“The morass is just getting … deeper and deeper and deeper”: Synthetic Media and News Integrity

Michael Davis and Monica Attard

Abstract

With the arrival of generative AI in 2022, waves of hype and hand-wringing struck the news industry. These initial responses have proved overblown, if not without foundation. The challenges and opportunities of synthetic media for news are real, if more humdrum than the hype would suggest. In this paper we draw from a two-phase qualitative study to explore how these challenges and opportunities have manifested in Australian newsrooms. We focus on: 1) How are newsrooms implementing generative AI in the production of synthetic media? 2) How do newsrooms perceive the potential impacts of synthetic media on news integrity? 3) How are perceived impacts on news integrity mediating the implementation of generative AI, particularly for synthetic media?

Industry surveys have shown that uptake of generative AI in Australian newsrooms is low relative to comparable markets. Similarly, in phase 1 we found almost no use of genAI to produce synthetic media for publication. This reflected apprehension over the limitations of generative AI and acute consciousness of threats to trust and news integrity. Phase 2 found some moderation of concern as low-risk opportunities had emerged, though applications in audience-facing content were still limited. Participants continued to express strong concerns about news integrity and audience trust.

We apply both a technological process lens and a normative lens focused on the concept of news integrity to interpret participant insights. We conclude that the limited uptake of generative AI in Australian newsrooms is driven by concerns about news integrity in a broad sense, going beyond journalistic standards to encompass the sociopolitical functions of journalism as well as concerns about continued platformisation of the media economy and an increasingly degraded information environment.

Keywords: Journalism, generative AI, synthetic media, news integrity, trust

Introduction

With the arrival of generative AI (hereafter, genAI) in 2022, waves of hype and handwringing struck the news industry. On the one hand, the technology heralded a new era of automation that would escalate production without increasing costs and deliver novel formats that would rejuvenate declining audiences. On the other, it would threaten jobs and undermine news quality. Meanwhile, increasingly sophisticated deepfakes would degrade political discourse, damage electoral integrity and accelerate the decline in public trust (Beckett & Yaseen, 2023; Ternovski et al., 2022). These polarised “utopian and dystopian portrayals” (Cools & Diakopoulos, 2024, p. 1) have proved overblown (Simon et al., 2023), if not without foundation. The challenges and opportunities of synthetic media for news are real, if both more humdrum and more profound than the hype suggests.

In this paper we draw from a two-phase empirical study into the impact of genAI on public-interest journalism in Australia to explore how these challenges and opportunities are being negotiated in Australian newsrooms.

AI has been making its way into news output for nearly two decades, mostly through automated reporting from structured data sources (Bäck et al., 2019, p. 11). In many cases, as Simon and Isaza-Ibarra (2023, p. 8) observe, "the technology has slowly moved into news production and distribution, often without readers (or journalists) really noticing." Discourse was polarised even in these early days, with one side championing the potential transformation of news production through technological innovation, and the other focused on industry disruption, particularly threats to journalists' jobs, as in discussions of "robot journalism" (Lindén & Dierickx, 2019).

In part, the polarised discourse surrounding AI must be understood in the "larger context of the digitization of media and public life" which has transformed journalism, "undercutting business models, upending work routines, and unleashing a flood of information alternatives to news" (Lewis, in Broussard et al., 2019). Questions about ethical practice, news values and journalistic purpose are also never far from mind. As Moran and Shaikh attest, debates within newsrooms about technology are embedded in broader conversations about the role and efficacy of journalism, and about where its boundaries lie. These centre on the question of how technologies "advance or hinder a particular normative vision for journalism" (R. E. Moran & Shaikh, 2022, p. 1757).

With the emergence of genAI, both the technological and the normative questions have been amplified. The potential transformation of production is seen as more significant than earlier AI technologies; but so too is the potential impact on news as an industry and sociopolitical institution. On one hand AI represents "the next level" of technical sophistication. On the other, AI is "fraught with myths, political connotations and emotional responses that stand in the way of an informed debate on AI, within and outside newsrooms" (Helberger et al., 2022, p. 1606).

In our study we find deep engagement within the news industry with both the technological and normative questions. Indeed, we find that the implementation of genAI in newsrooms is mediated largely by concerns about ethical practice and the sociopolitical functions of journalism, though resource limitations also play a role. For this reason, we apply both a technological process lens and a normative lens to investigate the implications of AI-generated synthetic media for the integrity of news. GenAI's technical capabilities and limitations are inseparable from normative questions about the desirability of its application in news, and so in examining both, we can build a fuller picture than by, for example, applying a classical technology acceptance model (Bagozzi, 2007), or an ethical analysis divorced from the economic and labour imperatives driving technological adoption.

While the normative lens considers common journalistic standards such as accuracy and fairness, our interviews reveal journalists are thinking about AI-generated synthetic media more broadly by framing it in terms of the sociopolitical functions of public-interest journalism and its critical importance in an increasingly degraded information environment. In this paper, we employ the concept of *news integrity* to capture this broader lens.

Given the novelty of generative AI, research into newsroom implementation is only emerging. In the Australian context, studies are limited. A report on the first phase of our research at the UTS Centre for Media Transition was the first comprehensive study of newsroom implementation in Australia (Attard et al., 2023). Thomson et al. (2024) observe the impact of genAI on visual journalism in seven countries including Australia, while a report from RMIT provides insights into audience as well as journalist perceptions of AI (Thomson et al., 2025).

The scholarly contribution of this study is not limited to Australia, however. Despite national differences in approaches to implementation and in industry and sociopolitical context, newsrooms worldwide face the same issues of news integrity as those revealed in our research.

Literature Review

Synthetic media

Whittaker et al. (2020, p. 91) define *synthetic media* as "all automatically and artificially generated or manipulated media," including but not limited to synthesised audio, virtual reality, and advanced digital-image creation. Squicciarini et al. (2024, pp. 15–16), use *synthetic content* to cover a similar range, defined as "digital output generated or modified by algorithms, typically AI techniques, such as machine learning," including text, audio, imagery or multimedia. They use *synthetic media* to refer to a subset of synthetic content that is intended for or available to audiences, with *deepfakes* a further subset of synthetic media. Similarly, Martin and Newell (2024, p. 448) refer to synthetic media as "synthetic outputs ... that are often (though not always) produced by generative AI systems and intended for people to consume," with *AI slop* referring to low-quality synthetic media. In general, synthetic media encompasses but is not exhausted by the new wave of generative AI technologies including GPTs (He & Fang, 2024, pp. 40–43), though they are the focus of our study.

Harris (2024, p. 131) observes that the distinction between synthetic and non-synthetic or human-produced media is not a "clear binary." Rather, genAI output could be thought to exist on a spectrum from lightly modified to fully synthesised. Barnes & Barraclough (2020, p. 214) note that most types and uses of synthetic media are benign. However, deepfakes, by their mere existence, "cast a shadow on the veracity of any given audiovisual record."

The terms *synthetic content* and *synthetic media* arise from, and remain primarily associated, with visual and auditory media rather than journalism (Feher, 2024, p. 353; Schell, 2024, p. 19). We found them very infrequently used amongst our interview subjects, who preferred *AI-generated news* or *content*. In scholarship, *automated* or *robot journalism* are common, as noted above. There is some use in industry-wide guidelines, however, such as the *Paris Charter on AI and Journalism* (2023) and the Partnership on AI's *Responsible Practices for Synthetic Media*, which goes beyond journalism to include other synthetic media (Partnership on AI, 2023, p. 3).

For newsrooms, the distinction between the use of genAI to create audience-facing synthetic content (or synthetic media, in Squicciarini et al.'s and Martin and Newell's terminology) and internal-only uses is critical. But *synthetic content* is an umbrella term that includes audience-facing content. Accordingly, throughout this paper we specify whether the use, content or media is *audience-facing* or *front-end* (as commonly used by our interviewees) or *internal* or *back-end* where it is not. For newsrooms, there is also an important distinction between synthetic media or content produced internally (whether for back-end or audience-facing uses) and *externally sourced synthetic media* used in a news story, such as a video of a breaking news event circulating on social media.

AI adoption in newsrooms

Surveys conducted since the arrival of ChatGPT reveal accelerating AI implementation in newsrooms globally. A 2023 survey found that almost half of newsrooms were actively working with genAI, though use was infrequent and confined to a small number of users (Roper et al., 2023, pp. 5–6). By 2025, a Thomson Reuters survey found 49 percent of journalists used AI daily (Radcliffe, 2025, p. 17). Usage in Australia is markedly lower. A 2025 report found 63 percent of journalists had not used genAI in their work during the previous year (Medianet, 2025, p. 48). Eighty-eight percent reported concern about the potential effects of genAI on the overall integrity or quality of journalism. Our study explores newsroom perspectives to identify the drivers of this limited adoption.

Globally, experimentation has mostly aimed at making newsroom workflows more efficient and scalable (Cools & Diakopoulos, 2024, p. 12). This includes by automating routine tasks—often those made necessary by other forms of technology, such as metatagging (Diakopoulos et al., 2024, p. 16)—or by augmenting human capabilities, for example, in large-scale data and document analysis (Radcliffe, 2025, p. 14). But it also includes both internal and audience-facing content creation and other editorial tasks. The Thomson Reuters survey found 30 percent of journalists used genAI for text creation and 21 percent for multimedia creation (Radcliffe, 2025, p. 18). A study by Møller and others (2025, p. 14) found that content creation has the lowest perceived potential for journalistic applications of genAI, with the highest in information analysis and content optimisation (e.g. SEO).

