



AI 2035: The Legal Profession and the Judiciary in the Age of Artificial Intelligence

Drafting with AI: Quality Control, Confidentiality, and Contract Risk Allocation

This session explores how transactional lawyers can effectively leverage AI in contract drafting while maintaining accuracy, protecting client confidentiality, and managing risk. Attendees will learn practical strategies for quality control, ethical use, and allocating AI-related risks in modern agreements.



Benjamin Altshul, Partner, Levenfeld Pearlstein

Benjamin Altshul is a partner at Levenfeld Pearlstein in the Real Estate Group, representing investors, developers, and corporate clients in commercial real estate transactions nationwide. Benjamin advises on acquisitions, dispositions, financing, and leasing across diverse asset classes. He is known for combining legal insight with a strong understanding of real estate economics. Benjamin also serves in leadership roles with the Jewish Council for Youth Services.



Michael Riesen, Partner, Smith Gambrell & Russell

Michael Riesen is a partner at Smith Gambrell & Russell and Executive Director of SGR Labs. Michael is a U.S. patent attorney and U.K. solicitor with formal training in artificial intelligence and machine learning. He advises clients on AI governance, intellectual property, and technology transactions. Michael works with Fortune 500 companies and emerging technology businesses. He bridges technical expertise with practical legal strategy.



Dev Thain, Program Manager, Spellbook

Dev Thain is the Head of Academic Partnerships at Spellbook, where he leads partnerships with law schools, bar associations, and legal institutions to drive adoption of AI-powered legal tools. He previously spent six years at Gartner advising enterprises and governments on transformation, technology strategy, and AI adoption. He holds a Master's in Learning Technology from the University of Oxford. Dev regularly speaks on practical applications of AI in legal practice, most recently at Harvard Law School and the American Bar Association.



Moderator: Judge Tracie Porter, Circuit Court of Cook County

Judge Tracie Porter currently sits in the Child Protective Division of the Circuit Court of Cook County. Judge Porter previously sat in the Traffic Division where she presided over minor traffic violations and Class A Misdemeanor matters. Immediately prior to her appointment as judge, Judge Porter practiced law in her own firm Tracie R. Porter, P.C. In her law practice, she represented clients in both transactional and litigation in matters involving business contracts, residential and commercial real estate, estate planning, probate, and entertainment law.

ATTENDEE MATERIALS

THE CHICAGO BAR ASSOCIATION PRESENTS:

AI 2035: The Legal Profession and the Law in the Age of Artificial Intelligence

Drafting with AI

Quality Control, Confidentiality, and Contract Risk Allocation

May 11, 2026 · Venue SIX10 · Chicago

Panel Benjamin Altshul · Michael Riesen · Dev Thain



THE
**CHICAGO
BAR
ASSOCIATION**

PANEL OVERVIEW

What We'll Cover

How lawyers can use AI drafting tools responsibly — protecting clients, preserving privilege, and allocating risk. Organized around three lenses that follow the drafting workflow itself.



INPUT

01

Quality Control

What goes into the draft — and how you verify what comes out of it.



PROTECTION

02

Confidentiality

How you safeguard client data throughout the workflow.



OUTPUT

03

Risk Allocation

Who bears the consequences when something goes wrong.

TIMED AGENDA

90 Minutes, Six Segments

May 11 · Venue SIX10 · 10:30 AM – 12:00 PM

TIME	TOPIC
10:30 – 10:35	Opening — Thesis & framework
10:35 – 10:50	Topic 1 — Baseline Reality. Where AI is actually being used; the three-stage playbook.
10:50 – 11:05	Topic 2 — Quality Control. Verification; <i>Mata v. Avianca</i> ; <i>In re Farris</i> .
11:05 – 11:20	Topic 3 — Confidentiality. Rule 1.6(e); <i>Heppner v. Wal-Mart</i> ; the Kovel bridge.
11:20 – 11:35	Topic 4 — Risk Allocation. Where liability sits; engagement letters; insurers.
11:35 – 11:45	Topic 5 — Strategic & Future-State. Junior-lawyer training; the 3–5-year outlook.
11:45 – 12:00	Q&A — Open discussion with the panel

Lawyers are not getting in trouble because they are using AI.

