The Future Lawyer: From AI-Enabled Practice to Contained Agentic Work

Author: Donald Feaver
Paper Type: Working Paper
Citation: Donald Feaver, *The Future Lawyer: From AI-Enabled Practice to Contained Agentic Work* (Working Paper, Leo Cussen Centre for Law, October 2026).
Series / Paper Number: Working Paper No. 2

Publication Details

Date: October 2026
Status: Working Paper – Not Peer Reviewed

Contact

Professor Donald feaver
Leo Cussen Centre for Law
dfeaver@leocussen.edu.au

Copyright and Disclaimer

© 2026 Leo Cussen Centre for Law. All rights reserved unless otherwise stated. Reproduction and distribution require permission except as permitted by law. The views expressed are those of the authors and do not necessarily represent Leo Cussen Centre for Law.

THE FUTURE LAWYER: FROM AI-ENABLED PRACTICE TO CONTAINED AGENTIC WORK

Working Paper

PROFESSOR DONALD FEAVER
LEO CUSSEN CENTRE FOR LAW

October 2026

ABSTRACT

The future lawyer will work alongside AI agents, and the profession needs a reliable way to contain their work. This working paper argues that the foundations for doing so already exist provided the structures of AI-assisted practice are deliberately and carefully designed.

It begins with the AI-enabled lawyer: one who can see legal work as a workflow, divide it into steps, and distinguish work that describes supplied material, which AI may do, from work that determines legal consequence, which the lawyer must do. It then turns to the AI-enabled firm, where that division becomes structural. Work is allocated through the workflow to the person holding authority: junior lawyers verify AI-assisted work against its sources; supervising lawyers allocate the work, verify its adequacy and review the determinations made; partners provide assurance; and a controlled environment records what is done.

AI agents complete sequences of work and make choices along the way. The same structure contains them. Specifications authorised in advance set what an agent may do; points of legal determination mark where it must stop; the environment records what it did; and a lawyer with standing adopts the result. Early results from an ongoing simulation research program show both the reach and the limit: an agent performed almost all the work of a legal matter but could not hold the lawyer's position. The paper also sets out where containment stops working. The future lawyer's task is to build that structure and hold it.

Keywords: legal profession; artificial intelligence; AI agents; supervision; professional responsibility; legal education; practical legal training

Drafting note (v4, October 2026). Figures from the simulation are marked ‡; they come from the first runs of an ongoing research program and will be updated as results settle.

I INTRODUCTION

Courts, regulators and professional bodies have settled on a single answer to generative artificial intelligence ('AI'). Lawyers may use it, but they must understand its limits, check what it produces and remain responsible for the work.[1] Courts have enforced that position, in some cases, by referring practitioners to their regulator.[2] At this stage of AI's development, that position is the current one. On its own, however, 'check and be responsible' is extremely vague. It does not tell a lawyer where AI may be used in a matter, which parts of the work it may do, which decisions must stay with a person, or how its output should be checked.[3] It does not tell a firm how to organise many lawyers doing this at once. And it says nothing about what 'responsible' means once an AI system can complete whole sequences of work on its own, making choices as it goes.

AI is arriving in legal practice in two overlapping phases. In the first, individual lawyers use AI to help with particular pieces of work: summarising documents, organising a chronology or producing a first draft. In the second, AI agents take on whole sequences of work. An AI agent is a system that works through a series of steps towards a goal, making choices along the way, rather than answering one question at a time.[4] How far agents are already in use depends on how they are defined. A recent foresight report for the solicitors' profession in England and Wales found no genuinely agentic AI in legal practice. An industry survey, by contrast, reports a minority of legal organisations already using systems they describe as agentic and a large majority expecting such systems to

---

[1] Law Society of New South Wales, Legal Practice Board of Western Australia and Victorian Legal Services Board and Commissioner, *Statement on the Use of Artificial Intelligence in Australian Legal Practice* (6 December 2024) ('*Joint Statement*'); Federal Court of Australia, *Generative Artificial Intelligence Practice Note (GPN-AI)*, 16 April 2026 ('*GPN-AI*'); Supreme Court of New South Wales, *Practice Note SC Gen 23: Use of Generative Artificial Intelligence*, 21 November 2024 ('*SC Gen 23*'); Supreme Court of Victoria, *Practice Note SC Gen 25: The Use of Artificial Intelligence by Court Users*, 14 May 2026 ('*SC Gen 25*').

[2] See, eg, *Mata v Avianca, Inc*, 678 F Supp 3d 443 (SD NY, 2023) ('*Mata*'); *R (Ayinde) v Haringey London Borough Council* [2025] EWHC 1383 (Admin) ('*Ayinde*') [case name: confirm]; *Valu v Minister for Immigration and Multicultural Affairs (No 2)* [2025] FedCFamC2G 95 ('*Valu*'); *JNE24 v Minister for Immigration and Citizenship* [2025] FedCFamC2G 1314 ('*JNE24*').

[3] Donald Feaver, 'The AI-Enabled Lawyer and Law Firm: A Best-Practice Discipline for Legal Work with Artificial Intelligence' (Unpublished Manuscript, 22 August 2026) Preface pt 1 ('AI-Enabled Lawyer').

[4] Donald Feaver, 'What Can Be Written Down: An AI Agent Completes a Legal Matter, With and Without Legal Research' (Working Paper, 2026) pt 1 ('What Can Be Written Down') [authorship: confirm]; see also Noam Kolt, 'Governing AI Agents' (2026) 101 *Notre Dame Law Review* 335.

be central to their work by 2030.[5] On either view, the second phase is arriving before the profession has settled how to manage the first.

This paper argues that the profession already has the means to contain agentic work, provided it builds the structures of AI-assisted practice deliberately. The argument moves in three steps. It begins with the AI-enabled lawyer and the capabilities that lawyer needs. It then shows how those capabilities, practised across a firm, produce a structure: a way of dividing legal work, allocating authority over it, controlling the technology and recording what was done. Finally, it explains AI agents and shows how the same structure can contain them. Agents can do almost all of the work. What they cannot do is hold the lawyer's position: the standing to decide, accountability for the result and the commitment to stand behind it. The future lawyer's task is to build the structure that keeps that position where it belongs, and to hold it.

The thread running through the argument is the division of labour. Professions have always divided expert work, both among themselves and within themselves.[6] Regulatory theory adds a distinction that matters here. A single institution can contain two divisions of labour at once: one that decides who performs the work, and another that decides who controls it and who answers for it.[7] AI is changing the first quickly. The question for the profession is whether the second can be held steady while the first moves.

Part II explains what this paper contributes and why its approach differs from much of the existing writing. Part III describes the AI-enabled lawyer. Part IV describes the AI-enabled firm and the structure it produces. Part V explains AI agents, sets out how the structure contains them, reports early results from an ongoing simulation research program and identifies where containment stops working. Part VI concludes.

---

[5] Law Society of England and Wales, *The Future of Agentic AI in Legal Practice* (Foresight Report, 2026) [month of publication: confirm]; Thomson Reuters Institute, *2026 AI in Professional Services Report* (Report, 2026).

[6] Andrew Abbott, *The System of Professions: An Essay on the Division of Expert Labor* (University of Chicago Press, 1988).

[7] Donald Feaver and Benedict Sheehy, 'The Shifting Balance of Power in the Regulatory State: Structure, Strategy, and the Division of Labour' (2014) 41 *Journal of Law and Society* 203 [issue number: confirm]; Donald Feaver and Benedict Sheehy, 'The Political Division of Regulatory Labour: A Legal Theory of Agency Selection' (2015) 35(1) *Oxford Journal of Legal Studies* 153.

II WHAT THIS PAPER CONTRIBUTES

Much has been written about AI and the legal profession. Most of it falls into one of two camps. One predicts transformation and disruption. The other lists risks and restates professional duties. This paper takes a different path, and it is worth explaining at the outset why it is distinctive.