There have been some notable attempts at full article generation using genAI, with what might generously be called mixed results (Farhi, 2023; Mahadevan, 2025). Some newsrooms have developed AI-generated newsreaders, complete with social-media profiles (Samosir, 2023). But many outlets are using genAI mainly for internal content-manipulation tasks like summarisation, transcription or information synthesis (Diakopoulos et al., 2024, p. 11). Others are experimenting with limited audience-facing content creation, subject to editorial scrutiny before publication. This includes headlines, social-media posts, article summaries, translations, data visualisations, and synthetic voice (Borchardt et al., 2024, pp. 66, 70, 86, 97, 131). As Simon notes, limitations arise from the complexity of the newsgathering, production and distribution processes, which is "messy and unpredictable" rather than "an assembly line of neatly defined components" that can be easily or fully automated (Simon, 2024, p. 20). For this reason, it can be difficult to blend automation tools into existing workflows (Gutierrez Lopez et al., 2023, p. 485).

To mitigate some of the weaknesses of consumer AI tools—including inaccuracy, hallucination, bias, and generic output, as well as legal and intellectual property concerns—well-resourced newsrooms have moved to develop customised, in-house AI models (Simon & Isaza-Ibarra, 2023, p. 10). One third of respondents to a global March 2025 survey reported their organisations were using AI tools trained on their own content (Center for News, Technology & Innovation, 2025, p. 38). These include archival search tools or proofreading tools trained on internal style guides (Borchardt et al., 2024, p. 74). Several newsrooms have incorporated audience-facing chatbots into their websites (Oliver, 2024; WashPostPR, 2024).

A long-running narrative accompanying moves towards automation is that it will free up journalists to do high-value work such as lengthy investigations and analysis (Meir, 2015; Tran, 2006). But widespread experimentation with genAI has so far yielded a relatively narrow range of beneficial uses, mostly in back-end, rather than audience-facing, production tasks (Cools & Diakopoulos, 2024, p. 11; Simon, 2024, p. 18). This is partly due to information-integrity problems such as inaccuracy and bias, but also a lack of news value in much AI output, including oversimplification, failure to highlight newsworthy information, or homogenisation of news content (Cools & Diakopoulos, 2024, p. 12). The work required to produce quality, newsworthy output—or to compile and edit low-quality output—may outweigh any potential time saving (Diakopoulos et al., 2024, p. 20; Radcliffe, 2025, p. 22; Simon, 2024, pp. 18–19; Simon & Isaza-Ibarra, 2023, p. 10). This is particularly the case with off-the-shelf products. But developing AI in house is very resource intensive, potentially for only modest productivity gains (Simon & Isaza-Ibarra, 2023, p. 10).

Some therefore question whether genAI heralds a new era of innovation, or is just another in a long line of hyped technologies accessible only to well-resourced newsrooms, leaving local outfits, and many in the global south, at a disadvantage (Ferrucci & Perreault, 2021; Min & Fink, 2021). While the accessibility of consumer AI has democratised the technology in newsrooms (Cools & Diakopoulos, 2024, p. 1), it is less useful than custom products, the cost of which may put them out of reach of many. Amid ongoing pressure to produce more content to satisfy the digital market, another key question is whether AI will merely facilitate a rise in low-quality content or "churnalism" rather than freeing up capacity for high-quality journalism (Simon, 2024, p. 19; Montaña-Niño, 2024, pp. 30–31; Golding & Murdock, 2022, p. 40).

News integrity

Concepts of integrity—including news, journalistic, editorial, media and more broadly, information integrity—are increasingly invoked in both industry and scholarly discourse amidst rising concern about the degradation of the online information ecosystem and the impact of digital platforms on news. Despite their popularity and broad application beyond these recent concerns, concepts of news integrity remain undertheorised. Here we tease out some essential elements of these concepts before looking in the next section at their relevance to genAI and synthetic media.

Integrity is often invoked in discussion of journalistic practice as a commitment to shared ideals and to the structures and practices that have evolved to promote them. As Borden and Tew (2007, p. 302) observe, "When journalists present news in a way that distorts the truth, their performance is at odds with the commitment to truthfulness that their role substantively requires." This normative commitment is what most clearly distinguishes journalism from other activities in the media marketplace (Borden & Tew, 2007, p. 303). Similarly, for Kieran (1998, p. 23) to accuse a journalist of bias, for example, "is to impugn his journalistic integrity in the deepest possible sense"—to claim that he is, "intentionally or otherwise, not adhering to the truth-respecting methods required for him to achieve the proper goal of journalism: arriving at the truth of the matter."

Both Borden and Tew and Kieran tie journalistic integrity to the sociopolitical functions of journalism. For the former, the proper goal of truth arises from the democratic function of the media as an "unofficial fourth estate." For Borden and Tew (2007, p. 303), journalistic standards of reliability, truthfulness, and independence, pursued through "a discipline of verification," provide an "epistemologically defensible" framework for creating and communicating knowledge that, ultimately, helps citizens participate meaningfully in the public sphere. Similarly, the SPJ Code of Ethics in the US founds the concepts of independence and integrity on the "highest and primary obligation of ethical journalism," which is to "serve the public" (Society of Professional Journalists, 2014). Many newsrooms' editorial policies also make explicit reference to integrity and its relation to serving the public interest and preserving trust (Riordan, 2014).

There is variety in how the public-service or public-interest value of news is articulated, but it typically includes what Hall (2025, p. 101) calls the three core democratic functions of news, namely: informing the public about public-interest issues, holding power to account, and providing a forum for public debate. Public trust in the news depends on the perception that these functions have not been undermined, for example by poor practice or by commercial or political pressures. Thus, the Peace Institute's *Media Integrity Matters* report (Petković & Mihajlović Trbovc, 2014, pp. 21–22) conceives of media integrity as encompassing the policies, structures and practices which "enable the media to serve the public interest and democratic processes," by providing "accurate and reliable information to citizens" and ensuring that citizens "have access to and are able to express a wide range of views and opinions without being exposed to bias and propaganda." Where these structures and practices are weak, the public can no longer trust that the news they read is accurate, reliable and free from bias or the influence of vested interests.

Hence, integrity relates not only to journalistic practice but also to the news itself. News that is produced with integrity has integrity in turn. Public trust thus extends not only to particular outlets but to the news they produce. As the Australian Press Council puts it, adherence to professional standards is vital to "the confidence consumers have in the integrity of news material being reported to them" (Australian Press Council, 2023, p. 3). One might say that the purpose of journalistic integrity as an integrity of process is to ensure news integrity as an integrity of product.

Because integrity is a feature of the product as well as the process, it can be undermined at any stage of news production and distribution, including after the news has been published and is no longer under the newsroom's control. This underlies longstanding concerns about the distribution of news on digital platforms, where, for instance, the integrity and continuity of a publisher's broader coverage can be lost in the torrent of atomised content (Wilding et al., 2018, p. 37). The structures and practices that maintain news integrity can also be weakened by market forces, for example, the loss of advertising revenues to digital platforms and diminishing consumer demand for traditional news, and consequent reductions in the journalistic workforce or in news coverage (Australian Competition and Consumer Commission, 2019, pp. 309–322). These industry effects bring into play the concept of information integrity (International Panel on the Information Environment et al., 2025, pp. 7–10), and the ability of the news industry to deliver public-interest news within a broader information environment in which it is playing a weakened role and over which it has limited control.

AI and news integrity

Emerging empirical studies of journalists' perspectives on genAI (Cools & Diakopoulos, 2024; Thomson et al., 2024), of metajournalistic discourse (Ananny & Karr, 2025) and emerging codes and guidelines on AI use (Becker et al., 2025; de-Lima-Santos et al., 2024) show that genAI adoption is attended by strong concerns about its potential impact on news integrity and public trust. Indeed, our participants regularly raised concerns about news integrity, suggesting that these concerns are mediating and constraining uptake in Australian newsrooms. Even where integrity is not explicitly invoked, we contend that concepts of integrity provide a useful framework for understanding how genAI technologies are being adopted and the attitudes of industry members towards them, encompassing journalistic practice, audience trust, and the sociopolitical functions of news.

Editorial standards

One of the leading concerns about genAI in news is the potential for newsrooms to unintentionally propagate the inaccuracies, hallucinations and bias that are notorious features of much genAI output (Jones et al., 2023, p. 4). This clearly threatens the integrity of news as a product, requiring an integrity of process to mitigate it. Thus, Zier and Diakopoulos (2024, p. 1) argue that careful editorial oversight is required to preserve journalistic integrity and the integrity of news or information. While Cools and Diakopoulos (2024, p. 5) focus on the importance of ethical principles, they frame these in terms of integrity, arguing that ethical principles can serve as "a compass for preserving the integrity of journalistic practices" as AI is implemented in news workflows.

