ios-share&referringSource=articleShare>.

The Report also proposed the creation of a new right (proposal 18.5) to de-index search results, similar to the GDPR right once called the 'right to be forgotten'. The proposal is for a right to de-index search results that are (inter alia) 'inaccurate, out-of-date, incomplete, irrelevant or misleading' (proposal 18.5(iv)). If this were extended to chatbots like ChatGPT, it could throw a spanner in the works of OpenAI and its AI competitors. Indeed, as much is already happening in the European Union.[14] The situation in the Australian market may turn out similarly messy.

Conclusion

ChatGPT is awesome. It has the potential to generate a great deal of good for humanity. But it also has the potential to generate a great deal of damage. That damage may be felt around the world, not just in those jurisdictions close to OpenAI.

It is entirely reasonable that Australians avail themselves over the remedies provided by Australian law when foreign companies cause them damage. That moral claim is even stronger in circumstances where foreign companies make a deliberate decision to be available to the global market, and so within the Australian market. OpenAI chooses to make ChatGPT available in Australia; its responsibility to comply with Australian law is a foreseeable and reasonable consequence of that decision.

On the subject of AI, Gary Marcus recently told *The New York Times* that '[w]e have a perfect storm of corporate irresponsibility, widespread adoption, lack of regulation and a huge number of unknowns'.[15] Civil litigation, like that following Mayor Hood's defamation dispute, may fill the gap left by legislators and other regulators. If the result is that OpenAI takes steps to make sure its technology minimises harm to individuals, then I support it.

Dr Michael Douglas is Senior Lecturer at UWA Law School and a Consultant at Bennett. Previous versions of this article were posted via LinkedIn, on the website of Bennett, and published by the Gazette of Law & Journalism. Readers of the Communications Law Bulletin are encouraged to support the GLJ.

13 Australian government, *Privacy Act Review: Report 2022* (2023) <https://www.ag.gov.au/sites/default/files/2023-02/privacy-act-review-report_0.pdf>.

14 See Jess Weatherbed, 'OpenAI's regulatory troubles are only just beginning', *The Verge*, 5 May 2023 < OpenAI's regulatory troubles are only just beginning>.

15 Cade Metz and Gregory Schmidt, 'Elon Musk and Others Call for Pause on A.I., Citing "Profound Risks to Society"'. *The New York Times*, 29 March 2023 <https://www.nytimes.com/2023/03/29/technology/ai-artificial-intelligence-musk-risks.html

CAMLA Young Lawyer Committee Event Report: Journalists' Privilege Webinar

Author: Kristina Hewetson, Associate at Baker McKenzie and CAMLA Young Lawyer Committee Member

On 31 May 2023, the CAMLA Young Lawyers Committee had the privilege of hosting a fantastic webinar panel comprised of three experts to discuss the complex web of journalists' privilege in Australia. Moderated by Imogen Loxton, Senior Associate at Ashurst and CAMLA YLC Member, the panel featured:

Dr Matthew Collins AM KC,
Barrister at Aickin Chambers.

Gina McWilliams
Senior Legal Counsel at News Corp.

Paul Farrell,
Investigative Reporter at the ABC.

Gina McWilliams and Matthew Collins commenced the discussion by providing an informative history of journalists' privilege, also known as "shield law". Prior to the introduction of that law, Australia was described as "the wild west", where journalists were imprisoned for protecting their sources. Shield laws were passed into the Evidence Acts in 2011, but have subsequently been subject to frequent scrutiny.

The panel described Australia's shield laws as "Swiss-cheese law" full of loopholes and inconsistent interpretations such as the varied definitions of "journalist" in different jurisdictions. We also traversed the controversial law of Journalist Information Warrants which permit government agencies to access the metadata of journalists in order to identify their confidential sources.

Paul Farrell provided his valuable insights into how a journalist makes the assurance of confidentiality with sources, in circumstances where the laws do not provide any guidance into how to establish the privilege. Paul advised the keen audience that the most important step is to establish clearly with your source the terms on which you are speaking. He also stressed that you must also be fully prepared to fulfil the serious promise of confidentiality. Paul also spoke about the value of encryption and the difficulty in keeping sources confidential, including balancing the need to make a contemporaneous note of the promise with the risk that this could expose the source.

When asked about the future, the panel discussed the need for a uniform law and extending the privilege to all persons engaged in acts of journalism, rather than only those engaged in the profession or occupation of journalism.

Paul Farrell and Gina McWilliams also provided interesting insights into the controversial raids on the ABC and a News Corp journalist's home. The panel spoke about how the administration of justice often trumps journalists' privilege, and the challenges this poses to public interest journalism.

Thank you to our amazing panel and moderator, and the CAMLA YLC event sub-committee (Imogen Loxton (Senior Associate, Ashurst), Anna Glen (Legal Counsel, ABC), Nicola McLaughlin (Legal Counsel, nbn) and Lucy Hughes (Senior Legal Counsel, Stan) for pulling the event together.