A. From Slogan to Specification

The paper's first contribution is to turn the profession's governing principle into something that can be specified, supervised and checked. Existing guidance states the duty: the lawyer is responsible. It rarely says where in the work that responsibility has to be exercised. This paper identifies the points in a piece of legal work at which a person's judgement is required, what kind of judgement each point calls for, who must exercise it and what record would show that it was exercised. That is what supervisors, insurers, regulators and courts need when something goes wrong, and what firms need in order to prevent it going wrong.[8]

B. A Line That Will Not Date

Second, the paper draws the line between human and machine work in a way that does not depend on what the technology can currently do. Most writing on legal AI is tied to the capability of particular systems, and dates as quickly as they do. The line proposed here asks a different question. It asks not whether a system can perform a piece of work, but whether that work decides a legal consequence for which someone must answer.[9] The answer to that question does not change when a new model is released. Nor does the paper's second anchor, standing, which depends on holding professional office rather than on skill.[10]

C. Containment Rather Than Alarm

Third, the paper offers a structural answer to agentic AI. It argues that agents can be governed through institutions the profession already has, namely supervision, the allocation of authority and the keeping of records, provided those institutions are built

[8] Donald Feaver, 'Beyond Verification: Professional Duty and Workflow Governance in AI-Assisted Legal Practice' (Working Paper, 2026) pts 1, 8 ('Beyond Verification') [authorship: confirm].

[9] Donald Feaver and Craig Collins, 'The Workflow as Anchor: Skill Architecture and the Division of Labour in AI-Assisted Legal Work' (Working Paper, 2026) pt 5.3.

[10] Ibid pt 5.4.

deliberately. That answer is neither alarmist nor promotional. It gives the leaders of firms, courts and regulators something they can act on now. It is also candid about where it stops working (Part V(F)).

D. Several Kinds of Evidence Together

Fourth, the paper brings together kinds of evidence that rarely appear side by side. They include a conceptual framework developed for practice and for legal education;[11] the first results of an ongoing research program in which an AI agent completes a legal matter;[12] an Australian survey of the supervisors who receive new lawyers into practice;[13] recent Australian and overseas decisions on AI in court; and research from regulatory theory, organisational studies and human factors. The conceptual argument does not rest on any single source. The empirical work illustrates it and tests it.

E. Why Now

Finally, the paper is timely. Australian courts issued or revised practice notes on generative AI between 2024 and 2026.[14] The arrangements for practical legal training are under review in New South Wales and nationally.[15] And in 2025 the regulator of solicitors in England and Wales authorised the first law firm to deliver regulated services through AI.[16] Decisions about how lawyers are trained, supervised and regulated are being made now. This paper is offered as a contribution to them.

F. How to Read the Paper

Two points about method. First, this is a working paper. It draws on a connected body of work by the authors: a book manuscript on AI-enabled practice, papers on legal skills and on the division of labour in AI-assisted work, a paper on professional duty and

[11] Feaver, 'AI-Enabled Lawyer' (n 3); Feaver and Collins (n 9); Donald Feaver, Craig Collins and Shirley Southgate, 'What Are Legal Skills? Legal Capability, Professional Formation and the Public Assurance of Competence' (Working Paper, Leo Cussen Centre for Law, 2026).

[12] Feaver, 'What Can Be Written Down' (n 4).

[13] Linda Baxter, Craig Collins and Donald Feaver, 'What Supervisors Want, What They Get: First-Year Lawyer Capability and AI Skills' (Working Paper, October 2026) ('Supervisors Survey').

[14] *GPN-AI* (n 1); *SC Gen 23* (n 1); *SC Gen 25* (n 1).

[15] Legal Profession Admission Board of New South Wales, *Interim Report on PLT Reform* (Report, 13 March 2026); Legal Services Council Admissions Committee and Law Admissions Consultative Committee, *Review of Practical Legal Training* (Consultation Paper, April 2026).

[16] Solicitors Regulation Authority, 'SRA Approves First AI-Driven Law Firm' (Media Release, 6 May 2025).

workflow governance, and the survey and simulation studies.[17] Those sources are cited so that readers can follow the argument to its foundations. Second, the argument is conceptual. The simulation and the survey are early-stage evidence. They show what the framework can do and where it is tested, but they do not prove it. Figures from the simulation come from the first runs of a research program that is continuing, and they will be updated as further results become available.

Table 1 sets out the few terms the paper relies on. Two are coined: *describing* and *determining* work, and *standing*. The others are used in their ordinary sense but given a precise meaning.

Table 1: Key terms used in this paper

Term	Meaning in this paper
Workflow	The sequence of steps through which a piece of legal work is framed, carried out, checked and authorised.
Task and step (operation)	A task is a recognisable unit of legal work, such as reviewing a contract or preparing advice. A step (or operation) is a smaller part of a task that can be separately assigned and checked, such as extracting dates from a set of emails.
Describing work	Work that reports what is in material supplied to it: extracting, organising, sorting into categories a lawyer has set, comparing, summarising or drafting from settled analysis. It may be assigned to AI and checked against the source.
Determining work	Work that decides a legal consequence: what the problem is, what a document means, how much a consideration weighs, what should be done and what position will be taken. It must remain with a lawyer.
Standing	The entitlement to make a decision and to answer for it. It belongs to a person who holds professional office, and is not a matter of skill.
AI agent	An AI system that works through a series of steps towards a goal, making choices along the way, rather than answering one question at a time.
Verification and assurance	Verification checks a particular fact, source or proposition against an authoritative basis. Assurance is the broader judgement that work is complete, fit for its purpose, properly supervised and authorised for use.
Containment	The arrangements that keep an agent's work within limits set in advance, stop it where a lawyer must decide, record what it did and leave the decision to rely on the work with a lawyer.

[17] Feaver, 'AI-Enabled Lawyer' (n 3); Feaver, Collins and Southgate (n 11); Feaver and Collins (n 9); Feaver, 'Beyond Verification' (n 8); Baxter, Collins and Feaver (n 13); Feaver, 'What Can Be Written Down' (n 4).

III THE AI-ENABLED LAWYER

A. What an AI-Enabled Lawyer Is

An AI-enabled lawyer is not defined by how close they are to the technology or how often they use it. The lawyer is defined by three abilities: to get real value from AI, to control how AI's contribution enters and leaves the work, and to use it to improve the handling of a matter without giving up legal judgement, authority or responsibility.[18] Two of these abilities pull against each other. *Mastery* is the ability to get useful help from the technology. *Control* is the discipline that keeps that help tied to the purpose of the work, the sources it relies on, the judgements a lawyer must make, supervision and authority.[19] Mastery without control speeds up unreliable work. Control without mastery forgoes real value. Mastery can be shown by deciding not to use AI at all, by narrowing what it is asked to do, or by rejecting output that would take too long to check. Efficiency is measured across the whole matter, not at the moment a draft appears.[20]

The evidence supports a measured view of what AI adds. In the first randomised controlled trial of AI assistance on legal tasks, law students given access to GPT-4 worked much faster, but the quality of their work improved only slightly and inconsistently.[21] A later trial using a reasoning model and a legal research tool built on retrieval from legal sources found significant gains in quality. The reasoning model deepened the analysis, and the research tool avoided introducing additional errors. The authors conclude that the value of AI depends on the nature of the legal work.[22] Outside law, a field experiment with management consultants described AI capability as a 'jagged technological frontier': strong on some tasks that appear difficult and weak on others that appear easy.[23] Mastery is therefore not a general facility with a tool. It is knowing, step by step, where AI helps and where it does not.

---

[18] Feaver, 'AI-Enabled Lawyer' (n 3) ch 1 pt 4.4.

[19] Ibid Preface pt 2, ch 1 pts 4.1–4.2.

[20] Ibid ch 1 pt 4.3.

[21] Jonathan H Choi, Amy B Monahan and Daniel Schwarcz, 'Lawyering in the Age of Artificial Intelligence' (2024) 109(1) *Minnesota Law Review* 147.

[22] Daniel Schwarcz et al, 'AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice' (2026) 3 *Journal of Law and Empirical Analysis* 220.

[23] Fabrizio Dell'Acqua et al, 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality' (2026) *Organization Science* (advance online publication) [volume and pages: confirm].

The definition places the lawyer's capability in integration: fitting AI's contribution into professional work. That requires three kinds of understanding that legal education and practice have mostly left unspoken. The lawyer must understand how legal work is structured, what legal skill consists of, and where the line falls between work that can be handed to AI and work that cannot.