GenAI is accompanied by concerns about loss of editorial control within the newsroom. For example, there is concern that eagerness to experiment might override ethical practice, particularly where it is driven by management (Gutiérrez-Caneda et al., 2024, p. 4; Møller et al., 2025, p. 16). There are also worries about the robustness of oversight measures, given the opacity of AI systems (Cools & Koliska, 2024, p. 666; Jones et al., 2022, p. 1736), a lack of AI literacy (Cools & Diakopoulos, 2024, p. 13; Jones et al., 2023, p. 4) and the pressures of the digital news cycle, which had already strained traditional verification processes before the advent of genAI (Hermida, 2015, pp. 39–41).

Editorial control in the digital information ecosystem

While the onus is on news publishers to ensure the integrity of the news they publish, once the news moves into the broader information ecosystem they can no longer do so and must instead rely on third parties that make use of that news to maintain its integrity. This includes AI systems that use the news as data for training or grounding generative models. Examples abound of genAI tools misrepresenting, misattributing or even hallucinating news stories (C. Moran, 2023). *News Integrity in the Age of AI* (European Broadcasting Union, 2025), a joint initiative of the European Broadcasting Union (EBU) and WAN-IFRA, responds to these issues, proposing five principles to "counter the misinformation crisis and protect the value of trusted news." These include requiring authorisation,

attribution and accuracy for news content in genAI models; fair recognition of the value of up-to-date, high-quality news; and ensuring AI harnesses the diversity or plurality of the news media, presumably by not limiting deals to powerful media organisations. Principle 10 of the *Paris Charter on AI and Journalism* offers a similar prescription, requiring access to journalistic content to be "conditional on respect for the integrity of the information and the fundamental principles of journalistic ethics."

Authenticity and trust

Issues of authenticity and trust arise even where oversight processes are robust. As Mike Ananny states, once synthetic content is incorporated into news, "we can't necessarily know if the news that we're reading was made by humans or made by machine learning models, or made by some mixture of those two things" (Avishai, 2023). Moran and Shaikh (2022, pp. 1766–1767) suggest concerns about authenticity rely on unquestioned assumptions about what "real journalism" is. But audiences value authenticity (Jones et al., 2023, pp. 4, 8; Wintterlin et al., 2020, p. 230), and, as Jones et al. (2023, p. 4) observe, audience expectations concerning authenticity and journalistic integrity are strongly linked to human creation of news. Yet studies have also found people view AI-generated text and chatbots as more objective and credible than humans (Lin & Lewis, 2022, p. 1635; Salas et al., 2023), while others have found transparency over AI use can decrease trust (Toff & Simon, 2024). That is, AI-generated content is credible, but paradoxically, journalists producing it are not.

Meanwhile, the proliferation of synthetic media on digital platforms, including deepfakes and AI slop, raises concerns about the capability and capacity of newsrooms to verify externally sourced material, particularly images, video and audio (Thomson et al., 2024, pp. 11–12), threatening information integrity and public trust (Cazzamatta & Sarısakaloğlu, 2025, p. 3) and causing collateral damage to news in the form of a liar's dividend (Chesney & Citron, 2018, p. 1758).

Economic impacts on news integrity

The potential impact of genAI on news integrity must therefore be understood against the backdrop of the broader media economy and the shift in journalism's place within it. Several emerging studies of genAI in newsrooms explore industry concerns over broader political and economic factors, or interpret these through a political economic lens. For example, Borchardt et al. (2024, pp. 23–24) highlight fears that as more users access news through chatbots, AI will exacerbate the problem of news visibility in atomised platform environments and further threaten revenue (Dodds et al., 2025, p. 6). Others see increasing dependence on technology companies for news production and distribution as an ongoing process of infrastructure capture that undermines journalistic autonomy (Simon, 2022, p. 1833; Sjøvaag, 2024, p. 247), especially as it is in many cases the same digital platform companies that are playing an outsized role in AI (Dodds et al., 2025, p. 6). Relatedly, discourses of efficiency and of "freeing up" journalists have been interpreted as supplanting labour by stealth (Matich et al., 2025, pp. 10–11). And some argue that casting genAI systems as tools for creativity or even as autonomous undermines the moral rights of journalists and other creators on whose work the systems have been trained (Montaña-Niño, 2024, p. 31).

Concerns about news integrity thus go hand in hand with impacts on journalistic labour. Automation may increase efficiency but decrease the role of human judgement (Cools & Koliska, 2024, p. 664). Journalists' concerns about authenticity, objectivity and voice have been interpreted as a form of boundary work to preserve independence and editorial control (Ananny and Karr, 2025, p. 13). At the organisational level, this can manifest in discussions about preserving brand integrity. But journalists—at least those who see themselves as observers or watchdogs rather than as mobilisers or entertainers (Møller et al., 2025, p. 15)—take the ethical implications of AI seriously because they take the sociopolitical functions of journalism seriously.

These considerations suggest that a broad view of news integrity—encompassing ethical journalistic practice, the sociopolitical functions of news, the impacts of the media economy and the relations between news and the broader information environment—is necessary for a comprehensive understanding of how newsrooms are implementing and responding to genAI. Taking these considerations into account, our research questions are as follows:

RQ1) How are newsrooms implementing genAI in the production of synthetic content?

RQ2) How do newsrooms perceive the potential impacts of genAI on news integrity?

RQ3) How are perceived impacts on news integrity mediating the implementation of genAI?

Methodology

In the study's first phase, conducted from July–October 2023, we interviewed 11 newsroom editors and one product lead from eight Australian media organisations. In the second phase, conducted from August–November 2024, we interviewed 13 news editors and six product leads from 14 news organisations, including the majority of those from phase one. In November 2024, we held a day-long workshop attended by many of those we interviewed, as well as some we did not. A table of participants is provided in Appendix A.

The study population was defined using criterion-based expert sampling (Etikan, 2016, p. 2), on the basis of expertise in newsroom editorial management or product development, and involvement in AI implementation or policy development. In phase one we focused exclusively on newsrooms producing public-interest journalism or "hard news" in different markets, models and media types to achieve sample variation. Phase two was broadened to test whether implementation differed at the margins of public-interest journalism, such as in factual lifestyle content. We also looked where possible to include at least one editor and one product lead from each newsroom, as these roles represent different imperatives within a newsroom's implementation process. Participants were approached directly or via the researchers' networks and selected based on willingness to participate. Further participants were then identified using snowball sampling.

Interviews were semi-structured. A set of general questions was posed to all participants, based on a literature review conducted in mid-2023 and updated as the project progressed. Others were aimed at specific newsrooms based on their characteristics. Further questions arose from participant responses. Questions covered uses of AI; implementation processes; practical limitations; risks to news integrity, journalistic ethics, and audience trust; legal risks; and risks for the industry and the broader information environment.

Sixteen participants attended the workshop, which was conducted under the Chatham House rule to encourage discussion. The workshop was split into three sessions, focusing on: (1) use cases and implementation, particularly relating to synthetic content generation; (2) principles and guideline development; and (3) cross-industry issues including the integrity of the broader information environment, and closer collaboration between newsrooms and AI companies, particularly for the purposes of mitigating risk. Session 1 was led by colleagues undertaking research on audience perceptions of AI in journalism, while sessions 2 and 3 were led by the authors. Sessions 1 and 2 were attended by news editors, content editors and product managers. In session 3, these were joined by two representatives from AI companies and two industry consultants.

The workshop was also semi-structured in approach. To facilitate discussion, the first session included a slide presentation on genAI use cases and audience perceptions of AI use that was compiled from our colleagues' prior research. Participants discussed whether they had implemented or considered any of these uses, and where they perceived risk. We also verbally shared general themes from our interviews. For the second session, participants were provided in advance with a handout of example AI guidelines and principles drawn from guidelines published by news

organisations and industry bodies in Australia, the UK and Europe. These were sorted into categories: journalistic principles (accuracy, impartiality, etc.); transparency; human oversight and accountability; use restrictions; evaluation and testing; and organisational and legal issues such as privacy, licensing and distributing risk, responsibility and liability. The handout also included questions for discussion. The third session was informed by the discussion in the two prior sessions.

The workshops were therefore an opportunity for multidimensional knowledge transfer, with the researchers sharing findings on implementation, audience perceptions and guideline development, and participants sharing with the researchers and with each other their own practical experiences and perceptions.

Both the interviews and the workshop sessions were recorded. Recordings were transcribed and coded thematically in NVivo. Broad themes were based on our research questions, including: AI implementation and use; constraints on implementation; and perceptions of risk, particularly to news integrity. Finer-grain codes were inferred inductively from the interview and workshop data. Coded data was then analysed on the basis of a four-way classification of our participants (see Appendix A) across three organisational variables: market (national, metropolitan, or regional); medium (television, radio, hardcopy newspaper or online); and model (public, commercial, or non-profit); as well as a single personal variable: professional role (news editor, factual content editor or product manager). For broadcasting and print, market generally reflects size, with national organisations the largest and regional the smallest, though subsidiary relationships complicate this relationship. Online-only outlets are generally small but have national reach, and in two cases are backed by international organisations. The views of the AI company representatives and industry consultants who attended the third workshop session have been excluded from the present sample.

The research has undergone ethics approval at the University of Technology Sydney (ETH21-5787-24-2) and conforms with all relevant requirements and guidelines. Participants were provided with information about the purposes and conduct of the study and about data retention and use. Written consent was obtained from all participants, and participants and their organisations have been de-identified. The attached table provides relevant information on participants, including their role and details about their organisation.

