

LADDER OF RISK

Match review rigor to legal risk

One of the most useful ways to implement AI governance is to match review rigor to the level of legal risk. This framework avoids both extremes and it asks the right question:

What level of human judgment does this task require?



WATCH FOR BIAS
It's critical to treat AI output as a hypothesis, not an answer. Automation and confirmation bias reinforce each other.



THE KEY QUESTION
The question is not whether to use AI but whether organizations can use it without surrendering judgment.

 RED Lawyer review required	AI can assist, but a lawyer must meaningfully review before reliance or release. Controls: Approval gates, primary-source verification, explicit sign-off		
 EXAMPLES • Court filings • Legal citations • Binding terms • Regulatory submissions	 ABA MODEL RULES 1.1 Competence 1.6 Confidentiality 3.3 Candor 5.1 Supervision 5.5 Unauthorized practice	 AGENTIC • Approval gates before sending or filing • Logging and audit trails • Exception routing to lawyer	 DOCUMENTATION • Full audit trail and sign-off recorded • Source verification logged • Reviewer name and date on file confirmed • Blackline comparison retained
 YELLOW Review strongly advised	Human review remains highly prudent where output can shape legal or business posture. Controls: Substantive review, fit-to-facts check, escalation of ambiguous outputs		
 EXAMPLES • Client communications • Policy drafts • Legal positions • Matter summaries	 ABA MODEL RULES 1.4 Communication 1.5 Fees 2.1 Adviser	 AGENTIC • Role-based access controls • Connector and data loss prevention (DLP) controls • Human review checkpoints	 DOCUMENTATION • Substantive review logged • Fit-to-facts check noted • Escalation decisions recorded
 GREEN Lower-risk assist recommended	Lawyer review might be optional, but verification still matters. Controls: Spot-checks, approved tools only, sensitive inputs prevented		
 EXAMPLES • Brainstorming • Research outlines • Nonsensitive summaries	 ABA MODEL RULES 5.3 Nonlawyer assist 8.4 Misconduct	 AGENTIC • Least-privilege permissions • Approved tools only • Quality assurance (QA) sampling	 DOCUMENTATION • Record spot-checked • Tool and prompt logged • QA sample results retained

▲ Human review intensifies as legal stakes rise. ▲

The bridge to agentic AI



GENERATIVE AI RISK
Produces the wrong answer (bad drafts, fake citations, incomplete reasoning)



AGENTIC AI RISK
Takes the wrong action (wrong file, bad routing, transmits sensitive data)

Embedded controls for both: Approval gates • Role-based access • Audit trails • Exception routing • DLP • Human-review checkpoints



The lawyer-in-the-loop is not replaced. In fact, the risks are magnified as AI gains the ability to act, not just draft.