B. Seeing the Work

The first is an understanding of workflow. There is one professional workflow, not a separate technological process running alongside it. AI enters, if at all, at chosen points in the ordinary sequence of taking instructions, establishing the facts, identifying the issues, analysing the law, drafting, advising, acting and keeping a record.[24] To choose those points, the lawyer has to see the work at the right level of detail. It helps to distinguish the legal service, the workflow that delivers it, the tasks within that workflow, and the smaller operations within each task.[25]

Analysis of legal automation at the level of whole tasks has a respectable history.[26] Its limitation is that a legal task is not a natural unit. 'Review the contract' or 'advise on termination' is shorthand for a bundle of quite different kinds of work. An estimate of how much of such a task can be automated says more about what happened to be in the bundle than about the work itself.[27] The better approach is to work backwards from the professional outcome rather than forwards from what the technology can do. The lawyer identifies what the client needs to understand, decide or achieve; then the document or action that will support it; then the tasks that will produce it; and only then the operations inside each task.[28]

Broken down in this way, the work reveals three things that are hidden when it is viewed task by task. First, some operations depend on a decision that has not yet been made, so a step that looks safe to hand to AI on its own is unsafe in its place in the sequence. Second, a single request to an AI system can merge several tasks into one fluent answer

[24] Feaver, 'AI-Enabled Lawyer' (n 3) Preface pt 3.

[25] Ibid Preface pt 4, table 1.

[26] Dana Remus and Frank Levy, 'Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law' (2017) 30 *Georgetown Journal of Legal Ethics* 501.

[27] Feaver and Collins (n 9) pt 2.1.

[28] Ibid; Feaver, 'AI-Enabled Lawyer' (n 3) ch 3 pt 3.

with no visible joins, so the lawyer cannot see what has gone wrong or what judgement has been quietly assumed. Third, judgement is not exercised once, at the end. It is exercised at identifiable points throughout the work.[29]

Lawyers have always worked through workflows, but the structure has mostly remained unspoken, carried in experience rather than written down.[30] Evidence from outside law suggests that AI works partly by capturing and passing on knowledge of this unspoken kind. In a large field study of customer-support staff, AI assistance improved the speed and quality of less experienced workers' output, while the most experienced saw small declines in quality.[31] For lawyers, the implication is that the unspoken structure of legal work has to be made explicit before AI can be controlled. The lawyer must be able to say where the technology may enter, what it may receive, what it may do and what human work must follow. Making that structure explicit is the first capability of the future lawyer. It carries a risk: a mapped workflow may be followed mechanically, like a checklist. The answer is proportion, together with the recognition that the same map serves as training scaffolding for a junior lawyer and as a supervision and record-keeping tool for an experienced one.[32]

C. Seeing the Skill

The second kind of understanding concerns legal skill. The profession has no agreed account of what legal judgement is. Commentators have observed that this gap becomes acute as AI absorbs the entry-level work through which judgement has traditionally been formed.[33] Frameworks for allocating work to AI reproduce the gap. They describe what AI can do in a precise, controlled vocabulary: extracting, organising, classifying, comparing, retrieving, transforming, drafting and critiquing. Each of these can be separately instructed, checked and rejected.[34] The work that remains with the lawyer, by contrast, is described afresh in ordinary language each time: 'assess

[29] Feaver and Collins (n 9) pt 2.2.

[30] Feaver, 'AI-Enabled Lawyer' (n 3) ch 1 pt 3.

[31] Erik Brynjolfsson, Danielle Li and Lindsey Raymond, 'Generative AI at Work' (2025) 140(2) *Quarterly Journal of Economics* 889 [issue number: confirm].

[32] Feaver and Collins (n 9) pt 2.3; Feaver, Collins and Southgate (n 11) pt 4.

[33] Thomson Reuters Institute, 'Honing Legal Judgment: The AI Era Requires Changes to How Lawyers Are Trained during and after Law School' (Web Page, 2026) [date: confirm].

[34] Feaver, 'AI-Enabled Lawyer' (n 3) ch 3 table 3.3; Feaver and Collins (n 9) pts 1.1, 3.1.

significance', 'resolve ambiguity', 'settle the substance'. Those descriptions may be accurate, but they are not categories. They cannot be compared across matters, linked to training, or used to tell a supervisor what to look for.

Supervisors experience the same gap from the other side. A survey of 106 Australian supervisors of first-year lawyers found that the qualities they most often found lacking were accuracy and attention to detail, named by 47 per cent, and practical judgement, named by 32 per cent. Neither appears in the lists of capabilities against which new lawyers are trained and assessed.[35] What supervisors find wanting is the part of legal work that the profession's own vocabulary does not name.

A theory of legal skill can supply the missing vocabulary. It describes a *skill architecture* as the particular combination of learned capacities needed to perform a legal task competently, safely and defensibly, and it proposes eleven of them, from noticing what matters to performing under pressure.[36] Almost every substantial piece of legal work draws on one of these in particular, complex problem-solving and judgement, so the broad label on its own tells a supervisor little.[37] The useful detail lies one level down. Complex problem-solving and judgement is made up of nine distinct thinking operations: framing the problem, analysing it, classifying facts and documents into legal categories, interpreting texts, bringing sources together, weighing considerations, choosing a strategy, reaching a judgement, and adapting what is known to a new situation.[38] At that level, it becomes possible to say that the work reserved for the lawyer in a particular step is the *interpretation* of a contractual term rather than the *classification* of a breach. That tells a supervisor where to look and what a failure would look like.

D. Where the Line Falls: Describing and Determining

Set the two vocabularies side by side and the division of labour between AI and lawyer can be stated rather than simply asserted. Several of the things AI can do correspond closely to particular thinking operations. Sorting material into categories the lawyer has set corresponds to classification, and comparing documents corresponds to analysis. But five thinking operations have no counterpart at all on the AI side: framing the

[35] Baxter, Collins and Feaver (n 13) pts 5.1, 6.1.

[36] Feaver, Collins and Southgate (n 11) pts 4–5, table 3.

[37] Feaver and Collins (n 9) pt 4.2.

[38] Feaver, Collins and Southgate (n 11) pt 7; Feaver and Collins (n 9) pt 4.3.

problem, interpreting texts, weighing considerations, choosing a strategy and reaching a judgement.[39] The two lists were developed separately, for different purposes and from different bodies of research. The fact that these five are left over is therefore a finding, not an assumption.[40]

What the five have in common supplies the line. Everything on the AI side works on material supplied to it and *describes* that material: it reports what a document says, how documents differ, what dates appear in a set of emails. Everything on the lawyer's side *determines* legal consequence: what the problem is, what a text means, how much a consideration weighs, what should be done and what position will be taken.[41] The line between describing and determining does not depend on what a system can do. Determining work is reserved to the lawyer not because it is difficult for a machine, but because it has legal effect and someone must answer for it. That matters because AI capability is uneven and changes with every release.[42] A line drawn around current capability would need to be redrawn constantly. This one does not. It also works at the margins. Applying categories that a lawyer has set is describing. Setting the categories is determining, because the categories carry the legal consequence.

One further element is not a skill at all. At the end of a matter, a lawyer adopts a position and authorises its use. That act is reserved to the lawyer because the lawyer holds an office and can be held to account. The paper calls this *standing*. The allocation of skills will shift as systems improve; standing will not.[43] Together, these ideas produce a simple four-part record for each step in a piece of legal work: the step itself, what AI may contribute, what the lawyer must determine or authorise, and what evidence shows that the lawyer actually did so.[44]

E. Ethics and Judgement, Located in the Work

The division of labour is also where professional ethics operates. The existing duties need no special supplement for AI, but they do need to be located. The standard of

[39] Feaver and Collins (n 9) pt 5.1.

[40] Ibid pt 5.2.

[41] Ibid pt 5.3.

[42] Dell'Acqua et al (n 23).

[43] Feaver and Collins (n 9) pt 5.4.