As noted, this research forms part of an ongoing study and only a subset of our findings is reported here. These have been selected solely on the basis of their relevance to this special issue. As noted above, an industry-targeted research report on phase one has previously been published, and some of those findings are included here (Attard et al., 2023).

Some quotes have been edited for clarity.

Findings

In phase one (July–Oct 2023), we found participants were cautiously optimistic about the opportunities brought by genAI. However, there was trepidation over how rapidly the next wave of disruption was approaching. While all participants thought genAI would have a momentous impact on the news industry and on journalism, there was uncertainty over precisely what it would be. A mantra of “no genAI in published content” served as the default short-term safeguard, reflecting apprehension over the limitations of genAI, a reluctance to undermine journalistic output, and an acute consciousness of the threats to trust and brand integrity. There was also significant concern about how to navigate the proliferation of online synthetic media, where increasing technical sophistication and a degraded information ecosystem amplify the need for robust verification processes and undermine the ability to undertake them.

Our second-phase investigations (Aug–Nov 2024) found some moderation of concern as experimentation had identified opportunities to enhance workflow. Despite this, implementation

remained limited and experimentation was carefully controlled, with the majority of organisations focused almost exclusively on back-end productivity and efficiency gains. A small number had experimented with audience-facing synthetic content, confined to a narrow range of low-risk applications. There was notable variation across the three organisational variables of market, medium and model. Larger, national organisations, particularly the public broadcasters and commercial radio networks, had progressed much further in experimentation and implementation than smaller, regional organisations.

Two main constraints on genAI implementation and experimentation emerged across the study: (1) a perception that the utility of genAI was limited, with cost often outweighing benefits; and (2) an overriding, principled focus on the integrity of news. Both constraints were clearly apparent across all organisations. However, while news integrity was a universal concern, the cost–benefit calculus yielded different results across markets, media and models in parallel with the differing levels of implementation.

Emerging uses of synthetic media in newsrooms

In both phases one and two we found that most news organisations see the biggest opportunities for genAI in back-end functionality, particularly in news gathering and production, reflecting other studies (Cools & Diakopoulos, 2024; Diakopoulos et al., 2024; Møller et al., 2025; Radcliffe, 2025). But even in phase one, participants were contemplating deeper investigation of the opportunities of genAI, including in front-end output, and were gathering the resources to begin experimenting. Many, especially larger national and metropolitan outfits, had formed working groups comprising editorial, product development and legal personnel to manage implementation and to develop AI policy.

> We have a huge technology, product and digital team here. ... We are really trying to understand how the tech works, what we might build in house, what we might use, what we might license. (P1-09)

By phase two, all participants had established such groups, though their formality, size and progress differed according to the size of the organisation, suggesting that resources are an important factor in newsrooms' ability to manage implementation. Across the board, implementation remained mostly experimental, focusing on low-risk opportunities with potential for good returns on investment such as increased efficiencies or audience expansion.

Audience-facing content

In phase one, no participant organisations had experimented with audience-facing synthetic content, and many, seeing only downside risk, had ruled it out in the near term. However, we found differences between and within organisations according to purpose, market and brand. While they still had an eye on potential opportunities, print outlets and public broadcasters were very cautious.

> The general policy is we don't want journalists using ChatGPT for their journalism. (P1-05)

In phase two, many of these organisations were still very wary of using AI for audience-facing content. A regional newspaper (P2-06) was not contemplating genAI to create synthetic content at all, even in areas such as data visualisation. However, some had begun experimenting, chiefly in digital content rather than news, and always with human oversight. For most, the scope of application was still limited to short texts and ideation.

> So there is news, which is the pointy end, and a very ... strict approach, ... whereas in content we accept that there the audience expectation is a little bit different. (P2-15)

The online lifestyle publisher was experimenting heavily with a wide range of efficiency-focused back-end use cases, but, somewhat contrary to our expectations, was as reticent as our other participants about using AI for audience-facing content (P2-04).

Several organisations were interested in exploring chatbots and other content delivery and personalisation uses. Given resource limitations, however, it was generally a lower priority than newsgathering and production. One public broadcaster had experimented with older types of AI in 2015 to develop a chatbot to deliver news and other information, but the project had stalled (P2-13). The organisation is now testing a genAI chatbot, confined to research and back-end tasks.

Despite the relatively limited implementation of genAI to produce synthetic content across all organisations, we found increased experimentation in several distinct areas. In audience-facing content these were largely limited to synthetic voice, image generation, short-text generation like headlines or alt text, and some translation. Much more common were back-end newsgathering and production tasks including transcription, summarisation, and idea generation.

Synthetic voice

Synthetic voice has emerged as a significant opportunity across different use cases. For the public broadcasters, improvements to accessibility and representation are a particular focus, as is connecting synthetic voice with the translation capability of genAI to serve Indigenous communities and migrant language groups.

In phase one, one of Australia's largest commercial radio networks was investigating how synthetic voice could be deployed for simple information services like short weather reports that otherwise require significant time for a journalist to produce.

> We are not talking about a developing situation like a cyclone coming into Cairns. It's 26 degrees and sunny, so a very short sentence. But ... there are actually quite a lot of touchpoints. Whereas if you could automate that process, and you've got 99 radio stations, you could be saving a good couple of hours of someone's time. (P1-09)

In phase two, this network had implemented audience-facing synthetic voice in the lower-risk areas of hyper-local weather reports and fuel-price updates (P2-16). This resulted in substantial time savings—especially important for Australian commercial radio broadcasters which are required by law to provide a certain amount of local content per day (Australian Communications and Media Authority, n.d.).

Another commercial radio network had developed a multi-faceted internal tool that can source content from around the world, draft scripts for short news bulletins in the distinct house styles of the network's various stations (reflecting market demographics), and synthesise those bulletins using cloned voices of their own journalists. The tool was still in testing and had not been used to publish audience-facing news content (P2-03).

Image generation

In phase one, some were contemplating synthetic image generation, though with little official testing. In phase two, more organisations had experimented in this area. Data visualisation was an opportunity in both internal analysis and audience-facing uses. Still, all were cautious about full-scale image generation for audiences.

> We've done some internal experiments with illustration for articles, seeing that as low risk. We haven't put that in front of audiences. (W-01)

Some organisations were more liberal with non-news uses of image generation and image modification, animation or extension rather than full generation. An editor at one of the public broadcasters, however, was clear that even image extension could undermine audience trust (W-05).

Headlines, short text generation and ideation

In phase one, short text generation was mostly a perceived opportunity rather than a subject of testing. In phase two, there had been much more experimentation, though application was still limited in audience-facing uses. This was the case across different media types and markets. Using AI to analyse a large set of images and to generate alt text was a common use case. Many were also using it for headlines, but there was reluctance to push too far.

Many organisations had also found a use for genAI in ideation. Again, it was perceived by all as an assistive technology, not a substitute for human creativity.

Transcription, summarisation and translation

In phase one, many organisations saw a potential application for genAI in transcription and translation. By phase two, many had implemented AI transcription tools, mostly in internal use, and had seen real efficiency gains. Public broadcasters and other organisations which produce content across different media types see strong value in automated transcription.

One public broadcaster had developed a customised LLM, principally for transcriptions, as off-the-shelf tools were inadequate.

> The in-house one was ... trained on our own content, and it performed a lot better when it came to nouns, Australian place names, Indigenous language, etc.; whereas, you know, an off-the-shelf [tool] that's built on a global language just doesn't perform quite as well. (P2-14)

Outside our interview cohort, one Australian broadcaster has implemented a tool that repurposes human-authored TV news scripts into online news stories (9News staff, 2024).

Demonstrated time savings also saw summarisation used across all participant organisations, principally for research. Use of genAI for translation was partly dependent on market and audience. As noted, the public broadcasters were experimenting with translation and synthetic voice in languages other than English. For commercial media, however, translation was still mostly viewed as a future opportunity.

Constraints on implementation

The limited scope of AI implementation even in phase two points to strong constraints based on: (1) a lack of perceived utility and value in AI tools, particularly for those not sufficiently well-resourced to develop in-house products; and (2) concerns for audience trust and news integrity. These are not unrelated: most participants saw the limitations of genAI as directly threatening the integrity of news, and audience trust as hinging on perceptions of authenticity. Indeed, we found that the implementation of genAI is mediated largely by concern for news integrity and trust, underpinned by broader cognisance of the sociopolitical role of journalism. Labour concerns were raised, but these were also often cast as a risk to news integrity, and most editors thought them misplaced in the short term.

Lack of utility and value

Many newsrooms have so far found limited use cases for AI. In particular, few saw value in using genAI to produce synthetic content, even in back-end tasks, as the need for human oversight might

outweigh efficiency gains. This was particularly pointed for smaller teams, including those that sit within larger organisations.

> Is that really where we're going to put our time into using those tools? ... If it's basically going to mean somebody's got to go back into it, go through it, check another source, make it two times the length of time that you're looking at for that? (W-12)

> Accuracy is the key point there, and I think ... in fast-moving newsrooms or small teams where you're really conscious you don't have a lot of ... resources to go back and check things beyond the rigorous fact checks you're already doing on stories, ... then that starts to impact trust for all of us. (P2-11)

Some observed that humans provide much more value in content creation. This connected with ideas about the value of originality and authorial voice and the sense that while AI is good at stringing words or pixels together, its output has a tendency to be bland and homogeneous.

