Institutional Ethics of AI in Legal Scholarship

Jeannie Marie Paterson

08.10.2026 doi: 10.59704/b65051f31807a528

Scholars interested in the regulation of AI will be familiar with the Collingridge dilemma, which suggests that the challenge of regulating technology arises from a tension between information and control. In the early stages of a technology's development, it can be difficult to predict its social impacts; by the time the risks are apparent, the technology may be so embedded in social structures that it is difficult to control.[1]

Something of the same dilemma may seem to face law schools in developing an institutional ethics for AI in legal scholarship. The dilemma of how and when to respond to the rise of generative AI in legal practice, the university and scholarly research would seem to match Collingridge's framing, as the uses of the tools are evolving so quickly that practice has not had time to settle, yet the consequences of not responding may erode enshrined research values.

The challenge should not be insurmountable. Collingridge did not suggest new technology could not be regulated. Rather, he developed a theory of decision-making under conditions of uncertainty, which allowed decisions to be "reversible, corrigible, and flexible"[2]. Universities and law schools thrive on debate and the reassessment of social practices, and this aptitude to encourage and manage debate provides a response to the challenge of integrity posed by AI for legal scholarship. This essay explores the challenge and the possible responses, including advocating the importance of consistency and transparency.

AI in Legal Practice and Legal Education

It is clear that legal practice is changing under the influence of emerging technology, particularly generative AI[3], although the longer-term trends are still to emerge. Law schools are responding differently to these changes. Some universities have imposed rules about secure assessment, such as invigilated exams representing a specified percentage of the grade[4]. Some law schools have partnerships with law tech providers to allow students to experiment with the products. Others have introduced a suite of tech-focused subjects and even mandated a unit in legal technology[5]. Largely, however, law schools remain in a period of uncertainty and experimentation.

There is a clear parallel with the Collingridge dilemma as law schools seek to understand the potential of the technology, the economics of adopting it, and its impact on legal practice. Law schools wish to prepare their students for legal practice, but largely they do not consider themselves trade schools responsible for teaching precise legal skills, such as the uses of particular technologies. Law schools commonly aspire to higher-order learning about the law: analysis, critique, judgment. Nonetheless, the advance of AI may mean that different manifestations of these skills are required or that students themselves wish to emphasise the technology in preparation for future work. Legg argues that while lawyers will need a new suite of tech-related skills, ethical judgment will remain a fundamental lawyering skill and should be

[1] David Collingridge, The Social Control of Technology (Frances Pinter Ltd, 1980).

[2] *Ibid.*

[3] Julian Webb and Jeannie Paterson, The Evolution of Legal Knowledge Work in an Age of Brilliant(?) Technologies: From Robo-Lawyer to Digital Law Clerk, ANU Journal of Law and Technology 6(1) (2025), p. 51-84.

[4] See, for instance, Gen AI – Academic Integrity and Assessment Reform, Australian Tertiary Education Quality and Standards Agency, https://www.teqsa.gov.au/guides-resources/higher-education-good-practice-hub/gen-ai-knowledge-hub/gen-ai-academic-integrity-and-assessment-reform, last accessed 10.08.2026.

explicitly taught as an aspect of any AI-enabled education[6].

At the same time, AI poses what Clay Shirky describes as "a menagerie of threats", increasingly backed up by research, that include "cognitive offloading, cognitive debt, cognitive surrender, automation bias, illusion of competence, moral disengagement"[7]. Law schools, therefore, need not only to retain assessment integrity but also to navigate new relationships with technology in their student body and those students' future workplaces.

The Dilemma of AI in Legal Research

AI also raises ethical and practical challenges for traditional legal research. Some legal scholars have become interested in research about AI, and this is important and timely. AI may be used as a tool for data analysis, which is consistent with its foundations and largely uncontroversial. Greater controversy arises in the use of AI by legal scholars to produce work. Generative AI can be used for many of the basic practices of scholarly work: brainstorming, preliminary research, summarising, proofreading. It can also be used (although the quality may vary) for higher-order tasks: generating ideas, drafting, writing and critique.

The use of AI, at least for more substantive purposes, may raise concern in the wider scholarly community, though some would object to any use at all on the ground that AI is then framing the fundamental questions to be addressed. There may be a fear of others getting ahead by using the tools in fraudulent ways, what Frazier and Rozenshtein describe as "scholarly deepfakes"[8]. Another concern goes to the essence of legal scholarly endeavour.

Legal research is far more than the assembling of information, and the capacity of AI for critical analysis is doubtful. It mimics rather than produces analysis, and its outputs are the result of statistical processes. As Julian Webb observes, generative AI "struggles with the nuances of real-world legal problem-solving, particularly complex relational problems where practical wisdom and contextual understanding play a significant role"[9]. The experience of being human may be relevant to output in the law. At least in the context of courts, Pasquale argues that "embodied judgment" may be the best approach for delivering "a relatable, respected, and legitimate justice system rather than an approach focused on the words and reasons of AI"[10]. It might be thought that, to the extent the law remains a profession and vocation governed by human embodied "judgment", the work of the legal scholar should be similarly grounded. What this commitment requires in practice is not straightforward and underlines the need for an institutional ethical framework.