They are getting in trouble because they are using AI without understanding how these tools actually work — and without bothering to review the output.

Every point in this presentation ties back to two existing obligations: the duty of competence in technology, and the duty to supervise nonlawyer assistants.

ORGANIZING FRAMEWORK

Input → Protection → Output

Not separate topics. Sequential stages of the same workflow — what you put in, how you protect it in transit, and who bears the consequences of what comes out.



01

INPUT · QUALITY CONTROL

What goes in

Verification workflows
Failure-mode awareness



02

PROTECTION · CONFIDENTIALITY

How it's protected

Vendor evaluation
Privilege preservation
Tool selection



03

OUTPUT · RISK ALLOCATION

Who is liable

Engagement letters
Insurance posture

GOVERNING AUTHORITY

Three Illinois Rules Govern Everything

Nothing we discuss today is a new ethical rule. These are existing rules applied to a new tool. Full text appears in Appendix A.



RULE 1.1

Competence in Technology

A lawyer's duty of competence includes keeping abreast of the benefits and risks associated with relevant technology. (Comment [8])



RULE 1.6(E)

Confidentiality of Information

A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation.



RULE 5.3

Nonlawyer Assistants

Lawyers are responsible for ensuring that the conduct of nonlawyer assistants is compatible with the lawyer's professional obligations. AI is the newest nonlawyer assistant you have ever hired.

PART 1

01

Baseline Reality

Where AI is actually being used — and where it is still falling short.

Five practical Use Cases Already in Play

01

Dictation & voice-to-text

Rapid capture of drafting ideas, client notes, and first-pass content.

02

Notetaking & summary

Automated capture of calls, meetings, and conference notes — with caveats.

03

Communication filter

Tone, clarity, and audience adaptation for client-facing written work.

04

Spellcheck on steroids

Grammar, consistency, defined-term tracking, and internal formatting.

05

Playbook first passes

First drafts of standard clauses against a defined negotiating playbook.

WHERE IT FALLS SHORT

Bespoke deal terms. Novel risk allocation. Anything requiring business-context judgment.

What's AI-ready vs. Still Too Risky

AI-READY

High-volume, repeatable contracts

- Standard NDAs
- Standard commercial leases
- Vendor and service agreements
- Playbook-driven first-pass drafts of familiar clauses
- Issue-spotting against a known playbook



TOO RISKY FOR UNSUPERVISED AI

Bespoke, novel, or high-consequence provisions

- Bespoke M&A provisions
- Novel indemnity structures
- Jurisdiction-specific regulatory language
- Limitation of liability and risk-allocation clauses
- Anything requiring business-context judgment



The Three-Stage Playbook

Playbooks unlock the highest value. Build trust in AI incrementally — do not start with AI writing your first draft.

STAGE 1

Issue Spotting

AI ROLE

Flags deviations from standard positions.

LAWYER ROLE

Reviews 100% of AI output.

STAGE 2

Attorney-Led Drafting

AI ROLE

Suggests language based on the playbook.

LAWYER ROLE

Selects, edits, and approves each provision.

STAGE 3

AI First Draft

AI ROLE

Generates initial draft of defined clause types.

LAWYER ROLE

Reduced review — only after Stages 1 & 2 yield documented accuracy rates.

What “Good Output” Actually Looks Like

AI output that reads as if a competent associate drafted it — still requires partner review.

- ✓ **Defined review checkpoints**
Built into the workflow, not left to individual discipline.
- ✓ **Internal consistency of terms**
Defined terms used consistently throughout the document.
- ✓ **Accurate cross-references**
Section numbers, exhibit labels, and schedule references all resolve.
- ✓ **Date and party-name consistency**
Parties and effective dates align across all provisions and exhibits.
- ✓ **Format and numbering integrity**
Outlines, margins, and numbering remain clean across edits.
- ✓ **No confident-sounding fabrication**
Citations and quotations trace back to the actual source.

PART 2

02

Quality Control

Trust, verification, and the failure modes you actually need to worry about.

Failure Modes: Catchable vs. Convincing

Two decisions tell the story of how AI failure modes have evolved. The first is easy to catch. The second is much harder.

CASE 1 · 2023

FABRICATED CASES

Mata v. Avianca

S.D.N.Y.