> I feel like, in terms of it generating content, and especially content that we would use, we're so far away from that just because we are the experts in that field. (W-08)

While product teams saw potential in a larger range of use cases and could meet resistance from editorial staff, they were acutely aware of what journalists need from AI tools to maintain editorial standards.

> In journalism, things need to be in certain places, and word order matters. It's far more precise than people give it credit for, when you're dealing with high-quality journalism. If you don't have those standards, you can get away with stuff. But if you do, it's going to be more work to get it to that standard. (P2-13)

Many observed a particular lack of value in consumer-level AI due to its poor accuracy, bias and other limitations. But developing internal tools is expensive and slow, even for larger organisations.

> We don't have $100 million spare ... to run around just trying all this stuff out. It takes heaps of time to do it properly and to have some faith in the integrity of the process. (W-05)

Newsrooms with more resources are able to invest in better models that reduce risk, allowing them to find more utility in AI. As well as larger budgets, larger, national organisations have large product teams and massive news archives, and tend to operate across different media. These factors both incentivise and permit greater application of resources to experimentation and implementation. Still, even the larger organisations in Australia lag behind large news organisations in Europe or the US in their level of experimentation and implementation, particularly in audience-facing uses.

Concerns about news integrity

In both phases, the constraints on experimentation and almost complete preclusion of audience-facing synthetic media content reflected deep concern over the potential for AI to undermine news integrity and audience trust, which mediated how participant organisations were thinking about potential uses and the safeguards needed to govern implementation. This was true across all participant organisations.

> For us, the key thing that will be top of mind every step of the way is safeguarding and retaining trust. And then that being at the centre of every decision we take in relation to genAI, but that not being at the cost of potential efficiencies or things that could actually help the audience. (P1-05)

> Integrity is so important to the journalism that we do. So, I struggle to see, as far as our storytelling goes, that we will be doing much with it for a little while, just because we're not ready. Integrity is very, very important. However, I think it would be very unwise to ignore it, as well. (P1-11)

Concepts of integrity and trust were deeply rooted in organisational culture, reflected in the observations of both product leads and editors. The lifestyle publisher, who had ruled out audience-facing content despite heavy experimentation in back-end uses, also reached for news integrity to explain this choice, noting they were members of the Australian Press Council and abided by its code of practice and consumer complaints scheme (P2-04).

Amongst all parrticipants, there was greater willingness to experiment and adopt where trust was not threatened, as in low-risk, back-end applications.

> The biggest takeaway for me is how audience trust is embedded in everything we do. People see the availability of AI tools that will build efficiencies in their workflow, but we also need to consider if this would have an impact on audiences. Ultimately, we want to use AI in a way that enhances our services and maintains editorial standards so audiences know they can trust all our content. (P2-18)

Most felt that, while risks need to be considered, fundamental journalistic standards—accuracy, impartiality, fairness and independence—remain unchanged. Instead, what is needed is guidance on how they apply to AI. Yet, while guidelines help in making decisions about use, what came through as most important was clarity about deeper principles and purposes.

> When you parse the various principles and guidance that’re being published by others, I think when you are deep into them and you can see the code that people are speaking, what sort of external references they're acknowledging, what kind of imperatives they're acknowledging, it's more useful. (W-05)

Authenticity and transparency

Often the conversation turned to questions of authenticity, including the risk that AI use would blur the boundaries between reality and representation. This was a concern even in non-news content. One participant (P2-12) noted they would never use AI to expand an image, “because if the photograph is documentary in nature, then you don't know what was beyond here, and you can't pretend to the audience that you did.”

Most participants insisted on the importance of transparency in maintaining trust. News integrity was at the heart of these concerns.

> I think it's inevitable that more AI tools are going to be adopted in journalism, but it's absolutely essential that we are upfront about our use of that, and that we communicate with our audiences about that. And because of the importance of trust in news organisations ... the audience has to know. They have to have have evidence and faith in the fact that if a news organisation is using Gen AI, ... they will tell you the ways in which they are using it and still

> guarantee the quality of the journalism. So I think, you know, slipping it in under the radar is not where the news industry should be going. (P2-01)

For regional newsrooms, notions of authenticity and transparency also tied back in to the value of journalists on the ground gathering news in local communities.

> If we're going to sort of get through this journey with genAI and still continue to provide what I think is an essential service to democracy and to anyone who lives in a democracy, then we've got to bring people with us, and they need to be able to distinguish between professional news outlets providing independent, reliable information that's been fact checked by, you know, professional journalists living in their community and understanding how things work, and a bot spitting something out based on an algorithm. (P2-01)

But given the "black box" nature of AI systems, some felt that it would not always be possible to explain to audiences how newsrooms were using AI, and thus to maintain trust in the integrity of their product.

> I can't tell you how those tools are actually working. So the explainability thing to me is a really big one, if or when we go down that path of audiences interacting with AI products and being able to explain to them really in ways that they would understand. I think that's going to be a real hurdle. (W-05)

External information pollution

In both phases, news integrity, authenticity and trust also arose as acute concerns in discussions about the effects of AI on the integrity of the broader information ecosystem, as an area that largely lies outside newsroom control. This was often tied to the potential for AI to pollute the information environment through spurious "pink slime" or misinformation and to generate a liar's dividend.

> I am more concerned about the dangers it poses for the news ecosystem as a whole. That's my major concern... (P1-01)

> If there's a whole lot of bullshit out there generated by AI, then that shifts the entire landscape into bullshit. (P2-01)

The flipside of this was the potential for quality news to become increasingly valuable in a degraded information ecosystem. But others felt that even if audiences seek out trusted news, it will be increasingly difficult to find in an atomised media environment.

> Outlets like [ours] have the ability to stand out in coming years. ... You know you can trust every single thing we say because we've done the legwork to establish and confirm what's occurred. (P1-05)

> As information sources splinter, the morass people will have to wade through to try to find reliable information is just getting deeper and deeper and deeper. (P1-01)

Verification

Likewise, nearly all participants were concerned about the pressure that external synthetic content would put on newsroom verification processes. Unlike many other concerns that accompanied the arrival of genAI, this had not moderated as we moved to phase two. All agreed that newsrooms need to be extra-vigilant, but many were concerned that verification processes, no matter how rigorously

applied, are not always effective, particularly given the spectrum of synthetic content ranges from cheap fakes to sophisticated deep fakes and subtly altered content.

> The possibilities of AI hoodwinking the media are now limitless and the fakes are coming at us all the time. (P1-10)

With prominent local cases of poor processes leading to the publication of inaccurate and offensive content (Dunstan & Ortolan, 2024), others were concerned about standards slipping due to the competitive pressure of the news environment.

For those working in regional newsrooms, however, verification was less of a problem, as their original news stories mostly feature people familiar to the journalists or present in their archives.

Workforce risks

The final areas of concern focused on market risks to news integrity, including job losses and fair use of news content. Some reported substantial concern about job losses from junior journalists; however, at the senior editorial level, all insisted that threats to news integrity and trust ruled out replacing journalists or fundamental reporting tasks. Rather than replacement, participants were thinking about augmentation. This was true even in radio, where the success of synthetic voice heightens the perceived threat to jobs. Similarly, on the lifestyle end of the news spectrum, where there is potentially more leeway to explore AI, we found strong commitment to improving workflows rather than reducing staff costs (P2-04).

> There's a lot of anxiety. But then when you get people using it, and they realise that it's got limitations and that it doesn't necessarily replace them, but can help them. Then you get eyes lighting up. (P2-13)

Regional outlets foresaw that AI might lead to some replacement of human work on time-consuming but low-value tasks, such as churning out stories from wire services or press releases in metropolitan newsrooms, but believed it could never replace the value of on-the-ground reporting (P2-03). However, some had heard concerning views from management about the potential to reduce headcount and needed to insist on the importance of maintaining journalist numbers to cover public-interest news. This touches again on the potential for revenue pressures to lead to more automation and a relaxing of editorial oversight.

Platform power

Many participants felt that the biggest threat AI poses to news integrity is not hallucinations or bias; though these are certainly of concern, they are mostly within newsrooms' editorial control. Rather, the biggest threat lay in the potential undermining of traffic as tech platforms increasingly sequester users inside "walled gardens" built on information scraped from news sites.

This concern was apparent already in phase one but had grown by phase two. Some product leads were eager for deals or other forms of collaboration to gain access to high-powered custom AI tools, as OpenAI was reportedly doing with newsrooms internationally (P2-13). And a few felt striking deals with AI companies offered an opportunity for news media to monetise their quality content, but a common view was that, in the long term, deals would only undermine the news business (P2-10). Many newsrooms were blocking AI scrapers, although there was a pervasive feeling that the horse has already bolted. For small newsrooms in particular, the power of tech companies means there is a significant bargaining imbalance, and there is concern that market developments would favour the larger news companies.