[44] Ibid pt 5.5.

competence and diligence in the *Legal Profession Uniform Law*, and the duties in the Conduct Rules, apply at different points in the workflow.[45] The duty of confidentiality is engaged when information is entered into a system, not only when output is sent out. The duty of independence is engaged when a lawyer decides whether to accept an analysis the system has produced. The duty of candour is engaged when material produced with AI is put before a court. And the duty of supervision is engaged in the process by which work is produced, not only in the final document.[46]

The Australian decisions bear this out, and scholarship on AI in litigation reads them in similar terms.[47] In *JNE24*, it was not enough to confirm that a cited case existed. The practitioner had to decide whether it supported the proposition, whether it was still good law and how it applied to the facts.[48] That is a failure of judgement about weight and application, not a failure to find the case. In *Murray*, the failure concerned the factual and documentary support for a filed document, and the Court required evidence of how the junior solicitor responsible had been supervised.[49] Research on AI legal tools explains why such failures recur. General-purpose models produced false legal information in a large share of answers to verifiable legal questions, and even specialised legal research tools produced incorrect or unsupported answers in a significant minority of cases.[50] Read through the division of labour, the decided cases are failures at the points where a lawyer had to determine something, hidden behind a document that looked finished.

This exposes the weakness of a nominal 'human in the loop'. Review offers little protection if the reviewer lacks the knowledge, information, time or authority to intervene.[51] The point is not confined to law. A study of 41 policies requiring human

[45] *Legal Profession Uniform Law Application Act 2014* (Vic) sch 1 ss 296–8 ('*Legal Profession Uniform Law*'); *Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015* rr 4.1.3, 4.1.4, 9.1, 17.1, 19.1, 37.1 ('*ASCR*').

[46] Feaver, 'Beyond Verification' (n 8) pt 4.2.

[47] Michael Legg, Vicki McNamara and Armin Alimardani, 'The Promise and the Peril of the Use of Generative Artificial Intelligence in Litigation' (2025) 48(4) *University of New South Wales Law Journal* 1196.

[48] *JNE24* (n 2) [20]–[25], [34].

[49] *Murray on behalf of the Wamba Wemba Native Title Claim Group v Victoria* [2025] FCA 731, [3]–[5] ('*Murray*').

[50] Matthew Dahl et al, 'Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models' (2024) 16(1) *Journal of Legal Analysis* 64; Varun Magesh et al, 'Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools' (2025) 22(2) *Journal of Empirical Legal Studies* 216.

[51] Feaver, 'AI-Enabled Lawyer' (n 3) ch 1 pt 6.

oversight of government algorithms found that they assume, without evidence, that people can oversee such systems effectively. In practice the policies give a false sense of security and allow responsibility to be shifted rather than discharged. The author proposes oversight by institutions rather than by individuals.[52] What the professional duties require is meaningful human legal control: the person present where the technology shapes the work must have, and must use, the skill and authority that point demands.[53] Ethics and judgement are therefore not a separate topic alongside AI capability. They are a question of where judgement sits in the work and who holds it.

F. Understanding Agents

The future lawyer also needs a working understanding of AI agents, and the discipline described so far supplies most of it. Before AI assists with any step, a careful lawyer settles nine matters: the task being advanced; the narrower operation to be assisted; why assistance is worthwhile; whether the operation is suitable; what information may be used and on what terms; what the system must not decide or assume; what output is required and how it must be traced to its sources; how the output will be checked; and who may adopt the result.[54] Read as a document, that note is a set of terms of delegation. It states the scope of the work, the permitted materials, the decisions the system may not make, the expected output, the plan for checking it and the person with authority to accept it. The questions that make a single step safe to hand to AI are the same questions that make a whole sequence of steps safe to delegate to an agent.[55] Understanding agents is therefore an extension of understanding workflow, not a separate technical skill. Part V returns to this.

G. Capability Is Shared

Finally, capability in this sense belongs to relationships between people, not to individuals alone. Whether a step should be performed by a system, a junior lawyer or a supervisor depends on what the step involves, what each person is capable of, what

[52] Ben Green, 'The Flaws of Policies Requiring Human Oversight of Government Algorithms' (2022) 45 *Computer Law & Security Review* 105681.

[53] Feaver, 'Beyond Verification' (n 8) pt 8.2.

[54] Feaver, 'AI-Enabled Lawyer' (n 3) ch 4 table 4.1; Feaver and Collins (n 9) pt 9.1.

[55] Feaver and Collins (n 9) pt 9.1.

authority each holds and what checking the work requires.[56] Capability, authority and responsibility are separate questions. A system may be able to compare two contracts, but it has no authority to decide which terms the client should accept. A junior lawyer may be able to prepare a preliminary analysis, but lack the authority to settle advice. A supervisor may hold the authority, but lack the knowledge, time or access to the sources needed to exercise it properly.[57] The individual lawyer's discipline cannot be sustained alone. It depends on how the firm allocates work, authority and control.

IV THE AI-ENABLED FIRM

A. From Lawyer to Firm

A lawyer may know how to use AI well and still work in a firm that has not decided which systems are approved, what information may be put into them, who supervises the work, or when a partner must approve the result.[58] The firm's task is different from the lawyer's. It must create conditions in which responsible use of AI is part of the ordinary way legal work is designed, supervised and approved.[59]

Research on AI in legal services shows what those conditions involve. A study combining case studies with survey data found that AI replaces people in some tasks, changes the work of those it does not replace, and creates new roles for legal experts in building and running AI-enabled services. Those roles typically sit in multidisciplinary teams, which traditional partnerships find harder to assemble.[60] AI also arrives in a particular kind of organisation: the law firm pyramid, in which the work of junior lawyers supports the partners, and the core of the partnership is increasingly reserved for those who control relationships with key clients.[61] Large firms do not yet expect that structure to change. Interviews with senior operating officers at ten of the largest firms in the United States found little expected change to billing or staffing, and an expectation that AI would move

[56] Feaver, 'AI-Enabled Lawyer' (n 3) ch 1 pt 5, table 1.1.

[57] Ibid ch 1 pt 6.

[58] Ibid ch 8 pt 1.

[59] Ibid.

[60] John Armour, Richard Parnham and Mari Sako, 'Augmented Lawyering' [2022] (1) *University of Illinois Law Review* 71.

[61] Marc Galanter and Thomas Palay, *Tournament of Lawyers: The Transformation of the Big Law Firm* (University of Chicago Press, 1991); Marc Galanter and William Henderson, 'The Elastic Tournament: A Second Transformation of the Big Law Firm' (2008) 60(6) *Stanford Law Review* 1867.

lawyers' time from gathering information to analysing it.[62] The argument here is that the pyramid may well survive in outline while what happens inside it changes considerably. Once many AI-enabled lawyers work together, the division of labour described in Part III becomes the central organising problem of the firm.

B. Two Divisions of Labour

It is tempting to allocate AI work by job title: juniors use the tools, supervisors check the output, partners sign. Allocating work by role, however, confuses the division of the work with the structure of employment, and invites a yes-or-no answer to what is really a question of design.[63] Yet roles do matter, because they carry authority. Regulatory theory offers a way through. Feaver and Sheehy show that the modern regulatory state contains two divisions of labour, one legal and one political, and that the relationship between them explains how control and accountability can come apart.[64] A firm has the same structure. One division allocates the performance of work. The other allocates authority over that work and accountability for it. The two must be kept in step.

The way to keep them in step is to fix the order in which work is allocated. The workflow settles what work exists. Breaking it down settles which steps are describing work and which are determining work. The determining steps, and the final act of standing behind the result, then decide which person must hold authority at each point. Four kinds of authority matter: authority to perform or assist with the work, to review it, to make the substantive legal or strategic decision, and to approve its communication, filing or other use.[65] An AI system may hold the first. The other three remain with people, allocated by the firm.

C. Junior Lawyers, Supervisors and Partners

Allocated in that order, the familiar levels of a firm take on definite content. Junior lawyers carry out safe AI-assisted work. 'Safe' can now be defined. It means describing work, done in an approved system on a defined set of source documents, checked item by item against those sources, and recorded. Junior lawyers also make some determinations

[62] Robert J Couture, 'The Impact of Artificial Intelligence on Law Firms' Business Models' (Harvard Law School Center on the Legal Profession, 25 February 2025).

[63] Feaver and Collins (n 9) pt 2.1.

[64] Feaver and Sheehy, 'The Shifting Balance of Power in the Regulatory State' (n 7).