A lawyer used ChatGPT for federal bankruptcy-tolling research. The AI fabricated six case citations — complete with docket numbers and quotes — that did not exist. The lawyer never verified. The brief was filed. Opposing counsel flagged the fabrications.

Sanctions: \$5,000 per attorney; required notice to every judge in the fabricated citations.

Relatively easy to catch — the cases themselves did not exist.

CASE 2 · 2025

REAL CASES, FABRICATED QUOTES

In re Farris

C.D. CAL.

A federal public defender used CoCounsel — a legal-specific AI tool. The cases cited were real. But the quoted passages attributed to those cases were fabricated. The AI pinned persuasive-sounding language to real opinions.

Consequences: No CJA fees for AI-related work; disciplinary referrals; replacement counsel appointed.

Much harder to catch — requires verifying the proposition and the pin cite.

The Verification Standard

Verify the proposition **and** the pin cite — not just the case name.



THE EXPERTISE ADVANTAGE

Experienced lawyers have pattern recognition

They can "feel" when something is off because years of practice have trained the instinct — a resource junior lawyers lack.



THE TRAINING CONCERN

AI-first training can stunt the instinct

Junior lawyers trained primarily on AI output may never develop that instinct. This is not just an efficiency question — it is a firm-level training obligation under Rules 5.1 and 5.3.

Workflows That Force Quality Control

Build quality control into the process — not into the person. Discipline fails; process does not.

01

Defined Checkpoints

Specify exactly when a human verifies — e.g., before filing, before sending to opposing counsel, before a final signature version.

02

Feedback Loops

Log errors, track patterns, iterate prompts. Every mistake should improve the next cycle, not just get corrected in place.

03

Deliberate Automation Decisions

Decide in advance what AI handles vs. what stays manual. Do not let the default be "whatever the junior associate reached for."

04

Training-Data & Bias Awareness

Understand what the model was trained on. Know where its blind spots are — especially for jurisdiction-specific or niche practice areas.

Trust is Earned, Not Assumed

At what point can lawyers begin to trust AI outputs without line-by-line review?

01

STAGE 1

Issue spotting

Lawyer reviews 100%.



02

STAGE 2

Attorney-led

Lawyer selects and edits all output.



03

STAGE 3

AI first draft

Reduced review — with documentation.

Trust is the endpoint of a documented process. Without recorded accuracy rates at Stages 1 and 2, there is no basis for reduced review at Stage 3.

PART 3

03

Confidentiality

What everyone worries about — but usually hand-waves.

Rule 1.6(e) and The Reasonableness Test

RULE 1.6(E)

“ A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client. ”

COMMENT [18] · REASONABLENESS FACTORS

A non-exhaustive list for evaluating whether safeguards are "reasonable":

- | | |
|---|--|
| 01 Sensitivity of the information | 02 Likelihood of disclosure absent the safeguard |
| 03 Cost of employing additional safeguards | 04 Difficulty of implementing the safeguards |
| 05 Extent to which safeguards adversely affect the representation | |

Heppner: Three Privilege Failures

Pro se litigant used consumer-grade Claude to help draft filings. Judge Rakoff (S.D.N.Y. 2025) denied privilege on three independent grounds — each of which is fixable.

01

No attorney-client relationship

Claude is not an attorney. There is no privileged relationship to invoke in the first place.

02

No expectation of confidentiality

Consumer terms of service permit data access for safety reviews, trust and safety work, and potential training.

03

Not at counsel's direction

Use was by the pro se litigant directly — not by a licensed attorney supervising the workflow.

Heppner is not a cautionary tale about AI. It is a cautionary tale about *consumer-grade AI used outside any attorney supervision*.

Preserving Privilege: Rakoff's Roadmap

The Heppner dicta is the clearest judicial signal to date on how firms should structure AI adoption to preserve privilege. Four conditions — all of which a well-run firm can meet.

01

Firm-licensed tool

Enterprise platform with a contract — not a personal consumer account.

02

Lawyer directing the use

Work performed under the supervision of a licensed attorney, not around one.

03

Purpose is legal advice

Tool is used to render legal advice to a client — not general research or personal use.

04

Access controls + audit trails

Logged, auditable use that can be reconstructed if ever challenged.

