

AI Agents for Lawyers

A Practical Guide to Agentic AI, Workflow Delegation, and the
Governance Legal Teams Need

Colin S. Levy
2026

About the Author

Colin S. Levy

Colin S. Levy is a legal technology advocate, writer, and advisor who works at the intersection of law, technology, and business. With experience spanning in-house legal roles, legal technology companies, and legal operations, he brings a practical perspective to how legal teams can adopt and govern emerging technologies responsibly.

Colin writes and speaks extensively on legal innovation, artificial intelligence in legal practice, and the evolving role of legal professionals in a technology-driven landscape. His work focuses on helping legal teams move beyond the hype cycle to make sound, informed decisions about the tools they use and the workflows they build.

He is the author of *The Legal Tech Ecosystem* and editor of the *Handbook of Legal Tech*, and a regular contributor to publications covering legal technology and operations. He advises organizations on responsible AI adoption, legal operations strategy, and the practical governance frameworks that make innovation sustainable.

How to Use This Guide

This guide is designed to be read cover to cover by a lawyer who is new to AI agents and to be used as a reference by lawyers who already are. It is short on purpose. The goal is to give you a working mental model and a set of artifacts you can adapt, not an exhaustive treatise. Depending on your role, different sections will matter most.

If you are new to AI agents. Start with Part One for the concepts and vocabulary, then read Part Two for use cases that will sound familiar from your own practice.

If you supervise other lawyers. Begin with Part Four on Professional Responsibility, then move to Part Five for the governance framework you will be expected to implement.

If you are evaluating or buying a tool. Skim Parts One through Three for context, then go directly to Part Seven for the vendor questionnaire and pilot readiness checklist.

If you are drafting a firm policy. Read Parts Four and Five together, then use the sample policy outline in Part Seven as a starting skeleton.

A Note on Scope

This guide is U.S.-centric and grounded primarily in the ABA Model Rules of Professional Conduct. State adoptions vary, and lawyers practicing internationally face additional regimes including the EU AI Act and the GDPR. Part Four flags both. Always consult your jurisdiction's specific rules and your own counsel before relying on any framework here.

Contents

Part One: From Generative AI to Agentic AI

- What Changes With Agents
- Key Terms at a Glance
- How Agents Actually Work
- Why Lawyers Should Care Now

Part Two: Use Cases for Legal Teams

- Contract Lifecycle Workflows
- Legal Research and Memo Drafting
- Intake, Triage, and Routing
- Compliance Monitoring
- Litigation and Discovery Support

Part Three: New Risks Agents Introduce

- Compounding Errors Across Steps
- Tool Misuse and Unauthorized Actions
- Confidentiality and Data Movement
- Auditability and the Paper Trail

Part Four: Professional Responsibility

- Supervision Under Rules 5.1 and 5.3
- Billing, Fees, and Rule 1.5
- Malpractice Exposure and Insurance
- State Variation in the United States
- Cross-Border Considerations

Part Five: A Governance Framework

- Scoping Authority and Boundaries
- Human-in-the-Loop Checkpoints
- Logging, Review, and Rollback
- Vendor Diligence for Agentic Tools

Part Six: Getting Started

- A Pilot Worth Running
- What to Measure
- A Final Word on Judgment

Part Seven: A Practical Toolkit

- Sample AI Agent Use Policy (Outline)
- Vendor Questionnaire for Agentic Tools
- Pilot Readiness Checklist

Human Checkpoint Decision Aid

Notes

Part One

From Generative AI to Agentic AI

What Changes With Agents

Generative AI answers questions. Agentic AI takes action. That single shift, from a system that produces text on request to a system that plans steps, calls other tools, and executes tasks across connected systems, is the most consequential change in legal technology since the move to cloud document management.

An agent does not just draft a clause when asked. It can read the inbound contract, compare it against your playbook, pull the counterparty's prior agreements from your CLM, draft redlines, route the file for review, and update the matter record, all in a single run. The lawyer's role moves from operator to supervisor.

Key Terms at a Glance

A short reference for the terms used throughout this guide. Refer back as needed.