> A lot of these things sound great in theory, but actually in practice, they're really, really difficult for small and medium sized publishers, you just don't get in the door. (P2-04)

Some argued that while there seemed to be a great deal of public concern about whether news media would use AI responsibly, there needs to be greater discussion about the responsibilities of tech companies. Again, discussion of these issues was driven by concern for news integrity.

> I think our industry needs to behave responsibly with respect to AI, but it's also a challenge across the tech titan ecosystem, and I think we're a ways away from that. (W-09)

Discussion

Returning to our first research question (RQ1), our findings show profound caution in Australian newsrooms, reflecting a recent industry survey that revealed relatively low AI adoption rates amongst Australian journalists (Medianet, 2025, p. 48). The vociferous concern we saw in the first phase of our research had moderated 12 months later, and experimentation had increased. However, the scope of implementation remained tightly governed and relatively narrow, focused on increasing the efficiency of back-end tasks such as transcription and summarisation, reflecting the findings of other recent studies (Cools & Diakopoulos, 2024, p. 12; Diakopoulos et al., 2024, p. 16). There had been very little experimentation with AI-augmented personalisation or delivery such as chatbots or even article summaries. Even within this narrow scope, there were questions about whether efficiency gains were outweighed by the need for continual verification and oversight. Very few organisations were experimenting with audience-facing synthetic content.

While implementation was relatively constrained across all our participant organisations, there were notable differences across the variables of market (national, metropolitan and regional), medium (print, TV, radio, online) and model (public, commercial, or non-profit). In print and online news outlets, use of genAI for audience-facing content is virtually non-existent, even amongst metropolitan and national outlets. Experimentation was also least advanced, particularly amongst those serving regional markets. Implementation is more advanced in radio, with synthetic voice emerging as the most likely audience-facing use case in Australia in the near term. This reflects a lengthy history of synthetic voice experimentation in radio (Furtáková & Janáčková, 2023, p. 95).

The large commercial radio networks, in particular, were well ahead of other participants in their willingess to test audience-facing uses. One had implemented synthetic voice for service information such as regional weather and fuel-price updates, and one had developed an end-to-end tool to search the web, script and synthesise news bulletins using synthetic voice, though it had yet to roll it out. There is still significant wariness about uses that might impact jobs or audience trust.

The public broadcasters, with a national reach across television and radio and large product teams, had also engaged in significant experimentation, with the prospect of some audience-facing uses on the horizon. However, these were focused on factual content rather than news, which had stricter parameters, and on serving linguistically diverse communities—a reflection of their public-service obligations.

These significant differences in implementation—within nationally low uptake rates—reflect variable resourcing and distinct organisational purposes. Smaller, regional print outlets have more constrained finances and very small product teams. They see few beneficial front-end use cases, highlighting the expertise of journalists in newsgathering and reporting, particularly on local issues. However, the low rate of adoption means that regional news organisations risk falling behind industry developments, potentially exacerbating sustainability concerns as audiences increasingly move online (Eder & Sjøvaag, 2025).

National and metropolitan publications, while better resourced than regional outlets, also have relatively small product teams and constrained finances, with a stronger focus on national and international coverage and investigative reporting. These outlets saw opportunities to optimise a variety of back-end tasks, including summarisation, transcription, and data analysis, but little opportunity for front-end production outside of data visualisation.

For regional commercial radio, AI is seen as an opportunity to deliver on its regulatory obligations to broadcast local content while minimising labour costs. And for metropolitan radio, AI was seen as an easy opportunity to synthesise press releases for broadcast, though none had yet put this into practice. We also observed a distinction between these and publicly funded organisations with legislated public-service mandates and commercial outlets, with the latter experimenting widely across back-end tasks but little in audience-facing content.

Looking at our second (RQ2) and third (RQ3) research questions, we found very little variation across our participant organisations. AI adoption was mediated in all newsrooms by concerns over the potential impact on brand integrity and audience trust should journalistic processes break down. Journalistic standards were thus seen as critical to counteract the perils of genAI (Cools & Diakopoulos, 2024). But the need for continual and robust oversight to ensure the integrity of the news product was often perceived as a drain on resources with little benefit, thus reducing the perceived utility of AI and constraining implementation. The lack of control over the functioning of AI tools exacerbated these concerns. Organisations with more advanced roll out of genAI were as concerned about news integrity as others, but better resourcing or the nature of the market or medium had opened up a greater range of cases which were perceived to be low risk, even in some audience-facing areas such as synthetic voice to generate service content.

Participants' reflections on the importance of ethical practice were also couched in an awareness of broader and deeper threats. They were sensitive to the reckless disruption of the media economy and information ecosystem by powerful AI companies, and the tension between securing deals, maintaining independence, and managing declining revenues, which might increase pressure for automation in editorial workflows (Borchardt et al., 2024, pp. 23–24; Simon, 2022, p. 1833; Sjøvaag, 2024, p. 247). Further, editors felt there to be a critical need for rigorous, original journalism, particularly in an environment polluted by misinformation and fraught political discourse, to counteract the blurring of the boundaries between reality and representation (Chesney & Citron, 2018; Montaña-Niño, 2024). While senior staff echo the long-running discourse around freeing journalists from the mundanity of digital workflows (Matich et al., 2025; Meir, 2015; Tran, 2006), they recognise that revenue and management pressures could weaken the safeguards against job losses, raised through concerns about risks to brand and news integrity. Perhaps the strongest of our participants' concerns, however, was the scraping of freely available but costly news to service the training and retrieval needs of AI platforms, which seem increasingly likely to undercut the public market for the news they have taken without compensation or attribution. Further, the inherent flaws of generative AI tools undermine the values of accuracy and reliability that underpin public trust in the news and sustain the industry that produces it.

In this sense, our study validates the utility of a broad conception of news integrity that encompasses both internal journalistic processes and adherence to editorial standards—what we have called process integrity—and the integrity of news as a product once it has been published into the information ecosystem—or what we have called product integrity. By including the latter aspects of news integrity we recognise that the ability of news to fulfil its democratic functions depends not only on journalistic process and editorial standards but also on external factors that are largely outside a newsroom's control. Even where news is produced to the highest standards, its integrity may be threatened as it is ingested as data for AI training and grounding and synthesised into generative output.

Despite these concerns, many of our participants were optimistic about the assistive opportunities of AI, augmenting workflows, facilitating time-consuming tasks and opening up new possibilities of analysis, ideation and even content creation—suitably constrained, of course, by editorial safeguards. The larger and better-resourced, in particular, are certainly experimenting and alive to AI's transformative potential, reflecting that "AI-infused journalism will be better and worse simultaneously, and in ways that only vaguely come into view as we see generative AI's early sprouts" (Dodds et al., 2025, p. 5).

Conclusion

In this study we reported on two phases of ongoing research into the implementation of generative AI in Australian newsrooms. We found that deep concerns over the integrity of news are driving relatively low adoption rates, constraining experimentation and potentially limiting the uptake of opportunities observed in overseas organisations. These concerns were apparent across all our participant organisations, suggesting that it is a significant constraint on implementation, in the Australian context at least. Variations in implementation rates were explained thus not by greater or lesser concern for the integrity of news, but largely by differences in market and resourcing, as well as the demands and opportunities brought by different media types and business models. All organisations were sensitive to the need to maintain audience trust and not undercut their own sustainability. Equally, however, they were concerned about the potential for genAI to threaten the integrity of news in areas outside their control. News companies perceive technological adoption as an additional strain on already-limited resources. But faced with the potential for that very technology to undermine the sustainability of the industry by pulling audiences away from news—even as it uses news to sustain itself—their concern does not seem misplaced.

No newsroom, on our count, is about to lay waste to the integrity of their product through reckless adoption of genAI. That is not to say that the threats are not there. There will no doubt be occasional acts of error and negligence. But we should not let that distract us from the ongoing undermining of the news and information ecosystem that may see us all end up in the same deepening morass.

Acknowledgements

Thanks to our capable research assistants on this project, Miguel D'Souza and Tamara Markus, and to the anonymous reviewers for their very helpful suggestions.

Funding

This research was funded by the Minderoo Foundation.

References

9News staff. (2024, March 7). *9ExPress*. 9News. https://www.9news.com.au/technology/9express/16480c33-636a-461f-9c4f-d0e2522c722a

Ananny, M., & Karr, J. (2025). How Media Unions Stabilize Technological Hype: Tracing Organized Journalism's Discursive Constructions of Generative Artificial Intelligence. *Digital Journalism*. https://doi.org/10.1080/21670811.2025.2454516

Attard, M., Davis, M., & Main, L. (2023). *Gen AI and Journalism*. UTS Centre for Media Transition. https://doi.org/10.6084/m9.figshare.24751881.v3

Australian Communications and Media Authority. (n.d.). *Local content on regional commercial radio*. Retrieved April 24, 2025, from https://www.acma.gov.au/local-content-regional-commercial-radio

Australian Competition and Consumer Commission. (2019). *Digital Platforms Inquiry Final Report*. https://www.accc.gov.au/about-us/publications/digital-platforms-inquiry-final-report

Australian Press Council. (2023, August). *Submission to the Department of Industry, Transport, Regional Development and Communications on the Exposure Draft of the Communications Legislation Amendment (Combatting Misinformation and Disinformation) Bill 2023 (sub. E3250)*. https://www.infrastructure.gov.au/sites/default/files/documents/acma2023-e3250-australian-press-council.pdf

Avishai, T. (Director). (2023, December 4). Synthetic Media: AI and Journalism (No. 6) [Broadcast]. In *Knowing Machines*. https://engelberg-center-live.simplecast.com/episodes/synthetic-media-ai-and-journalism

Bäck, A., Diakopoulos, N., Granroth-Wilding, M., Haapanen, L., Leppänen, L. J., Melin, M., Moring, T. A., Munezero, M. D., Siren-Heikel, S. J., Södergård, C., & Toivonen, H. (2019). *News Automation: The rewards, risks and realities of "machine journalism."* World Association of Newspapers and News Publishers, WAN-IFRA.