The Kovel Bridge

Enterprise AI as an agent of the attorney. Under *Kovel*, privilege has long extended to third parties — accountants, translators, litigation support vendors — acting at the direction of counsel. Enterprise AI, properly structured, fits the same framework.

THE ANALOGY

Accountant

Translates financial records into a form counsel can use to render legal advice.



Translator

Converts a client's statements into a language counsel can work with.



Enterprise AI

Assists counsel in producing, reviewing, or organizing work product — at counsel's direction.



Three-Tier Confidentiality Framework

Not all AI tools carry the same confidentiality profile. Match the tier to the sensitivity of the work.

TIER	TOOL TYPE	CONFIDENTIALITY PROFILE
FREE / PUBLIC	ChatGPT Free, Claude Free, Gemini Free	No confidentiality protections. Data may be used for training. Never use with client information.
PAID / CONSUMER	ChatGPT Plus, Claude Pro, CoCounsel	Improved terms, but still consumer-grade. Review ToS carefully. Data may still be retained or accessible.
ENTERPRISE	Firm-licensed platforms with BAAs, access controls, and audit trails	Strongest protections. Contractual commitments on data handling. Required for client-matter work.

PART 4

04

Risk Allocation

Where liability sits today — and how to address it in the engagement.

Where Liability Sits Today

Liability sits with the lawyer. **Period.**

Tool vendors' terms universally disclaim liability for output accuracy. Sanctions are not for using AI — they are for failing to verify.

MATA · SANCTIONS

Failing to verify — not using AI

\$5,000 per attorney

Required notice to every judge in the fabricated citations

Sanctions were for failing to verify, not for using AI

FARRIS · CONSEQUENCES

Legal-specific tool — didn't matter

No CJA fees awarded for AI-related work

Disciplinary referrals to state bars

Replacement counsel appointed

CoCounsel is legal-specific — still yielded sanctions

Engagement Letter Considerations

An explicit conversation about AI use belongs at the outset of every engagement — both the firm's use and the client's use. The engagement letter is where that conversation is memorialized.

01

Scope of AI use

Define what "AI-assisted" means in this engagement — research, drafting, review, summarization.

02

Client opt-out mechanism

A clear path for the client to decline AI use for their matter, with corresponding fee and timing impacts.

03

Firm's use and client's use

Address both directions. A client feeding drafts into free-tier ChatGPT can destroy privilege over firm work product.

04

Indemnity provisions

Who bears the risk for AI-introduced errors? Who controls the tools used? Align the allocation with who can manage the risk.

The Client-Side Problem

THE SCENARIO

A client runs your draft through free-tier ChatGPT

...on their own, to "check" the language, shorten it, or translate it.

HOW TO ADDRESS IT

Raise AI use directly at the engagement kickoff — do not assume the client knows the risk.

Memorialize expectations about client-side AI use in the engagement letter.

Offer enterprise-tier alternatives where appropriate — route clients who want AI assistance into a supervised tool, not to a consumer chatbot.

Re-raise the topic when new personnel on the client side enter the matter.

What Firms Need

No new ethical rules are required. What is required is documented institutional discipline.

01

A documented AI policy

Approved tools, permitted use cases, review standards, and training requirements — in writing.

02

A six-month review cycle

The technology and vendor landscape moves too fast for annual review. Revisit every six months.

03

Training under Rules 5.1 & 5.3

Supervisory lawyers have an affirmative obligation to ensure AI conduct is compatible with professional obligations.

04

Insurance posture

Expect malpractice carriers to require documented AI policies as a condition of coverage or favorable premiums.

PART 5

05

Strategic & Future-State

What this means for training, judgment, and competitive positioning.

Training and Pattern Recognition

THE EXPERIENCED-LAWYER ADVANTAGE

Pattern recognition is a career-built asset

Experienced lawyers can sense when AI output is off — a provision reads strangely, a citation feels too convenient, a definition doesn't match the deal. Years of reading contracts have trained an instinct that AI cannot replace.



THE JUNIOR-LAWYER TRAINING CONCERN

AI use at the start of a career can stunt instinct

Junior lawyers trained primarily on AI output may never develop the reflex that catches Farris-type errors — real citations with fabricated quotes. Firms have an affirmative training obligation under Rules 5.1 and 5.3 to build that reflex deliberately.