Term	What It Means
AI Agent	An AI system that can plan multi-step tasks, use external tools, and take actions across connected systems to achieve a defined goal.
Tool Use	An agent's ability to call APIs, search databases, send messages, or operate other software as part of completing a task.
Orchestration	The layer that decides which model, tool, or sub-agent runs at each step and in what order.
Human in the Loop (HITL)	A checkpoint where a person must review or approve before the agent proceeds. The default posture for legal work.
Autonomy Level	How much an agent can do without human approval. Ranges from suggest-only to fully autonomous execution.
Guardrails	Hard limits on what an agent can access or do, enforced by configuration rather than by the model's own judgment.
Audit Log	A complete, time-stamped record of what an agent observed, decided, and did. Essential for review and defensibility.
Multi-Agent System	A setup in which several specialized agents collaborate, often with one orchestrator delegating to others.

How Agents Actually Work

Under the hood, an agent is a language model wrapped in a loop. Given a goal, the model decides what to do next, calls a tool, observes the result, and decides again. It continues until the goal is met or a stopping condition is reached.

Three things distinguish an agent from a chatbot. First, it has access to tools, such as a document store, a contract platform, a calendar, or an email system. Second, it can take multiple steps without being prompted between each one. Third, its actions have effects in the real world. A chatbot that hallucinates produces a wrong answer. An agent that hallucinates can send the wrong email, file the wrong document, or sign the wrong version.

The Core Distinction

Generative AI produces output. Agentic AI produces outcomes. The risk profile changes accordingly, and so should the controls around it.

Why Lawyers Should Care Now

Agentic features are already shipping inside the platforms legal teams use every day. Contract lifecycle management vendors, research providers, and e-discovery platforms have moved well past simple prompt boxes. The 2025 ACC and Everlaw survey found that generative AI use in corporate legal departments more than doubled in a single year, jumping from 23 percent in 2024 to 52 percent in 2025.¹ Agentic capabilities are the next wave layered on top of that adoption, and they are arriving whether or not legal is at the table.

The right question is not whether to allow agents but how to scope, supervise, and document their use so that the work product remains defensible and the client remains protected.



Part Two

Use Cases for Legal Teams

Contract Lifecycle Workflows

Contracts are the most natural starting point for agentic workflows because the steps are repeatable, the documents are structured, and the systems involved expose APIs. A contract agent can intake a third-party paper, classify it, compare it to your playbook, draft redlines, surface deviations for human review, and route the file through approval. Throughout, the lawyer's role shifts from cycling through clauses to reviewing flagged issues and confirming positions.

Legal Research and Memo Drafting

Research agents can decompose a question, run multiple searches across primary and secondary sources, retrieve and read relevant authorities, and assemble a draft memo with citations. Used well, this collapses days of associate time into hours. Used carelessly, it produces the same fabricated citations that have already led to sanctions in courts across the country, beginning with *Mata v. Avianca, Inc.* (S.D.N.Y. 2023) and continuing through a growing line of cases involving both general-purpose and purpose-built legal AI tools.² A 2024 Stanford RegLab study found that leading legal-specific AI research tools, including Lexis+ AI and Thomson Reuters' Westlaw AI-Assisted Research, produced incorrect or misgrounded responses on between roughly 17 and 33 percent of queries tested.³ Verification is non-negotiable.

Intake, Triage, and Routing

Many in-house legal teams drown in inbound requests that do not require legal expertise to categorize. An intake agent can ask clarifying questions, classify the request, attach self-serve resources where appropriate, and route true legal issues to the right attorney with context already gathered. The downstream effect is that lawyers spend their time on the matters that actually need them.

Compliance Monitoring

Compliance agents can watch for regulatory changes in defined jurisdictions, map them against an organization's policies and obligations, and flag what may need attention. The value is in the surfacing, not the conclusion. A flagged update is the start of a human review, not the end of one.

Litigation and Discovery Support

In discovery, agents can coordinate review across large document populations, propose responsive tags, draft privilege descriptions, and assemble chronologies. The traditional review pyramid compresses, but the need for skilled supervision grows. Errors in this context have direct consequences for clients and for the lawyers responsible.

Part Three

New Risks Agents Introduce

Compounding Errors Across Steps

An agent that is 95 percent reliable on a single step is far less reliable across ten steps. The math is unforgiving. Each step depends on the prior step being correct, and small errors propagate and amplify. A misread of a contract type at step one can cascade into the wrong playbook, the wrong redlines, and the wrong approval path.

The Multiplication Problem

Reliability across a chain is the product of reliability at each link, not the average. Designing an agent workflow means designing where the human checkpoints sit so that errors cannot quietly compound.