[65] Feaver, 'Beyond Verification' (n 8) pt 7.

under delegation, because a junior lawyer is a legal practitioner and an officer of the court, and so holds a measure of standing in their own right.[66] Supervising lawyers allocate and verify. They settle the workflow, decide which steps AI may assist with, judge whether checked material is good enough for the next step, and supervise how the work was produced as well as what it says. Partners, or the lawyer responsible for the matter, provide assurance. They approve any use of the work that affects advice, filings or communication with others, and they hold the standing on which that use depends.

This definition of safe work matches what supervisors say they want. Asked how AI is changing their expectations of first-year lawyers, 71 per cent of supervisors said new lawyers need to know when AI use is appropriate, 58 per cent that they need to record or disclose it, 58 per cent that they need stronger judgement because AI output may be unreliable, and 56 per cent that they need to understand the risks to confidentiality and professional responsibility. Only a third pointed to greater productivity.[67] Supervisors want judgement about AI, not fluency with it.

Table 2 shows how one task in a supply dispute, the preparation of a chronology of deliveries from about 120 emails, would be allocated.[68]

Table 2: Allocating one task across the firm: a delivery chronology

Step	What AI may do	What a lawyer must determine	Who checks	Who approves use
Find references to delivery dates	Extract them from the nominated emails	Whether the set of emails is complete and reliable	Junior lawyer	Junior lawyer, as working material
Group emails about revised delivery dates	Organise them	Whether an extension was agreed, and whether it is legally effective	Junior lawyer against the emails; supervisor on what they mean	Supervising lawyer
Record explanations given for the delay	Extract and organise them	Whether each explanation is material, disputed or legally significant	Junior lawyer; supervisor	Supervising lawyer

[66] Feaver, 'What Can Be Written Down' (n 4) pt 8.3.

[67] Baxter, Collins and Feaver (n 13) pt 7.2, table 7.

[68] Adapted from Feaver and Collins (n 9) pt 7.1; Feaver, 'AI-Enabled Lawyer' (n 3) ch 7 pt 5, ch 8 pt 9.

Step	What AI may do	What a lawyer must determine	Who checks	Who approves use
Produce a draft chronology	Organise and draft	Whether the sequence is a sound basis for later analysis	Supervising lawyer	Supervising lawyer, for internal use only
Rely on the chronology in advice or correspondence	Nothing	Whether to adopt it for that use	Responsible lawyer	Partner or responsible lawyer

The picture so far may read as a description of a large firm. Most Australian lawyers do not work in one. The survey's respondents came mostly from small and medium practices, with 41 per cent in small firms or sole practices and only seven from large firms.[69] In a small practice the three levels may be one or two people. The functions remain even when the roles merge. A sole practitioner who uses AI still has the work performed, checks it against its sources, decides what it means, approves its use and can record each step. The structure scales down because it is built from functions, not job titles.[70]

D. The Problem of Many Hands

Levels of review add assurance only if each level checks something different. In *DPP v GR*, written submissions containing non-existent cases and incorrect references had passed through several lawyers. Each agreed with the conclusion while assuming that someone else had checked the foundations.[71] Responsibility was present on paper and absent in practice. Political philosophers call this the problem of many hands: when many people contribute to an outcome, it becomes hard to identify anyone who is responsible for it, and computer systems make the difficulty worse.[72] The same risk arises in a firm when the person who corrected the AI output is treated as if they had also supplied the judgement, the authority and the assurance needed for the work to be

[69] Baxter, Collins and Feaver (n 13) pt 3.2.

[70] Feaver, 'AI-Enabled Lawyer' (n 3) ch 8 pts 3–4.

[71] *Director of Public Prosecutions v GR* [2025] VSC 490, [61]–[80] ('*DPP v GR*'); Feaver, 'Beyond Verification' (n 8) pt 5.4.

[72] Dennis F Thompson, 'Moral Responsibility of Public Officials: The Problem of Many Hands' (1980) 74(4) *American Political Science Review* 905; Helen Nissenbaum, 'Accountability in a Computerized Society' (1996) 2(1) *Science and Engineering Ethics* 25.

used.[73] The columns of Table 2 are the remedy. They separate checking against sources, judging legal significance and approving use, and assign each to the person who holds the relevant authority.

E. How Much Supervision

Not every step needs the same supervision. Giving a step to AI makes sense only where the cost of the assistance plus the cost of checking it is less than the cost of doing the work without AI.[74] Extracting dates from a defined set of emails is slow to do by hand and quick to check. Assessing the weight of legal authority is the reverse: quick for a system to produce and slow to check, because checking it means doing the work. Supervisors already feel this cost. Forty per cent of those surveyed said AI increases their supervision burden because its output must be checked.[75] The intensity of supervision should follow how easily a failure could be detected, rather than how difficult the task is. Advice that is competent and current on its face, but rests on instructions that have since changed, looks exactly like correct advice.[76] At the level of the firm, these considerations become graduated levels of control, set by the consequences of error, the difficulty of checking and how easily a mistake could be reversed.[77]

F. The Controlled Environment

All of this happens within a tightly controlled technological environment, and the controls come before anyone writes a prompt. The firm must approve systems and, separately, approve uses. A platform suitable for some confidential work may be unsuitable for a particularly sensitive matter. A drafting tool that is suitable for internal work may not be authorised to send anything to a client.[78] Confidentiality, legal professional privilege and privacy must be addressed before material is put into a system, and any restrictions a client has imposed must be visible to the lawyer at the point of use rather than buried in an engagement letter.[79] Some workplaces have chosen

[73] Feaver, 'AI-Enabled Lawyer' (n 3) ch 7 pt 4.5.

[74] Feaver and Collins (n 9) pt 6.

[75] Baxter, Collins and Feaver (n 13) pt 7.2.

[76] Feaver and Collins (n 9) pt 7.3.

[77] Feaver, 'AI-Enabled Lawyer' (n 3) ch 8 pt 3, table 8.1.

[78] Ibid ch 8 pt 6.

[79] Ibid ch 8 pts 6, 8.

the simplest control: 17 per cent of the supervisors surveyed said first-year lawyers in their workplace are not permitted to use AI tools at all.[80] But an environment that does not meet lawyers' needs will be bypassed. In a 2026 international survey, a third of lawyers, accountants and compliance professionals said they had used AI that their organisation had not approved, and the proportion was higher among those who thought their organisation was moving too slowly.[81]

The controlled environment is where the division of labour becomes enforceable, and where oversight by individuals gives way to oversight by the institution, as the research on human oversight recommends.[82] Permissions determine who may do which steps, with which material. Logging records what was actually done. That second function becomes decisive when agents arrive.

G. Forming the Next Generation

The structure raises a problem it cannot solve on its own. Junior lawyers have traditionally developed judgement by doing the document review, chronologies, preliminary research and first drafts that AI now absorbs.[83] The Australian survey shows the change under way. Forty per cent of supervisors report first-year lawyers doing AI-assisted work, a third report that some traditional junior tasks are reducing or changing, and 12 per cent say AI is reducing, or is expected to reduce, the number of first-year lawyers they employ.[84] International surveys show the same concern. Nearly half of professionals worry about AI's effect on the development of independent judgement, and legal professionals expect the time needed to reach trusted judgement to lengthen by almost two years.[85] A survey of nearly 900 lawyers in the United Kingdom found that 69 per cent identified checking and verifying sources as a gap in junior lawyers' skills.[86]

Supervisors do not expect judgement to be complete on arrival. A majority place professional judgement (57 per cent) and autonomy (66 per cent) at the end of the period

[80] Baxter, Collins and Feaver (n 13) pt 7.2.

[81] Thomson Reuters Institute, *Future of Professionals Report 2026* (Report, June 2026).

[82] Green (n 52).

[83] Feaver, 'AI-Enabled Lawyer' (n 3) ch 1 pt 6, ch 8 pt 11.

[84] Baxter, Collins and Feaver (n 13) pts 4.1, 7.2, 7.5.

[85] Thomson Reuters Institute, *Future of Professionals Report 2026* (n 81).