Judgment Matters More, Not Less

AI is designed to be persuasive and confident — even when wrong. That makes human judgment more valuable. The lawyers who thrive will be those who can distinguish confident-sounding output from correct output.

Clarity of thought wins

If you can communicate clearly as a human being, you can use these tools effectively. The skill is not technical — it is clear thinking.

Commoditization is inverted

AI lowers the floor on mechanical drafting. That does not commoditize judgment — it makes the judgment layer more visible and more valuable.

Confidence is not correctness

Tools will keep getting better at being wrong in convincing ways. The lawyer's job is to hold the line between plausible and accurate.

Best-In-Class In Three To Five Years

The firms and legal teams that compound their AI advantage over the next few years will share a recognizable set of traits.

01

Staged playbooks with accuracy rates

Issue spotting → attorney-led drafting → AI first draft, with documented accuracy rates at each stage.

02

Compounding client-service advantages

Faster turnaround, more consistent work product, better internal handoffs — all reinforcing one another.

03

Integrated feedback loops

Every AI error improves the next prompt cycle. Institutional learning, not individual learning.

04

A hospitality mindset

Freed-up time is reinvested in strategic advice and client relationships — not absorbed into higher hourly volume.

Three Takeaways

1

The ethical rules haven't changed.

Rule 1.1 (competence), Rule 1.6(e) (confidentiality), and Rule 5.3 (supervision) already govern everything we discussed. No new rules needed — just disciplined application of existing ones.



2

Verify the output. Read the terms. Direct the use.

Every case we discussed — Mata, Farris, and Heppner — failed on at least one of these. Every one of those failures was avoidable.



3

Quality control is where the upside lives.

The firms and lawyers who build verification into their AI workflows — not as an afterthought, but as the centerpiece — will deliver better work, faster, and more consistently.



PANELIST

Benjamin M. Altshul



Partner

Levenfeld Pearlstein, LLC

REAL ESTATE · AI IN
LEGAL PRACTICE

Benjamin Altshul is a partner in the Real Estate Group at Levenfeld Pearlstein, where he represents sophisticated real estate investors, developers, and corporate clients in complex commercial real estate transactions. His practice focuses on acquisitions and dispositions, financing, leasing, and the negotiation of risk allocation in real estate contracts, with particular experience in multifamily and shopping center asset classes.

Benjamin is an active adopter of generative AI in legal practice, building AI-integrated drafting and knowledge management systems for transactional work and regularly educating colleagues on practical applications of generative AI in the legal profession.

He serves as President of the Board of the Jewish Council for Youth Services and has been recognized as an Illinois Super Lawyers Rising Star and in The Best Lawyers in America: Ones to Watch.

PANELIST

Michael J. Riesen



Partner

Smith, Gambrell & Russell,
LLP

INTELLECTUAL PROPERTY ·
EXECUTIVE DIRECTOR, SGR
LABS

Michael Riesen is a Partner and Atlanta Section Head of the Intellectual Property practice at Smith, Gambrell & Russell, and serves as Executive Director of SGR Labs, the firm's initiative developing technology and AI solutions for legal work.

A U.S. patent attorney and U.K. Solicitor, Michael counsels Fortune 500 companies, emerging-growth ventures, and universities on IP strategy, technology transactions, and AI-related regulatory and commercial issues. His practice includes negotiation and drafting of complex technology, SaaS, and AI agreements, development of corporate AI policies, and AI audit programs.

Michael has been recognized in The Best Lawyers in America for IP Litigation and the IAM Patent 1000, and regularly speaks on AI in legal practice at universities and public and private institutions.

PANELIST

Dev Thain



**Head of Academic
Partnerships**

Spellbook

LEGAL AI STRATEGY ·
ACADEMIC & BAR
PARTNERSHIPS

Dev Thain is the Head of Academic Partnerships at Spellbook, where he leads partnerships with law schools, bar associations, and legal institutions to drive adoption of AI-powered legal tools.

He previously spent six years at Gartner advising enterprises and governments on transformation, technology strategy, and AI adoption. Dev holds a Master's in Learning Technology from the University of Oxford.