Principles-Based Regulation and Ethical Problem Solving

In institutional contexts, ethical frameworks operate as a form of soft law regulation, providing guidance on the parameters of what is acceptable as well as an aspiration to good research practice. Additionally, ethical frameworks may provide evidence of compliance with legal requirements, such as in the law of tort, and transgressions may breach existing laws, such as privacy and data protection laws.

Precise rules are important in some contexts. For example, speed limits in suburban areas are set as a rule with an exact maximum speed. In other contexts, however, guidance on ethical or legal conduct is expressed as a principle. Principles need to be adapted to the context and require a judgment about what is right, but they can guide that decision and may be supplemented with factors to consider. For example, in determining whether a professional exercised due care

[5] See, for instance, Shape Your Future With NUS Law, National University of Singapore, 15.08.2024, https://law.nus.edu.sg/media/shape-your-future-with-nus-law/, last accessed 10.08.2026.

[6] Michael Legg, Better Than a Bot – Instilling Ethical Judgement Into the Lawyers of the Future in the Age of AI, Griffith Law Review 33(3) (2024), p. 273-293.

[7] Clay Shirky, AI Is a "Wicked Problem." That's the University's Specialty, The Chronicle of Higher Education, 04.08.2026, https://www.chronicle.com/article/ai-is-a-wicked-problem-thats-the-universitys-specialty?bc_nonce=4jrrouz95956vrcwltvmts&cid=reg_wall_signup, last accessed 10.08.2026.

[8] Kevin Frazier and Alan Rozenshtein, Large Language Scholarship, Minnesota Legal Studies Research Paper No 25-26, SSRN, 04.04.2025, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5200768, last accessed 10.08.2026.

and skill in providing a service, courts consider a range of contextual factors including the skill of the professional, established practice and the circumstances of the client. The ethical standards for the legal professional are expressed in a similar manner: as open-textured principles that require legal and ethical judgment to apply. Webb explains that legal ethics cannot be a substitute for the lawyer's personal moral compass[11]. This requirement for judgment about what is the right way to act is most closely associated with virtue theory or "practical wisdom"[12].

Principles-based governance thus provides a degree of flexibility and adaptivity. Critics may say this adaptivity comes at the cost of certainty. But certainty in the fast-moving realm of AI is a mirage in any event. Certainly, the approach will be contested, which means being subject to challenge, dissent and critique. This is not a fatal flaw. Debate and the airing of diverse views are the hallmarks of the university as a place of scholarly endeavour. A pluralist approach is not destabilising if the end — integrity — is kept in mind. Ethical decision-making is always difficult but can be guided by an understanding of the end goal — be it efficiency or virtue.

Guiding Principles

In formulating ethical frameworks for legal scholarship, indicative guiding principles are already available and suitable to support the objective of research integrity while still allowing scholars to experiment with the possibilities of generative AI. The risks of harm from AI are increasingly well recognised, including bias, surveillance, error and cognitive offloading. Equally, there is no shortage of statements of ethical approaches to AI with aspirations to beneficence, fairness and accountability[13]. Courts and legal regulators have similarly adopted guidance, which might guide without binding law schools, and importantly include a commitment to values of integrity, honesty and diligence[14]. These standards resonate with the research integrity frameworks themselves, which are usually premised on higher-order values of honesty, rigour, transparency, fairness, respect and accountability[15].

These frameworks can provide the baseline principles in considering ethical governance of AI in legal scholarship. In the university environment, institutional consistency might be added to these lists as a core value. Here the commitment might be that if law schools are giving guidance to students, then we too should follow it. Of course, scholars are not being certified as ready for practice. But if we ask students to forgo technology in the classroom or for study, then it is reasonable to reflect on our own practices in teaching and research.

I would also affirm the foundational role of transparency, although not all would agree.

The Value of Transparency

As Coghlan, Miller and Paterson argue, "[a]lthough transparency is not necessarily or always an ethical good, it is associated with more basic ethical ideas such as justice and respect for autonomy sufficiently frequently that it is often treated as a key ethical principle in AI Ethics."[16] In the context of legal scholarship, transparency on AI use means including a frank disclosure of the ways in which AI was used to contribute to the manuscript. The need for transparency arises from a commitment to honesty in the process of research and scholarly writing in a time when AI can contribute substantively to that process, and practices vary between institutions and scholars.

Frazier and Rozenshtein are critical of disclosure as a response to concerns about AI use in legal scholarship, arguing it may "fail to advance scholarly values" and distract from genuine "authorial responsibility"[17]. However, the role of transparency is not to displace responsibility. Transparency, properly understood, is a

[9] Julian Webb, Generative Artificial Intelligence in Legal Practice: Epistemic Expertise, Unethical AI, and the Role of Virtue, Minds and Machines 36 (2026).