Worked Example: A Contract Intake Goes Sideways

Consider an intake agent that receives a vendor agreement, classifies the contract type, selects the matching playbook, drafts redlines, and routes the file for approval. At step one, the agent reads the document title "Master Services Agreement" and classifies it as a standard MSA, missing a clause on the third page that converts it into a data processing agreement. At step two, the agent loads the MSA playbook, which contains no DPA-specific positions on cross-border transfers, sub-processor approval, or breach notification. At step three, the agent generates redlines that look complete because they cover every clause in the MSA playbook, but they say nothing about the DPA terms because those terms were never on the checklist. At step four, the agent routes the file to the commercial attorney rather than the privacy attorney, because routing depends on classification. The commercial attorney, trusting the prior steps, reviews the redlines without re-reading the underlying contract. The deal closes. Six months later, a data incident raises a question no one can answer cleanly. None of the four steps was unreasonable in isolation. The compounding is what produced the harm.

Tool Misuse and Unauthorized Actions

When an agent has the ability to send email, edit documents, file with courts, or charge a credit card, the question of what it is allowed to do becomes a question about authority. The relevant controls are not just technical. They include who decided the agent could take this action, what scope of authority it was granted, and how that grant is documented. These are familiar questions for lawyers in other contexts, and they apply here directly.

Confidentiality and Data Movement

Agents call tools, and tools live somewhere. Each call may transmit privileged or confidential information to a new system. Data flow diagrams that were accurate last quarter may no longer reflect what an agentic workflow actually does. Data processing agreements, subprocessor lists, and conflict checks all need to keep pace.

Auditability and the Paper Trail

If something goes wrong, can you reconstruct exactly what the agent did, when, and why? Logging that is sufficient for debugging is not always sufficient for legal review. The right standard is whether a partner, a court, or a regulator could follow the trail without needing the vendor to translate.

Part Four

Professional Responsibility

Supervision Under Rules 5.1 and 5.3

ABA Formal Opinion 512, issued in July 2024, is the authoritative starting point for any lawyer's analysis of generative AI use, and its reasoning extends naturally to agentic tools.⁴ The opinion identifies six rule areas implicated by AI use: competence (Rule 1.1), confidentiality (Rule 1.6), communication with clients (Rule 1.4), candor toward the tribunal (Rules 3.1 and 3.3), supervision of subordinates and non-lawyer assistance (Rules 5.1 and 5.3), and reasonable fees (Rule 1.5).

Rules 5.1 and 5.3 are the rules that hit hardest in the agentic context. They require lawyers with managerial authority to make reasonable efforts to ensure that the conduct of subordinate lawyers and non-lawyer assistants conforms to the rules of professional conduct. Bar guidance has increasingly treated AI tools as a form of non-lawyer assistance, which means the supervisory framework applies. An agent that drafts, files, or transmits work product is, for purposes of supervisory analysis, analogous to a paralegal or vendor whose work the supervising lawyer remains accountable for.

What This Means in Practice

If you would not let a first-week paralegal send the email, file the brief, or sign the document without review, the agent should not do it either. The supervisory standard does not soften because the assistant is software.

Billing, Fees, and Rule 1.5

Model Rule 1.5 prohibits unreasonable fees. Formal Opinion 512 makes clear that lawyers may not bill clients for time the lawyer did not actually spend, even when AI tools produce work product that would have taken hours to draft by hand.⁵ The principle is straightforward: efficiency gains from agentic workflows belong, at least in part, to the client unless the engagement letter says otherwise.

The opinion also addresses cost recovery. Out-of-pocket AI charges may be passed through to clients only if the engagement agreement permits it and the charges are reasonable. Overhead-style allocations of subscription costs are generally not recoverable absent specific disclosure. Lawyers using agentic tools should review their engagement letters and billing guidelines now, before the first client raises the question, and consider whether alternative fee arrangements better reflect the way the work is actually being done.

Worked Example: A Time Entry Before and After

Before agents, an associate spends 4.5 hours pulling counterparty precedent, comparing against the playbook, and drafting redlines for an inbound MSA. The bill reads: "Review and redline MSA against playbook, including precedent comparison, 4.5 hours." After agents, the same workflow takes the associate 45 minutes of supervised review on top of an agent-prepared draft. Billing the same 4.5 hours is not permitted under Rule 1.5. The defensible approach is to bill the actual time, note the use of an AI tool if the engagement letter requires it, and let the efficiency flow through. Some firms are shifting these workflows to fixed-fee or capped arrangements precisely so that the value conversation does not happen one time entry at a time.