He regularly speaks on practical applications of AI in legal practice, most recently at Harvard Law School and the American Bar Association.

APPENDIX A

Illinois Rules of Professional Conduct

RULE 1.1 · COMPETENCE

Competent representation

A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.

Comment [8]: To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements.

RULE 1.6(E) · CONFIDENTIALITY

Safeguarding information

A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.

RULE 5.3 · NONLAWYER ASSISTANTS

Supervisory responsibilities

With respect to a nonlawyer employed or retained by or associated with a lawyer: **(a)** a partner, and a lawyer who individually or together with other lawyers possesses comparable managerial authority in a law firm, shall make reasonable efforts to ensure that the firm has in effect measures giving reasonable assurance that the person's conduct is compatible with the professional obligations of the lawyer; **(b)** a lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer.

APPENDIX B

Cases Discussed

RULE 1.1 · COMPETENCE

Mata v. Avianca, Inc. — No. 22-cv-01461 (S.D.N.Y. June 22, 2023)

Plaintiff's counsel used ChatGPT for federal bankruptcy-tolling research. The AI fabricated six case citations — complete with docket numbers and quotations — that did not exist. Counsel filed the brief without verification. Sanctions: \$5,000 per attorney; required notification to every judge referenced in the fabricated citations.

2025 · REAL CASES, FABRICATED QUOTES

In re Farris — No. 2:22-cr-00082 (C.D. Cal. Jan. 2025)

Federal public defender used CoCounsel, a legal-specific AI tool. Cited cases were real, but quoted passages attributed to those cases were fabricated. Consequences: CJA fees denied for AI-related work; disciplinary referrals; replacement counsel appointed.

2025 · PRIVILEGE & CONSUMER AI

Heppner v. Wal-Mart Stores, Inc. — No. 24-cv-09286 (S.D.N.Y. Mar. 2025)

Pro se litigant used consumer-grade Claude to draft filings. Judge Rakoff denied privilege, finding: (1) Claude is not an attorney; (2) no reasonable expectation of confidentiality under Anthropic's consumer terms; and (3) use was not at counsel's direction. Dicta outlined a path for preserving privilege: firm-licensed tool, lawyer direction, legal-advice purpose, and access controls with audit trails.

ATTENDEE MATERIALS

THE CHICAGO BAR ASSOCIATION PRESENTS:

AI 2035: The Legal Profession and the Law in the Age of Artificial Intelligence

Drafting with AI

Quality Control, Confidentiality, and Contract Risk Allocation

May 11, 2026 · Venue SIX10 · Chicago

Panel Benjamin Altshul · Michael Riesen · Dev Thain



THE
**CHICAGO
BAR
ASSOCIATION**

PANELISTS



Benjamin M. Altshul

Partner

Levenfeld Pearlstein, LLC

REAL ESTATE · AI IN LEGAL PRACTICE



Michael J. Riesen

Partner

Smith, Gambrell & Russell, LLP

INTELLECTUAL PROPERTY · SGR LABS



Dev Thain

Head of Academic Partnerships

Spellbook

LEGAL AI STRATEGY · ACADEMIC
PARTNERSHIPS

PANEL OVERVIEW

What We'll Cover

How lawyers can use AI drafting tools responsibly — protecting clients, preserving privilege, and allocating risk. Organized around three lenses that follow the drafting workflow itself.



INPUT

01

Quality Control

What goes into the draft — and how you verify what comes out of it.



PROTECTION

02

Confidentiality

How you safeguard client data throughout the workflow.



OUTPUT

03

Risk Allocation

Who bears the consequences when something goes wrong.

Lawyers are not getting in trouble because they are using AI.

They are getting in trouble because they are using AI without understanding how these tools actually work — and without bothering to review the output.

Every point in this presentation ties back to two existing obligations: the duty of competence in technology, and the duty to supervise nonlawyer assistants.

ORGANIZING FRAMEWORK

Input → Protection → Output

Not separate topics. Sequential stages of the same workflow — what you put in, how you protect it in transit, and who bears the consequences of what comes out.