[86] LexisNexis UK, *The Mentorship Gap* (Report, 2026).

of supervised practice, not at its start.[87] Sixty per cent reported shortfalls in capabilities they themselves expect to develop only during supervised practice, and three-quarters attributed the gap to limited exposure to realistic legal work before admission.[88] This is consistent with what is known about how expertise develops. Judgement emerges through responsible engagement with real situations, and newcomers learn by taking a real but bounded part in the work of experienced practitioners.[89] If AI absorbs the bounded work, the route to judgement narrows. A study of robotic surgery found exactly this: when technology allowed senior surgeons to operate without a trainee, the approved routes to learning failed, and only a minority of trainees reached competence, largely by learning in ways their institutions did not sanction.[90]

The legal evidence is more encouraging, and it points to design. In a randomised trial, law students who used AI to summarise legal materials performed no worse on a later test of comprehension, and better on a later problem-solving task, even though neither group then had access to AI. But students who began with strong reasoning often produced weaker work when they used AI to revise it. The authors conclude that AI's effect on legal reasoning depends on when and how it is used.[91] Field evidence from other work points the same way: AI assistance helped newer workers most.[92] Formation is a design problem, not an inevitable casualty.

The answer is not to preserve manual work for its own sake but to design workflows for development. A junior lawyer might review a representative sample of documents, set the categories, use AI across the larger set, check the output and explain the errors and their legal significance to a supervisor.[93] The order matters, for the reason the trial evidence suggests: the lawyer forms a view before seeing the system's. The four-part

[87] Baxter, Collins and Feaver (n 13) pt 4.5, table 3.

[88] Ibid pts 5.4, 6.3.

[89] Stuart E Dreyfus, 'The Five-Stage Model of Adult Skill Acquisition' (2004) 24(3) *Bulletin of Science, Technology & Society* 177; Jean Lave and Etienne Wenger, *Situated Learning: Legitimate Peripheral Participation* (Cambridge University Press, 1991).

[90] Matthew Beane, 'Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail' (2019) 64(1) *Administrative Science Quarterly* 87.

[91] Nicholas Bednar et al, 'Artificial Intelligence and Human Legal Reasoning' (2027) 121 *Northwestern University Law Review* (forthcoming).

[92] Brynjolfsson, Li and Raymond (n 31).

[93] Feaver, 'AI-Enabled Lawyer' (n 3) ch 8 pt 11.

record described in Part III also serves as a training tool, because it tells a supervisor exactly which judgement a junior lawyer is being asked to exercise at each point.[94] Medical education offers a developed model for extending authority step by step. Under it, defined units of professional work are entrusted to a trainee once the trainee has shown the competence to perform them without supervision.[95]

The survey adds one more lesson. The gaps that cost supervisors most were not technical. They concerned the supervisory relationship itself: communicating with supervisors, managing time, knowing when to ask for help, and applying feedback.[96] Being supervised is a capability, and it can be taught. Australian scholarship has long called for supervision itself to be treated as a skill to be developed.[97] As Part V shows, knowing when to stop and escalate is also the capability that matters most when work is delegated to an agent. The same points at which an agent must stop and refer a decision to a person can be written into a junior lawyer's workflow, so that the question of when to ask for help has a clear answer.

H. The Structure That Results

Taken together, the AI-enabled lawyer and the AI-enabled firm produce a structure with four parts. The first is a mapped workflow that separates describing work from determining work. The second is an allocation of authority over each determination to a particular person. The third is a controlled environment that permits, restricts and records the use of AI. The fourth is a system of supervision and formation that keeps people capable of doing, and judging, the work themselves. This is the structure that Part V tests against AI agents.

V CONTAINING AGENTIC WORK

A. What Agents Change

When AI assists with a single step, the lawyer exercises control each time: deciding, instructing, receiving, checking and verifying. When work is delegated to an agent,

[94] Feaver and Collins (n 9) pts 1.1, 3.2.

[95] Olle ten Cate, 'Entrustability of Professional Activities and Competency-Based Training' (2005) 39(12) *Medical Education* 1176.

[96] Baxter, Collins and Feaver (n 13) pts 5.3, 8.3.

[97] Jeff Giddings and Michael McNamara, 'Preparing Future Generations of Lawyers for Legal Practice: What's Supervision Got to Do With It?' (2014) 37(3) *University of New South Wales Law Journal* 1226.

control is exercised in advance. The lawyer authorises a specification, and the system works within it.[98] The judgement involved is the same, covering what sources may be used, what the system may not decide and how the result will be checked. But it is made once, in advance, for every occasion on which the agent runs. This is how supervision has always developed as trust is established. A partner does not approve every letter a senior associate sends. The partner settles the terms on which letters may be sent, and checks a sample.

Agency theory explains why control in advance is necessary, and why it is hard. Kolt identifies in AI agents the classic problems of delegation: the agent knows things the principal does not, exercises discretion, and may not act with undivided loyalty. He argues that the usual remedies, which are incentives, monitoring and enforcement, may fail where agents make decisions that cannot easily be interpreted, at great speed and scale.[99] Delegation also brings a temptation. Feaver and Sheehy show that governments choose forms of agency that maximise their control while minimising their legal and political accountability for what the agency does.[100] A firm delegating to AI agents faces a similar incentive: to keep the benefit of the work while spreading the responsibility for it thinly. The argument of this Part is that the structure described in Part IV, and the lawyer's standing within it, closes off that option.

B. Testing the Structure: An Ongoing Research Program

The authors are conducting a research program that tests the framework by asking an AI agent to complete legal matters under controlled conditions.[101] The first studies have used fictional commercial disputes. For example, a food manufacturer had taken late delivery of a packaging line. The supplier said the delivery date had been changed by agreement in a telephone call, and that an exclusion clause in the contract defeated the claim in any event. The matter was broken down into nine tasks and 41 steps. For each step, the workflow stated what AI could safely do and, where judgement was required, what had to be decided, which skills that decision drew on and which authorities set the

[98] Feaver and Collins (n 9) pt 9.2.

[99] Kolt (n 4).

[100] Feaver and Sheehy, 'The Political Division of Regulatory Labour' (n 7).

[101] Feaver, 'What Can Be Written Down' (n 4) pt 3.

standard.[102] A single AI agent was given the file and the workflow and instructed to complete the matter without any human decision. It played both the AI assistant and the lawyer. Where a decision belonged to the supervising partner or the client, it simulated that decision and labelled it as simulated.

The agent completed the matter. It produced an issues map for the partner, a letter of advice to the client, an email for the client to send to the supplier and a note of how AI had been used. It recorded 47 separate judgements.‡[103] It noticed things nobody had planted in the file. One was that the contract's commissioning period meant the client's planned trial with a supermarket customer might have been lost however punctually the line was delivered. Another was an inconsistency in the date of one document that the scenario's authors had missed. It went back to earlier tasks multiple times to correct itself, stopped the line of legal research it could not carry out safely without access to sources, and declined to certify its own letter as fit to send. In a second run scenario with access to legal sources, it re-examined the 26 legal statements in its letter. Six changed, and every change went against the client.‡[104]

These are early results. They come from similar matters and early-stage runs of the program. Further runs in fresh sessions, blind scoring of the agent's work against points deliberately planted in the file, independent coding of the agent's records, and comparison with junior lawyers working on the same file are part of the continuing program.[105] The results are reported here because they show the framework at work, not because they settle any question about how often, or how reliably, agents perform.

C. Simulated Judgement and What It Cannot Supply

At each point where the workflow required judgement, the agent made a determination. It decided what the client's real problem was, how much weight a diary note written at the time could bear against a later recollection, whether the claim was worth pursuing, and what the email to the supplier could properly say. Its records show reasoning that can be examined and disputed. In one, it took a weakness in the client's position, which

[102] 'Harbour Foods: Matter Completion Workflow' (Unpublished, 27 September 2026) ('Harbour Foods Workflow'); Feaver, 'What Can Be Written Down' (n 4) pt 3.1.

[103] Feaver, 'What Can Be Written Down' (n 4) pts 4–6.

[104] Ibid pt 7.1, table 3.