Malpractice Exposure and Insurance

Professional liability insurers have begun asking about AI use on renewal applications, and the questions are getting more specific. Carriers want to know which tools are in use, what the firm's policies require, what training has been provided, and how output is verified. A complete and truthful answer is not just a matter of underwriting cooperation. Misrepresentations on a renewal application can affect coverage when a claim arises.

The malpractice exposure from agentic tools is best understood as an amplification of existing exposures rather than a new category. The same errors that lead to malpractice claims today, missed deadlines, incorrect advice, unauthorized actions, breach of confidentiality, can all be produced at greater speed and scale by an agent operating without adequate supervision. The mitigation is the same as it has always been: clear scope, competent supervision, contemporaneous documentation, and a willingness to slow down when something does not look right.

State Variation in the United States

ABA Formal Opinion 512 is persuasive authority, not binding rule text. Each state sets its own rules of professional conduct and issues its own ethics guidance. Several states have moved early and the positions diverge in ways that matter to practitioners. California's State Bar issued practical guidance on generative AI in November 2023 emphasizing confidentiality and supervision.⁷ The Florida Bar followed in January 2024 with a formal opinion addressing client consent, billing, and oversight obligations.⁸ New York, New Jersey, the District of Columbia, and several others have issued their own opinions or working group reports. The differences are not always cosmetic. Some jurisdictions require affirmative client disclosure for material AI use; others do not. Some address fee arrangements explicitly; others do not. Always check the rule text and most recent guidance in every jurisdiction where you practice before relying on a national framework.

Cross-Border Considerations

Lawyers and legal teams operating in or serving clients in the European Union face additional layers of regulation that the Model Rules do not reach. The EU AI Act, which entered into force in 2024 with phased application through 2026 and beyond, classifies certain AI systems used in administration of justice and democratic processes as high-risk and imposes documentation, transparency, and human-oversight obligations on providers and deployers.⁹ The General Data Protection Regulation continues to govern personal data processed by agentic tools, including the Article 22 limits on decisions based solely on automated processing.¹⁰ The United Kingdom, Canada, Singapore, and others are pursuing their own approaches, with regulators increasingly focused on agentic systems specifically. A workflow that is compliant under U.S. ethics rules may still create exposure under foreign privacy or AI-specific law when the data, the user, or the affected individual is outside the United States. Cross-border matters warrant a separate review by counsel familiar with the relevant regimes.

Part Five

A Governance Framework

Scoping Authority and Boundaries

Treat an agent's deployment the way you would treat the onboarding of a new junior professional. Define what it can access, what it can do, what it must escalate, and what is outside its scope entirely. Write it down. Review it on a schedule. The clearer the scope, the easier the supervision.

Human-in-the-Loop Checkpoints

Not every step needs a human review, but the steps with irreversible or high-stakes effects always do. A useful framing is to ask, for each action the agent might take, whether the consequence can be undone in five minutes by the same person who would have caught the mistake. If not, it needs a checkpoint.

Autonomy by Risk

Level	Description	Appropriate Use
Suggest	Agent proposes; human executes every action.	New tools, sensitive matters, novel workflows.
Approve	Agent prepares actions; human approves each batch before execution.	Routine but consequential work such as contract redlines.
Notify	Agent acts; human is notified and can intervene after the fact.	Low-stakes, easily reversible actions such as triage tagging.
Autonomous	Agent acts without human notification.	Reserve for narrow, well-tested, low-risk steps with strong logging.

Logging, Review, and Rollback

Three capabilities determine whether you can defend an agentic workflow. First, complete logging of inputs, decisions, tool calls, and outputs. Second, regular human sampling of completed runs to detect drift and surface issues that did not trip an automatic flag. Third, a way to undo or correct what the agent has done when something goes wrong. If a workflow lacks any of these, slow it down until it has them.

Vendor Diligence for Agentic Tools

Standard SaaS diligence is necessary but not sufficient. For agentic tools, ask specifically about which actions the agent can take in your environment, how its behavior is tested and monitored, what telemetry is retained, whether your data is used to improve the model, and how incidents are surfaced. Get the answers in writing and tie them to the contract.

Part Six

Getting Started

A Pilot Worth Running

Pick a workflow that is repetitive, well-bounded, and high-volume but low-stakes. NDA triage is the canonical example because the inputs are structured, the playbook is stable, and the consequences of error are recoverable. Run it in approve mode for thirty days. Measure how often the agent's recommendation matches the human's. Use the gap as a learning tool, not a verdict.