Bagozzi, R. (2007). The Legacy of the Technology Acceptance Model and a Proposal for a Paradigm Shift. *Journal of the Association for Information Systems*, *8*(4), 244–254. https://doi.org/10.17705/1jais.00122

Barnes, C., & Barraclough, T. (2020). Deepfakes and synthetic media. In R. Steff, J. Burton, & S. R. Soare (Eds.), *Emerging Technologies and International Security: Machines, the State, and War*. Routledge. https://doi.org/10.4324/9780367808846

Becker, K. B., Simon, F. M., & Crum, C. (2025). Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News Organisations. *Digital Journalism*, *0*(0), 1–21. https://doi.org/10.1080/21670811.2024.2431519

Beckett, C., & Yaseen, M. (2023). *Generating Change: The Journalism AI report*. Polis, London School of Economics and Political Science. https://www.journalismai.info/s/Generating-Change-_-The-Journalism-AI-report-_-English.pdf

Borchardt, A., Simon, F., Zachrison, O., Bremme, K., Kurczabinska, J., Mulhall, E., & Johanny, Y. (2024). *Trusted journalism in the age of generative AI*. European Broadcasting Union. https://ora.ox.ac.uk/objects/uuid:8c874e2e-34de-4813-ba23-84e6300af110

Borden, S. L., & Tew, C. (2007). The Role of Journalist and the Performance of Journalism: Ethical Lessons From "Fake" News (Seriously). *Journal of Mass Media Ethics*, *22*(4), 300–314. https://doi.org/10.1080/08900520701583586

Broussard, M., Diakopoulos, N., Guzman, A. L., Abebe, R., Dupagne, M., & Chuan, C.-H. (2019). Artificial Intelligence and Journalism. *Journalism & Mass Communication Quarterly*, *96*(3), 673–695. https://doi.org/10.1177/1077699019859901

Cazzamatta, R., & Sarısakaloğlu, A. (2025). Mapping Global Emerging Scholarly Research and Practices of AI-supported Fact-Checking Tools in Journalism. *Journalism Practice*, 1–23. https://doi.org/10.1080/17512786.2025.2463470

Center for News, Technology & Innovation. (2025). *What It Means to Do Journalism in the Age of AI: Journalist Views on Safety, Technology and Government*. https://innovating.news/2024-journalist-survey/

Chesney, R., & Citron, D. K. (2018). Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3213954

Cools, H., & Diakopoulos, N. (2024). Uses of Generative AI in the Newsroom: Mapping Journalists' Perceptions of Perils and Possibilities. *Journalism Practice*, *0*(0), 1–19. https://doi.org/10.1080/17512786.2024.2394558

Cools, H., & Koliska, M. (2024). News Automation and Algorithmic Transparency in the Newsroom: The Case of the Washington Post. *Journalism Studies*, *25*(6), 662–680. https://doi.org/10.1080/1461670X.2024.2326636

de-Lima-Santos, M.-F., Yeung, W. N., & Dodds, T. (2024). Guiding the way: A comprehensive examination of AI guidelines in global media. *AI & Society*. https://doi.org/10.1007/s00146-024-01973-5

Diakopoulos, N., Cools, H., Li, C., Helberger, N., Kung, E., Rinehart, A., & Gibbs, L. (2024). *Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem*. https://doi.org/10.13140/RG.2.2.31540.05765

Dodds, T., Zamith, R., & Lewis, S. C. (2025). The AI turn in journalism: Disruption, adaptation, and democratic futures. *Journalism*, *0*(0). https://doi.org/10.1177/14648849251343518

Dunstan, J., & Ortolan, M. (2024, January 31). An AI-generated image of a Victorian MP raises wider questions on digital ethics. *ABC News*. https://www.abc.net.au/news/2024-02-01/georgie-purcell-ai-image-nine-news-apology-digital-ethics/103408440

Eder, M., & Sjøvaag, H. (2025). Falling behind the adoption curve: Local journalism's struggle for innovation in the AI transformation. *Journal of Media Business Studies*, *0*(0), 1–19. https://doi.org/10.1080/16522354.2025.2473301

Elbeyi, E., Bruhn Jensen, K., Aronczyk, M., Asuka, J., Ceylan, G., Cook, J., Erdelyi, G., Ford, H., Milani, C., Mustafaraj, E., Ogenga, F., Yadin, S., Howard, P. N., Valenzuela, S., Brulle, R., Jacquet, J., Lewandowsky, S., & Roberts, T. (2025). *Information Integrity about Climate Science: A Systematic Review*. International Panel on the Information Environment (IPIE). https://doi.org/10.61452/BTZP3426

Etikan, I. (2016). Comparison of Convenience Sampling and Purposive Sampling. *American Journal of Theoretical and Applied Statistics*, *5*(1), 1. https://doi.org/10.11648/j.ajtas.20160501.11

European Broadcasting Union. (2025, May 5). *Media outlets worldwide join call for AI companies to help protect news integrity*. https://www.ebu.ch/news/2025/05/media-outlets-worldwide-join-call-for-ai-companies-to-help-protect-news-integrity

Farhi, P. (2023, January 17). CNET used AI to write articles. It was a journalistic disaster. *The Washington Post*. https://www.washingtonpost.com/media/2023/01/17/cnet-ai-articles-journalism-corrections/

Feher, K. (2024). Exploring AI media. Definitions, conceptual model, research agenda. *Journal of Media Business Studies*, *21*(4), 340–363. https://doi.org/10.1080/16522354.2024.2340419

Ferrucci, P., & Perreault, G. (2021). The Liability of Newness: Journalism, Innovation and the Issue of Core Competencies. *Journalism Studies*, *22*(11), 1436–1449. https://doi.org/10.1080/1461670X.2021.1916777

Furtáková, L., & Janáčková, Ľ. (2023). AI in Radio: The Game Changer You Did Not Hear Coming. In M. Prostináková Hossová, M. Martovič, & M. Solík (Eds.), *Marketing Identity: AI – The Future Of Today. Proceedings from the International Scientific Conference*. University of Ss. Cyril and Methodius. https://mmidentity.fmk.sk/wp-content/uploads/2024/10/MM_2023_eng.pdf

Golding, P., & Murdock, G. (2022). The Political Economy of Contemporary Journalism and the Crisis of Public Knowledge. In S. Allan, *The Routledge Companion to News and Journalism* (2nd ed., pp. 36–45). Routledge. https://doi.org/10.4324/9781003174790-5

Gutierrez Lopez, M., Porlezza, C., Cooper, G., Makri, S., MacFarlane, A., & Missaoui, S. (2023). A Question of Design: Strategies for Embedding AI-Driven Tools into Journalistic Work Routines. *Digital Journalism*, *11*(3), 484–503. https://doi.org/10.1080/21670811.2022.2043759

Gutiérrez-Caneda, B., Lindén, C.-G., & Vázquez-Herrero, J. (2024). Ethics and journalistic challenges in the age of artificial intelligence: Talking with professionals and experts. *Frontiers in Communication*, *9*. https://doi.org/10.3389/fcomm.2024.1465178

Hall, C. J. (2025). Platform journalism on YouTube: A democratic functions approach to analysing journalism on digital platforms. *Australian Journalism Review*, *47*(1), 97–115. https://doi.org/10.1386/ajr_00178_7

Harris, K. R. (2024). Synthetic Media Detection, the Wheel, and the Burden of Proof. *Philosophy & Technology*, *37*(4), 131. https://doi.org/10.1007/s13347-024-00821-0

He, X., & Fang, L. (2024). Regulatory Challenges in Synthetic Media Governance: Policy Frameworks for AI-Generated Content Across Image, Video, and Social Platforms. *Journal of Robotic Process Automation, AI Integration, and Workflow Optimization*, *9*(12), Article 12.

Helberger, N., van Drunen, M., Moeller, J., Vrijenhoek, S., & Eskens, S. (2022). Towards a Normative Perspective on Journalistic AI: Embracing the Messy Reality of Normative Ideals. *Digital Journalism*, *10*(10), 1605–1626. https://doi.org/10.1080/21670811.2022.2152195

Hermida, A. (2015). Nothing but the truth: Redrafting the journalistic boundary of verification. In M. Carlson & S. C. Lewis (Eds.), *Boundaries of Journalism*. Routledge.