[105] Ibid pt 11, app C.

was that the client had left the supplier's emails unanswered, and applied it to the client's own email, which therefore could not leave the supplier's reliance on the exclusion clause unanswered.[106] In another, it rejected its simulated partner's suggestion to remove the analysis of causation from the letter, because without it the client could not make an informed choice.[107]

This is simulated judgement: the agent performs, in writing, the thinking operations that the division of labour reserves to the lawyer. It does not undermine the line between describing and determining. It confirms why that line was drawn without reference to what systems can do. The agent can produce a determination, but what it produces is a proposal about legal consequence. It has no effect until someone with standing adopts it. Simulated judgement also covers only the part of judgement that can be put into words. It does not reach what lies outside the file, such as the client's appetite for risk or the history and feel of a commercial relationship. The agent could reason about these things, but it could not know them.[108]

Nor could the agent supply the lawyer's position. It could act out the partner's approval and the client's instructions, but it could not give them. It had no standing to decide and no accountability if it was wrong. It also lacked commitment: the act of putting one's name to a view and being bound to correct it if it turns out to be wrong, an obligation the law already recognises.[109] These are not features of the work, so no better model will supply them. They belong to the person who holds the office.[110]

D. Four Elements of Containment

The structure described in Part IV contains agentic work through four elements. Each corresponds to something a well-run firm already does for work done by people. Together, they keep an agent's work within limits set in advance, stop it where a lawyer must decide, record what it did, and leave the decision to rely on the work with a person who can answer for it.

---

[106] Ibid pt 4.2, table 2.

[107] Ibid pt 6.1.

[108] Ibid pts 8, 9.3.

[109] Ibid pt 8.5; *Myers v Elman* [1940] AC 282 [pinpoint: confirm]; *ASCR* r 19.2.

[110] Feaver, 'What Can Be Written Down' (n 4) pt 8.6.

1 - Specification: What the Agent May Do

The first element is a specification, authorised in advance by a person with authority to approve it. It is the nine-point design note described in Part III, written once for a class of work rather than composed afresh each time.[111] It sets the scope of the work, the sources the agent may use, the decisions it may not make, the form its output must take, how that output will be checked and who may accept it. Because the usual remedies for problems of delegation may not work well against agents operating at speed, the specification has to carry more of the load than the instructions given to a human assistant.[112] Over time, a firm can build a library of specifications for recurring kinds of work, such as preliminary contract review, chronologies in disputes or due diligence, each already approved for delegation.[113] That library carries a warning. A specification applied without adapting it to the facts, sources and risks of the particular matter is the checklist problem described in Part III, repeated at industrial scale.

2 - Stopping Points: Where the Agent Must Stop

The second element is a set of stopping points. The line between describing and determining marks where they fall. A continuous run of describing work can be delegated as a block. A step that determines legal consequence ends the block, and a lawyer must decide before the work goes on. A second rule limits how long a block may run. Errors carry forward within a block, so a block may continue only as long as each output can still be traced back to its source documents. Where that trail is lost, the block must stop.[114]

The simulation shows stopping points at work. The agent stopped the line of research it could not do safely, went back to correct itself, and declined to certify its own letter. It also showed where a real person must step in: the partner's decisions, the client's instructions, and, at the opening of the file, the approval of the tool, the search for conflicts of interest and the disclosure of costs. In the simulation each of these was performed by simulation and labelled as such. In practice they are the points at which

[111] Feaver and Collins (n 9) pts 9.1, 9.4.

[112] Kolt (n 4).

[113] Feaver and Collins (n 9) pt 9.4.

[114] Ibid pt 9.3.

supervision and sign-off must actually occur.[115] The same idea applies to people. Knowing when to ask for help is one of the gaps that costs supervisors most.[116] The stopping points written into an agent's specification can equally be written into a junior lawyer's workflow.

3 - The Record: What the Agent Did

The third element is a record. Delegation promises a better record than work done by people, which rarely documents the judgement exercised. A delegated workflow can record every step automatically.[117] But there is an important qualification. An agent's own account of its reasoning is a reconstruction, not a trace of what actually happened.[118] Research on language models confirms that the reasoning they state can be plausible yet misleading: it can omit the factors that actually drove an answer, and rationalise answers that were pushed towards error.[119] A note written by the AI about its own use would be no protection in a real matter.[120] The record must therefore be captured by the firm's controlled environment, independently of what the agent says about itself, using measures of the kind proposed for governing AI agents generally: identifying when an agent is acting, monitoring it as it works, and keeping logs of its activity.[121] The record has value only when a person with standing makes it or adopts it.

4 - Adoption: Who Stands Behind the Result

The fourth element is adoption. Work done by an agent becomes professional work only when a lawyer with standing adopts it and authorises its use. Adoption depends on verification, and agentification changes what verification involves in four ways.

It checks decisions, not just documents. When a person does the work, a reviewer reads the finished document and infers the judgement behind it. When an agent does the work, each judgement can be recorded: the condition that had to be met before the work could

---

[115] Feaver, 'What Can Be Written Down' (n 4) pts 6, 10.1.

[116] Baxter, Collins and Feaver (n 13) pt 5.3.

[117] Feaver and Collins (n 9) pt 9.5.

[118] Feaver, 'What Can Be Written Down' (n 4) pt 9.3.

[119] Miles Turpin et al, 'Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting' (Conference Paper, Conference on Neural Information Processing Systems, 2023).

[120] Feaver, 'What Can Be Written Down' (n 4) pt 10.1.

[121] Alan Chan et al, 'Visibility into AI Agents' (Conference Paper, ACM Conference on Fairness, Accountability, and Transparency, 3–6 June 2024).

go on, the decision made, stated so that it could turn out to be wrong, the alternative rejected, the skill involved, the authority applied, and whether a source could settle the point.[122] The reviewer can examine the decisions directly. The evidence of judgement lies in the reasoning, not in the labels. Where the workflow named the skill and authority for each decision, the agent's records simply repeated them, so it was the reasoning that showed judgement being exercised.[123]

It looks hardest at what favours the client. Without access to legal sources, the agent's reasoning drifted towards its client. Its errors were not invented cases, which a citation check would catch, but plausible arguments that happened to be wrong, which a citation check would pass.[124] Research on language models suggests a possible explanation. Models trained on human feedback tend to give answers that match what the user already believes rather than answers that are true, apparently in part because people rate agreeable answers more highly.[125] Whether a similar pull operates towards the party an agent is acting for has not been tested, and should be treated as a hypothesis. The practical lesson does not depend on the explanation. Every legal proposition in agentic work should show on its face whether it has been verified, and the reviewer should test hardest the arguments that help the client, because those are the most likely to have been accepted without a source.[126]

It is independent of the system that did the work. In the simulation the same model did the work and checked it, so an error shared by both roles could not be caught, and the safeguard of the lawyer forming a view before seeing the system's was observed only in form.[127] *JNE24* points the same way. There one AI system had been used to check submissions researched with another, and the Court treated verification as a matter of legal judgement rather than a technical check.[128]

---

[122] Feaver, 'What Can Be Written Down' (n 4) pt 3.3.

[123] Ibid pts 9.1–9.2.

[124] Ibid pt 7.1.

[125] Mrinank Sharma et al, 'Towards Understanding Sycophancy in Language Models' (Conference Paper, International Conference on Learning Representations, 2024).

[126] Feaver, 'What Can Be Written Down' (n 4) pts 10.1–10.2.

[127] Ibid pt 9.5.

[128] *JNE24* (n 2) [14], [20]–[25]; Feaver, 'Beyond Verification' (n 8) pt 5.2.

It depends on a verifier who could have done the work. A lawyer cannot answer for work they cannot assess, and cannot supervise an agent without knowing where judgement sits in the work and what it involves.[129] The agent in the simulation could be instructed only because the workflow had been mapped and the skills named. The same knowledge is what a lawyer needs to judge whether the agent's work is sound. Research on automation in other fields shows why this is hard to sustain. When most of a task is automated and people are left responsible for the rest, they lose the practice that keeps their skills sharp, so the more advanced the system, the more critical the human contribution becomes.[130] People who monitor automated systems also tend to trust them too much, and that tendency grows with workload and time pressure.[131] Clinical evidence suggests the effect is real. After routine exposure to AI assistance in colonoscopy, specialists' detection rate in procedures performed without AI fell from 28.4 per cent to 22.4 per cent.[132] The legal trial evidence counsels against treating such decline as inevitable.[133] But the survey adds a pointed observation. Checking AI output is accuracy and attention to detail by another name, and that is the capability supervisors most often find lacking in new lawyers.[134] Firms that delegate a workflow must still form lawyers who could perform it.