What to Measure

Adoption metrics tell you whether people are using the tool. They do not tell you whether it is working. Track outcomes that matter: agreement rate between the agent and the human, cycle time from intake to disposition, escalation frequency, and the number of issues caught in human review that the agent missed. Watch the trend, not the snapshot.

A Final Word on Judgment

Agents are powerful because they remove friction. That is also what makes them dangerous. The friction in legal work is often where judgment lives, where a lawyer pauses, looks again, and asks whether the obvious answer is the right one. Designing agentic workflows well means preserving the moments that matter while clearing away the ones that do not. The goal is not to take lawyers out of the loop. It is to put them where they add the most value.

The Throughline

Generative AI asked lawyers to verify outputs. Agentic AI asks lawyers to govern systems. The shift in skill is from editor to supervisor, and the shift in posture is from reviewing drafts to designing controls.



Part Seven

A Practical Toolkit

The artifacts in this part are starting points, not finished products. Adapt them to your organization's structure, regulatory environment, and risk appetite. Each is meant to be short enough to actually use.

Sample AI Agent Use Policy (Outline)

- 1. Purpose and Scope.** State which agentic tools and workflows the policy covers, who must follow it, and how it relates to existing technology, confidentiality, and records policies.
- 2. Approved Tools.** Maintain a current list of agents approved for use, the data sources each may access, and the lawyer or administrator responsible for each. New tools require review before use.
- 3. Permitted and Prohibited Uses.** Identify categories of work where agents may be used (for example, contract intake, research drafts, document review tagging) and categories where they may not (for example, final advice to clients, court filings without independent review, matters involving specifically restricted data).
- 4. Confidentiality and Data Handling.** Specify what client and firm information may be entered into which tools, what must be redacted, and what must never be entered. Reference applicable data processing terms.
- 5. Human Review Requirements.** Define the default autonomy level for each approved workflow and the specific actions that require human approval before execution.
- 6. Verification and Citation.** Require independent verification of factual assertions, citations, and quoted material in any work product produced or assisted by an agent.
- 7. Client Disclosure.** Set the firm's position on when and how clients are informed of AI agent use, consistent with Model Rule 1.4 and applicable state guidance.
- 8. Billing.** State how time and costs related to agentic tools are recorded and billed, consistent with Model Rule 1.5 and engagement letter terms.
- 9. Logging and Records.** Require retention of agent activity logs sufficient to reconstruct what the agent did on any given matter, and identify the retention period.
- 10. Incident Reporting.** Provide a clear path for reporting suspected errors, hallucinations, unauthorized actions, or confidentiality concerns. Identify who reviews and who decides on remediation.
- 11. Training and Acknowledgment.** Require initial and periodic training for all users, with a written acknowledgment of the policy on file.
- 12. Review Cycle.** Set a fixed review interval (for example, every six months) and identify who owns the review.

Vendor Questionnaire for Agentic Tools

Use this list as a baseline for vendor diligence. Get answers in writing and tie material answers to the contract.

Capabilities and Authority

- What specific actions can the agent take in our environment, by default and at maximum?
- Which actions require human approval, and is that configurable per workflow?
- Can authority be scoped by user, matter type, data source, or client?

Reliability and Testing

- How is the model evaluated before changes ship to our tenant?
- What benchmarks or task-level metrics do you publish or share under NDA?
- What is your rollback procedure if a release introduces regressions?

Data Handling

- Where is our data stored and processed, and by which subprocessors?
- Is our content used to train, fine-tune, or evaluate any model? Can we opt out?
- How is privileged or confidential content segregated from other tenants?
- What happens to our data on termination, and on what timeline?

Logging and Auditability

- What does the audit log capture for each agent run (inputs, tool calls, outputs, decisions)?
- Can we export logs on demand in a structured format?
- How long are logs retained, and can we configure retention?

Security and Compliance

- Which security certifications do you hold (SOC 2 Type II, ISO 27001, others)?
- How are credentials and tool access tokens stored and rotated?
- What is your incident response process and notification timeline?

Professional Responsibility Fit

- How does your tool support our supervisory obligations under Model Rules 5.1 and 5.3?
- What disclosures or documentation do you provide to support client communication?
- Do you maintain documentation aligned to ABA Formal Opinion 512 considerations?