Jones, B., Jones, R., & Luger, E. (2022). AI 'Everywhere and Nowhere': Addressing the AI Intelligibility Problem in Public Service Journalism. *Digital Journalism*, *10*(10), 1731–1755. https://doi.org/10.1080/21670811.2022.2145328

Jones, B., Jones, R., & Luger, E. (2023). *Generative AI & Journalism: A Rapid Risk-Based Review*. University of Edinburgh. https://www.research.ed.ac.uk/en/publications/generative-ai-amp-journalism-a-rapid-risk-based-review

Kieran, M. (1998). Objectivity, impartiality and good journalism. In M. Kieran (Ed.), *Media Ethics* (1st ed., pp. 23–36). Routledge.

Lin, B., & Lewis, S. C. (2022). The One Thing Journalistic AI Just Might Do for Democracy. *Digital Journalism*, *10*(10), 1627–1649. https://doi.org/10.1080/21670811.2022.2084131

Lindén, T. C.-G., & Dierickx, L. (2019). Robot Journalism: The Damage Done by a Metaphor. *Unmediated*, *2*, 152–155.

Mahadevan, A. (2025, March 20). An Italian newspaper launched a generative AI experiment. It's not going well. *Poynter*. https://www.poynter.org/tech-tools/2025/il-foglio-newspaper-generated-artificial-intelligence/

Martin, A., & Newell, B. (2024). Synthetic Data, Synthetic Media, and Surveillance. *Surveillance & Society*, *22*(4), 448–452. https://doi.org/10.24908/ss.v22i4.18334

Matich, P., Thomson, T. J., & Thomas, R. J. (2025). Old Threats, New Name? Generative AI and Visual Journalism. *Journalism Practice*, *0*(0), 1–20. https://doi.org/10.1080/17512786.2025.2451677

Medianet. (2025). *2025 Australian Media Landscape Report*. https://engage.medianet.com.au/2025-media-landscape-report

Meir, N. (2015, June 15). Automated earnings stories multiply. *The Associated Press*. https://www.ap.org/the-definitive-source/announcements/automated-earnings-stories-multiply/

Min, S. J., & Fink, K. (2021). Keeping Up with the Technologies: Distressed Journalistic Labor in the Pursuit of "Shiny" Technologies. *Journalism Studies*, *22*(14), 1987–2004. https://doi.org/10.1080/1461670X.2021.1979425

Møller, L. A., Cools, H., & Skovsgaard, M. (2025). One Size Fits Some: How Journalistic Roles Shape the Adoption of Generative AI. *Journalism Practice*, 1–22. https://doi.org/10.1080/17512786.2025.2484622

Montaña-Niño, S. (2024). Automated Journalistic Assemblages. A Conceptual Approach to the Normative and Ethical Debates on AI Implementation in Newsrooms. *Problemi Dell'informazione*, *1*. https://doi.org/10.1445/113227

Moran, C. (2023, April 6). ChatGPT is making up fake Guardian articles. Here's how we're responding. *The Guardian*. https://www.theguardian.com/commentisfree/2023/apr/06/ai-chatgpt-guardian-technology-risks-fake-article

Moran, R. E., & Shaikh, S. J. (2022). Robots in the News and Newsrooms: Unpacking Meta-Journalistic Discourse on the Use of Artificial Intelligence in Journalism. *Digital Journalism*, *10*(10), 1756–1774. https://doi.org/10.1080/21670811.2022.2085129

Oliver, L. (2024, November 1). This chatbot helps tell the story of how women are affected by drug trafficking in Paraguay. *Reuters Institute News*. https://reutersinstitute.politics.ox.ac.uk/news/chatbot-helps-tell-story-how-women-are-affected-drug-trafficking-paraguay

Paris Charter on AI and Journalism. (2023, November 10). https://rsf.org/sites/default/files/medias/file/2023/11/Paris%20Charter%20on%20AI%20and%20Journalism.pdf

Partnership on AI. (2023, February 27). *PAI's Responsible Practices for Synthetic Media*. https://partnershiponai.org/

Petković, B., & Mihajlović Trbovc, J. (Eds.). (2014). *Media integrity matters: Reclaiming public service values in media and journalism* (1st ed). Peace Institute, Institute for Contemporary Social and Political Studies.

Radcliffe, D. (2025). *Journalism in the AI Era* (TRF Insights). Thomson Reuters Foundation. https://www.trust.org/wp-content/uploads/2025/01/TRF-Insights-Journalism-in-the-AI-Era.pdf

Riordan, K. (2014). *Accuracy, Independence, and Impartiality: How legacy media and digital natives approach standards in the digital age*. Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/our-research/accuracy-independence-and-impartiality-how-legacy-media-and-digital-natives-approach

Roper, D., Henriksson, T., Hälbich, K., & Martin, O. (2023). *Gauging Generative AI's impact on newsrooms*. World Association of News Publishers (WAN-IFRA). https://wan-ifra.org/insight/gauging-generative-ais-impact-in-newsrooms/

Salas, A., Rivero-Calle, I., & Martinón-Torres, F. (2023). Chatting with ChatGPT to learn about safety of COVID-19 vaccines – A perspective. *Human Vaccines & Immunotherapeutics*, *19*(2). https://doi.org/10.1080/21645515.2023.2235200

Samosir, H. (2023, July 13). AI news readers are becoming more common across Asia. Will Australia follow suit? *ABC News*. https://www.abc.net.au/news/2023-07-13/artificial-intelligence-news-readers-becoming-common-in-asia/102591790

Schell, K. (2024). *AI transparency in journalism: Labels for a hybrid era*. Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2025-01/RISJ%20Fellows%20Paper_Katja%20Schell_MT24_Final.pdf

Simon, F. M. (2022). Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy. *Digital Journalism*, *10*(10), 1832–1854. https://doi.org/10.1080/21670811.2022.2063150

Simon, F. M. (2024). *Artificial Intelligence in the News How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena*. Tow Center for Digital Journalism. https://journalism.columbia.edu/news/tow-report-artificial-intelligence-news-and-how-ai-reshapes-journalism-and-public-arena

Simon, F. M., Altay, S., & Mercier, H. (2023). Misinformation reloaded? Fears about the impact of generative AI on misinformation are overblown. *Harvard Kennedy School Misinformation Review*. https://doi.org/10.37016/mr-2020-127

Simon, F. M., & Isaza-Ibarra, L. F. (2023). *AI in the News: Reshaping the Information Ecosystem?* Oxford Internet Institute. https://www.oii.ox.ac.uk/wp-content/uploads/2023/08/Minderoo_Report_Simon_Ibarra.pdf

Sjøvaag, H. (2024). The business of news in the AI economy. *AI Magazine*, *45*(2), 246–255. https://doi.org/10.1002/aaai.12172

Society of Professional Journalists. (2014). *SPJ Code of Ethics*. https://www.spj.org/spj-code-of-ethics/

Squicciarini, M., Valdez Genao, J., & Sarmiento, C. (2024). *Synthetic Content and AI Policy: A Primer*. UNESCO. https://policycommons.net/artifacts/17958669/synthetic-content-and-its-implications-for-ai-policy/18857919/

Ternovski, J., Kalla, J., & Aronow, P. M. (2022). The Negative Consequences of Informing Voters about Deepfakes: Evidence from Two Survey Experiments. *Journal of Online Trust and Safety*, *1*(2). https://doi.org/10.54501/jots.v1i2.28

Thomson, T. J., Thomas, R. J., & Matich, P. (2024). Generative Visual AI in News Organizations: Challenges, Opportunities, Perceptions, and Policies. *Digital Journalism*. https://doi.org/10.1080/21670811.2024.2331769

Thomson, T. J., Thomas, R., Riedlinger, M., & Matich, P. (2025). *Generative AI & journalism*. RMIT University. https://doi.org/10.6084/m9.figshare.28068008

Toff, B., & Simon, F. M. (2024). "Or They Could Just Not Use It?": The Dilemma of AI Disclosure for Audience Trust in News. *The International Journal of Press/Politics*, *0*. https://doi.org/10.1177/19401612241308697

Tran, M. (2006, August 18). Robots write the news. *The Guardian*. https://www.theguardian.com/news/blog/2006/aug/18/robotswriteth

WashPostPR. (2024, November 7). The Washington Post Launches "Ask The Post AI," a New Search Experience. *The Washington Post*. https://www.washingtonpost.com/pr/2024/11/07/washington-post-launches-ask-post-ai-new-search-experience/

Whittaker, L., Kietzmann, T. C., Kietzmann, J., & Dabirian, A. (2020). "All Around Me Are Synthetic Faces": The Mad World of AI-Generated Media. *IT Professional*, *22*(5), 90–99. https://doi.org/10.1109/MITP.2020.2985492

Wilding, D., Fray, P., Molitorisz, S., & McKewon, E. (2018). *The Impact of Digital Platforms on News and Journalistic Content*. UTS Centre for Media Transition. http://hdl.handle.net/10453/159124

Wintterlin, F., Engelke, K. M., & Hase, V. (2020). Can transparency preserve journalism's trustworthiness? Recipients' views on transparency about source origin and verification regarding user-generated content in the news. *Studies in Communication and Media*, *9*(2), 218–240. https://doi.org/10.5771/2192-4007-2020-2-218

Zier, J., & Diakopoulos, N. (2024, October 26). Labeling AI-Generated News Content: Matching Journalist Intentions with Audience Expectations. *Proceedings of the Computation and Journalism Symposium 2024*. https://cplusj2024.github.io/

Appendix A

Table of participants