Adoption must also be real rather than nominal. Elish describes how blame for a failure in a complex automated system can fall on the person nearest to it, who had little control over how it behaved. That person becomes a 'moral crumple zone', absorbing the impact in order to protect the reputation of the technology.[135] A partner who carries formal accountability for a delegated workflow, without the knowledge, access or authority to

---

[129] Feaver, 'What Can Be Written Down' (n 4) pt 10.3.

[130] Lisanne Bainbridge, 'Ironies of Automation' (1983) 19(6) *Automatica* 775.

[131] Raja Parasuraman and Dietrich H Manzey, 'Complacency and Bias in Human Use of Automation: An Attentional Integration' (2010) 52(3) *Human Factors* 381; Christopher D Wickens et al, 'Complacency and Automation Bias in the Use of Imperfect Automation' (2015) 57(5) *Human Factors* 728; Kate Goddard, Abdul Roudsari and Jeremy C Wyatt, 'Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators' (2012) 19(1) *Journal of the American Medical Informatics Association* 121.

[132] Krzysztof Budzyń et al, 'Endoscopist Deskilling Risk after Exposure to Artificial Intelligence in Colonoscopy: A Multicentre, Observational Study' (2025) 10(10) *Lancet Gastroenterology & Hepatology* 896.

[133] Bednar et al (n 91).

[134] Baxter, Collins and Feaver (n 13) pt 8.5.

[135] Madeleine Clare Elish, 'Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction' (2019) 5 *Engaging Science, Technology, and Society* 40.

shape it, is in exactly that position. The specification is what prevents it. A partner who has approved the specification has exercised judgement over the work, rather than merely lending it a name.

E. The Partner's Role: Assuring Systems

The partner's assurance role changes accordingly. Under agentification, it becomes assurance of systems as well as of work: approving specifications, setting the points at which a person must decide, deciding how much of the agent's work to sample, and revising the specification as experience builds up.

Regulatory practice has begun to adopt this structure. In May 2025 the Solicitors Regulation Authority authorised the first law firm in England and Wales to deliver regulated services entirely through AI, a service for recovering unpaid debts.[136] The conditions attached correspond closely to the four elements described here. The system is not autonomous and takes each step only with the client's approval. It is configured not to propose case law. And named solicitors remain accountable for the firm meeting professional standards.[137] Excluding case law removes the step at which the trail back to sources most often breaks. Requiring client approval at each step places a stopping point at each boundary. Naming accountable solicitors keeps standing with people.

F. The Limits of Containment

The claim that the profession already has the means to contain agentic work is a strong one, and it has limits. Six deserve to be stated plainly.

Agents that act in the world. The structure works best where an agent's output is reviewed before anything happens. Agents that send communications, file documents or carry out transactions can create consequences before any review occurs. Containment then requires that every external action be a stopping point by design, as the conditions imposed on the first AI-driven firm in England and Wales require.[138]

Environments the firm does not control. The argument assumes that the firm controls its technological environment. Much of it is in fact controlled by vendors, who may make

[136] Solicitors Regulation Authority (n 16).

[137] Ibid; 'Legal Regulator Approves First AI-Driven Law Firm', *Global Legal Post* (online, May 2025) [date: confirm]; 'SRA Approves "AI-Driven" Law Firm', *Legal Cheek* (online, May 2025) [date: confirm].

[138] Solicitors Regulation Authority (n 16); 'Legal Regulator Approves First AI-Driven Law Firm' (n 137).

material decisions about confidentiality and system behaviour while the lawyer remains professionally responsible for the consequences.[139] The firm's control is therefore partial, and must be supplemented through contract, due diligence and the choice of systems.

Speed and scale. Supervision by sampling assumes that the volume of work allows meaningful samples to be taken and reviewed. Agents working at speed and scale may outrun it.[140] Monitoring will increasingly need to be built into systems, rather than left to after-the-event review.[141]

Commercial pressure. The law firm pyramid depends on the leverage provided by junior lawyers' work.[142] As agents take on that work, firms will face exactly the temptation described earlier: to keep the benefit while spreading the accountability.[143] Professional regulators may need to require firms to make their allocation of authority over agentic work visible.

The erosion of capability. Containment depends on people who are capable of the work they verify. If firms allow the routes to judgement to close, verification will become nominal and containment will fail with it.[144]

The evidence base. The evidence offered here is early. The simulation covers one matter in an ongoing program. The survey reports the perceptions of supervisors drawn from a sample that was not random.[145] The framework itself has not yet been tested as a tool that independent lawyers can apply consistently, and the authors have proposed a study to do so.[146]

None of these limits defeats the argument. Each marks a place where the structure needs reinforcement, and where firms, regulators and courts should look most closely.

---

[139] 'The Lawyer in the Loop Is Not Enough', *Law Technology Today* (American Bar Association, September 2026) [author and date: confirm].

[140] Kolt (n 4).

[141] Chan et al (n 121).

[142] Galanter and Palay (n 61); Galanter and Henderson (n 61).

[143] Feaver and Sheehy, 'The Political Division of Regulatory Labour' (n 7).

[144] Bainbridge (n 130); Beane (n 90).

[145] Baxter, Collins and Feaver (n 13) pt 3.4.

[146] Feaver and Collins (n 9) pt 8.

VI CONCLUSION

The principle that the lawyer remains responsible becomes useful once it is specified. This paper has traced how that can be done. The AI-enabled lawyer sees legal work as a workflow, divides it into steps, and separates the describing work that AI may do from the determining work that a lawyer must do. The AI-enabled firm turns that division into a structure: work allocated through the workflow to the person holding authority, levels of review that each check something different, a controlled environment that permits, restricts and records, and a system of formation that keeps people capable of the work. AI agents can perform almost all of a legal matter, including the part of judgement that can be written down. The same structure contains them, through a specification authorised in advance, stopping points where a lawyer must decide, a record captured by the firm rather than by the agent, and adoption by a lawyer who stands behind the result.

What remains with the lawyer, at the end of each stage, is the same. It is not a set of tasks too hard to automate. It is a position: the standing to decide, accountability for the result and the commitment to stand behind it. At each stage, the division of labour that allocates the performance of work has moved. The division that allocates authority and accountability has had to stay where it was.

Regulatory theory suggests that systems fail where their underlying commitments are not translated coherently into the arrangements that are meant to deliver them.[147] The profession's duties are the commitments. The workflow, the firm's allocation of authority and its controlled environment are the arrangements through which those duties are carried out. That has consequences for each institution involved.

For legal education and practical legal training, now under review, it means teaching legal workflows and the specific skills of judgement; teaching what it is to be supervised; and teaching AI as a matter of professional responsibility rather than tool proficiency.[148] For firms, it means building the allocation of authority, the controlled environment and the arrangements for approving specifications, and preserving the work through which

[147] Benedict Sheehy and Donald Feaver, 'Designing Effective Regulation: A Normative Theory' (2015) 38(1) *University of New South Wales Law Journal* 392; Donald Feaver and Benedict Sheehy, 'Designing Effective Regulation: A Positive Theory' (2015) 38(3) *University of New South Wales Law Journal* 961.

[148] Baxter, Collins and Feaver (n 13) pt 9.1; Legal Profession Admission Board of New South Wales (n 15); Legal Services Council Admissions Committee and Law Admissions Consultative Committee (n 15).

new lawyers learn. For regulators, it means locating existing duties at the points in the work where they apply, and expecting supervision of the process by which work is produced as well as of the finished product.[149] For courts, it suggests that when AI-assisted material is before them, the useful questions are where AI entered the work, who determined what, and what record shows it.[150]

The research program on which this paper draws is continuing, and its findings will be reported as they settle. The central claim, however, does not depend on them. AI changes who, or what, does legal work. It does not change who must answer for it. The future lawyer's task is to build the structure that keeps those two things aligned, and to hold it.

[149] Feaver, 'Beyond Verification' (n 8) pt 8.

[150] Ibid pts 5, 7.