Pilot Readiness Checklist

Before launching any agentic pilot, confirm each of the following:

- Workflow is repetitive, well-bounded, and produces effects that can be reversed if wrong.
- A named lawyer owns the pilot and has authority to pause or terminate it.
- Approved tool is on the firm's vetted list and has executed contractual terms in place.
- Data sources the agent will touch have been mapped and approved.
- Default autonomy level is set to suggest or approve, not notify or autonomous.
- Audit logging is enabled and the export format has been tested.
- Verification protocol for any factual or citation output is documented.
- Affected lawyers and staff have completed training and signed the policy acknowledgment.
- Client disclosure approach is decided and consistent with engagement terms.
- Success and failure metrics are defined in writing before the pilot starts.
- Review cadence is on the calendar (weekly during pilot, monthly thereafter).
- Incident reporting channel is live and known to users.

Human Checkpoint Decision Aid

When deciding whether a step in an agentic workflow requires a human checkpoint, run through these four questions. If the answer to any one of them is yes, the step needs a human in the loop before execution.

Question	If Yes, Why It Matters
Is the action irreversible within five minutes by the same person?	The cost of error is permanent. A checkpoint preserves the option to correct course.
Does the action affect a client, court, regulator, or counterparty externally?	External actions create reputational, ethical, and contractual exposure that warrant review.
Does the step require legal judgment or interpretation of ambiguous facts?	Judgment is the lawyer's job and the value the client is paying for.
Would the firm be required to disclose the action under engagement, court, or regulatory rules?	If disclosure is required, the lawyer must know what was done and why.

Notes

References below are provided for reader follow-up and are accurate to the best of the author's knowledge as of publication. Citations should be verified independently before being relied on in client work.

1. Association of Corporate Counsel and Everlaw, 2025 Generative AI in Corporate Legal Departments survey, reporting that generative AI use among surveyed corporate legal departments rose from 23 percent in 2024 to 52 percent in 2025.
2. *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023) (sanctioning attorneys for submitting a brief containing fabricated case citations generated by ChatGPT). For subsequent appellate treatment of similar conduct see *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024) (referring attorney to grievance panel after she submitted a reply brief containing a non-existent case generated by ChatGPT). Numerous additional instances have since been reported across federal and state courts, and many federal judges have issued standing orders addressing AI use in filings.
3. Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, and Daniel E. Ho, *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, Stanford RegLab and Stanford HAI (May 2024), peer-reviewed and published in the *Journal of Empirical Legal Studies* (2025). The study found that leading legal AI research tools, including Lexis+ AI and Thomson Reuters' Westlaw AI-Assisted Research, produced incorrect or misgrounded responses on between roughly 17 and 33 percent of queries tested, even with retrieval-augmented generation.
4. ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512: Generative Artificial Intelligence Tools (July 29, 2024), addressing competence, confidentiality, communication, candor, supervision, and fees in the context of lawyer use of generative AI tools.
5. ABA Formal Opinion 512 at section discussing Model Rule 1.5, noting that lawyers may not bill clients for time not actually expended and addressing the treatment of AI-related costs.
6. Model Rules of Professional Conduct R. 1.1 (competence), 1.4 (communication), 1.5 (fees), 1.6 (confidentiality), 3.3 (candor toward the tribunal), 5.1 (responsibilities of partners, managers, and supervisory lawyers), and 5.3 (responsibilities regarding non-lawyer assistance). State adoptions vary; consult your jurisdiction's rules.
7. State Bar of California, Standing Committee on Professional Responsibility and Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* (issued November 2023), addressing confidentiality, competence, supervision, communication, and billing in the context of generative AI use.
8. The Florida Bar, Ethics Opinion 24-1 (issued January 2024), addressing lawyer responsibilities when using generative AI, including confidentiality obligations, oversight of AI-generated work, billing practices, and lawyer advertising considerations. Other states with published guidance include New York, New Jersey, Pennsylvania, Texas, and the District of Columbia, among others.

9. Regulation (EU) 2024/1689 of the European Parliament and of the Council on Artificial Intelligence (the EU AI Act), which entered into force August 1, 2024, with provisions applying on a phased timeline. Annex III identifies AI systems intended to be used by a judicial authority or on its behalf to assist in researching and interpreting facts and the law as high-risk systems subject to specific obligations.

10. Regulation (EU) 2016/679 (General Data Protection Regulation), Article 22, providing data subjects the right not to be subject to a decision based solely on automated processing that produces legal or similarly significant effects. Use of agentic tools to process personal data of EU residents implicates the broader GDPR framework, including lawful basis, purpose limitation, and data minimization principles.