[10] Frank Pasquale, The Non-Delegable Duty to Think: Judicial Legitimacy and the Limits of Generative AI, Cornell Legal Studies Research Paper No 26-03, SSRN, 26.03.2026, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6473658, last accessed 10.08.2026.

[11] Julian Webb, Generative Artificial Intelligence in Legal Practice: Epistemic Expertise, Unethical AI, and the Role of Virtue, Minds and Machines 36 (2026).

baseline element of accountability, performing several roles[18].

At this point it might be rhetorically asked whether an author should disclose spell check or editing software[19]. Scholars are familiar with spell check, and thus it may go without saying. At present, the scholarly community is less familiar with generative AI, and practices are varied. Moreover, use of generative AI may be going further than merely checking; it is possibly contributing, and contributions are usually disclosed, such as through research assistants or helpful colleagues. Such acknowledgments are not merely kind; they recognise the distribution of labour in producing the work, without deflecting the author's overall responsibility for the work.

Although it may have a performative element, transparency may prompt the author to reflect on their own role in producing the output. In pondering why lawyers continue to include AI-hallucinated references in submissions to courts, Legg, McNamara and Alimardani point to the behavioural biases that may allow hallucinated content to slip through despite now widespread knowledge of the risk[20]. These tendencies include automation bias and verification drift, which essentially refer to a tendency to believe the apparently accurate outputs produced by AI and a drift away from good intentions to exercise active authorial control over the work. Similar tendencies may affect legal scholars, and transparency may provide a practical reminder of their own responsibilities.

Additionally, transparency benefits the reader, who may learn from the disclosure and reflect on their own use. In this way, transparency may contribute to a broader community of practice on ethical use.

Conclusion

Responding to the influx of AI into scholarly writing, in the legal domain and elsewhere, poses real challenges. It will be neither possible nor appropriate to develop a set of determinative rules covering all possible uses. Ethical frameworks do not operate in this way. Instead, they should be principles-based, allowing them to accommodate both the evolution of the technology and developments in scholarly practice taking place in the light of shared commitments to research integrity.

I used AI (Claude) to proofread this chapter.

[12] Julian Webb, Generative Artificial Intelligence in Legal Practice: Epistemic Expertise, Unethical AI, and the Role of Virtue, Minds and Machines 36 (2026).

[13] See, for instance, OECD AI Principles Overview, OECD, https://oecd.ai/en/ai-principles, last accessed 10.08.2026.

[14] See, for instance, Statement on the Use of Artificial Intelligence in Australian Legal Practice, Victorian Legal Services Board, 06.12.2024, https://lsbc.vic.gov.au/news-updates/news/statement-use-artificial-intelligence-australian-legal-practice, last accessed 10.08.2026.

[15] See, for instance, Australian Code for the Responsible Conduct of Research 2018, Australian Research Council, https://www.arc.gov.au/australian-code-responsible-conduct-research-2018, last accessed 10.08.2026.

[16] Simon Coghlan et al., Good Proctor or "Big Brother"? Ethics of Online Exam Supervision Technologies, Philosophy and Technology 34 (2021), p. 1581-1606.

ABOUT THE AUTHOR

Jeannie Marie Paterson

Jeannie Marie Paterson is a Professor of Law at the University of Melbourne and the co-founding director of the Centre for Artificial Intelligence and Digital Ethics (CAIDE).

SUGGESTED CITATION

Paterson, Jeannie Marie: Institutional Ethics of AI in Legal Scholarship, Verfassungsblog, 08.10.2026, doi: 10.59704/b65051f31807a528.

Licensed under CC BY-SA 4.0.

Langzeitarchiviert bei IntRechtDok unter: doi.org/10.17176/20261009-143454-0.

[17] Kevin Frazier and Alan Rozenshtein, Large Language Scholarship, Minnesota Legal Studies Research Paper No 25-26, SSRN, 04.04.2025, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5200768, last accessed 10.08.2026.

[18] Jeannie Paterson, Understanding AI Transparency and Literacy: Lessons From Consumer Credit Regulation, UNSW Law Society Court of Conscience 19 (2025), p. 107; Jeannie Paterson, Misleading or Deceptive AI: Why Transparency Provides the Baseline for Responsible AI in Consumer Transactions, UNSW Law Society Court of Conscience 18 (2024), p. 113.

[19] Kevin Frazier and Alan Rozenshtein, Large Language Scholarship, Minnesota Legal Studies Research Paper No 25-26, SSRN, 04.04.2025, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5200768, last accessed 10.08.2026.

Michael Legg, Vicki McNamara, and Armin Alimardani, The Promise and the Peril of the Use of Generative Artificial Intelligence in Litigation, UNSW Law Journal 48(4) (2025), p. 1196-1235.