01

INPUT · QUALITY CONTROL

What goes in

Verification workflows
Failure-mode awareness



02

PROTECTION · CONFIDENTIALITY

How it's protected

Vendor evaluation
Privilege preservation
Tool selection



03

OUTPUT · RISK ALLOCATION

Who is liable

Engagement letters
Insurance posture

What's AI-ready vs. Still Too Risky

AI-READY

High-volume, repeatable contracts

- Standard NDAs
- Standard commercial leases
- Vendor and service agreements
- Playbook-driven first-pass drafts of familiar clauses
- Issue-spotting against a known playbook



TOO RISKY FOR UNSUPERVISED AI

Bespoke, novel, or high-consequence provisions

- Bespoke M&A provisions
- Novel indemnity structures
- Jurisdiction-specific regulatory language
- Limitation of liability and risk-allocation clauses
- Anything requiring business-context judgment



Failure Modes: Catchable vs. Convincing

Two decisions tell the story of how AI failure modes have evolved. The first is easy to catch. The second is much harder.

CASE 1 · 2023

FABRICATED CASES

Mata v. Avianca

S.D.N.Y.

A lawyer used ChatGPT for federal bankruptcy-tolling research. The AI fabricated six case citations — complete with docket numbers and quotes — that did not exist. The lawyer never verified. The brief was filed. Opposing counsel flagged the fabrications.

Sanctions: \$5,000 per attorney; required notice to every judge in the fabricated citations.

Relatively easy to catch — the cases themselves did not exist.

CASE 2 · 2025

REAL CASES, FABRICATED QUOTES

In re Farris

C.D. CAL.

A federal public defender used CoCounsel — a legal-specific AI tool. The cases cited were real. But the quoted passages attributed to those cases were fabricated. The AI pinned persuasive-sounding language to real opinions.

Consequences: No CJA fees for AI-related work; disciplinary referrals; replacement counsel appointed.

Much harder to catch — requires verifying the proposition and the pin cite.

Three-Tier Confidentiality Framework

Not all AI tools carry the same confidentiality profile. Match the tier to the sensitivity of the work.

TIER	TOOL TYPE	CONFIDENTIALITY PROFILE
FREE / PUBLIC	ChatGPT Free, Claude Free, Gemini Free	No confidentiality protections. Data may be used for training. Never use with client information.
PAID / CONSUMER	ChatGPT Plus, Claude Pro, CoCounsel	Improved terms, but still consumer-grade. Review ToS carefully. Data may still be retained or accessible.
ENTERPRISE	Firm-licensed platforms with BAAs, access controls, and audit trails	Strongest protections. Contractual commitments on data handling. Required for client-matter work.

Engagement Letter Considerations

An explicit conversation about AI use belongs at the outset of every engagement — both the firm's use and the client's use. The engagement letter is where that conversation is memorialized.

01

Scope of AI use

Define what "AI-assisted" means in this engagement — research, drafting, review, summarization.

02

Client opt-out mechanism

A clear path for the client to decline AI use for their matter, with corresponding fee and timing impacts.

03

Firm's use and client's use

Address both directions. A client feeding drafts into free-tier ChatGPT can destroy privilege over firm work product.

04

Indemnity provisions

Who bears the risk for AI-introduced errors? Who controls the tools used? Align the allocation with who can manage the risk.

The Client-Side Problem

THE SCENARIO

A client runs your draft through free-tier ChatGPT

...on their own, to "check" the language, shorten it, or translate it.

HOW TO ADDRESS IT

Raise AI use directly at the engagement kickoff — do not assume the client knows the risk.

Memorialize expectations about client-side AI use in the engagement letter.

Offer enterprise-tier alternatives where appropriate — route clients who want AI assistance into a supervised tool, not to a consumer chatbot.

Re-raise the topic when new personnel on the client side enter the matter.

Three Takeaways

1

The ethical rules haven't changed.

Rule 1.1 (competence), Rule 1.6(e) (confidentiality), and Rule 5.3 (supervision) already govern everything we discussed. No new rules needed — just disciplined application of existing ones.



2

Verify the output. Read the terms. Direct the use.

Every case we discussed — Mata, Farris, and Heppner — failed on at least one of these. Every one of those failures was avoidable.



3

Quality control is where the upside lives.

The firms and lawyers who build verification into their AI workflows — not as an afterthought, but as the centerpiece — will deliver better work, faster, and more consistently.

