FIA Law & Compliance Conference 2026

Panel: AI Applied for Compliance & Legal Professionals
Date: April 16, 2026
11:15am – 12:15pm

AI in Legal and Compliance Work: Ethical Duties and Practical Guardrails

Author: Edward So, Partner, Hogan Lovells US

Panelists

- Moderator: Joseph Ajibesin, Executive Director and Counsel, Morgan Stanley
- Evan Barnes, Assistant General Counsel, Clear Street
- Ian McGinley, Partner, Sidley
- Jeff Le Riche, Partner, Husch Blackwell
- Edward So, Partner, Hogan Lovells US

CLE Written Materials Submission

AI in Legal and Compliance Work: Ethical Duties and Practical Guardrails

By: Edward So[1]

AI has moved from experimentation into regular use across many legal, risk, and compliance workflows.[2] The question is no longer whether AI may affect legal practice. It already does. That is especially true in financial and commodity businesses, where legal and compliance teams are under constant pressure to summarize large information flows, translate regulatory developments quickly, and reduce friction in documentation and internal advice. The right question, then, is not whether AI belongs in practice. It is how to use it without compromising competence, confidentiality, supervision, accuracy, or client trust. The American Bar Association's 2024 formal opinion on generative AI recognizes that lawyers are already using these tools for legal research, contract review, due diligence, document review, regulatory compliance, and drafting.[3] But the same opinion makes the central point that should frame any discussion on the subject – at least for now, lawyers' use of AI is analyzed principally under existing professional duties rather than a separate AI-specific body of ethics rules.[4]

The key point is simple: for now, AI-related ethics issues are addressed primarily through existing professional duties rather than a separate AI-specific code. The same duties that already govern email, cloud storage, outsourcing, and electronic research now govern AI use as well. The most relevant are competence, confidentiality, supervision, communication with clients, candor to tribunals, and reasonable fees. The New York City Bar's 2024 opinion makes the same point in practical terms, describing AI guidance as "guardrails," not a new code.[5] Similarly, Florida's 2024 opinion permits use of generative AI, but only if lawyers protect client information, provide competent and accurate services, supervise appropriately, and avoid improper billing.[6]

1. Competence still comes first

Model Rule 1.1 requires competent representation.[7] Under the ABA's technology competence standard, that includes understanding the benefits and risks of relevant technology.[8] ABA Opinion 512 says lawyers need not become AI engineers, but they do need a reasonable understanding of the

---

[1] Edward So is a partner in Hogan Lovells' New York office, where he leads the firm's derivatives practice in the Americas; the views expressed here are solely those of the author.

[2] Thomson Reuters' 2025 Generative AI in Professional Services Report states that "About half" of professionals in legal, tax, accounting & audit, corporate risk & fraud, and government "use GenAI in some fashion," and adds that specialized generative AI tools are being "baked into professional workflows." The same report says corporate professionals are more likely to use generative AI for "compliance and reporting purposes," and that more than 40% of active users use generative AI at least daily, while just under three-quarters use it at least weekly.

[3] ABA Standing Comm. on Ethics & Pro. Resp., Formal Op. 512, Generative Artificial Intelligence Tools (July 29, 2024) ("**ABA Opinion 512**").

[4] For purpose of this article, "AI" in this article refers to generative artificial intelligence, meaning tools that can create new content—such as text, images, audio, video, or code—in response to user prompts.

[5] *See* N.Y.C. Bar Ass'n, Comm. on Pro. Ethics, Formal Op. 2024-5, *Ethical Obligations of Lawyers and Law Firms Relating to the Use of Generative Artificial Intelligence in the Practice of Law* (2024) at 2.

[6] The Florida Bar, Ethics Opinion 24-1 (Jan. 19, 2024).

[7] Model Rules of Pro. Conduct ("**Model Rule**") r. 1.1 (Am. Bar Ass'n 2023).

[8] Model Rule r. 1.1 cmt. 8; *see* ABA Opinion 512 at 2.

specific tool's capabilities and limitations.[9] That includes understanding drawbacks of AI, including the risk of hallucinations, bias, incomplete analysis, and misuse of confidential information. It also means, when using third party AI platforms (*e.g.*, ChatGPT or Grok) understanding the vendor terms, privacy policy, and who can access inputs and outputs.

From a legal and compliance perspective, 'competence' has an operational meaning. A lawyer can use AI to create a first draft of an internal issues list, summarize a batch of documents, or organize a chronology. But the lawyer may not hand off legal judgment to the machine. In a futures, derivatives, or commodity setting, that means AI may help with the first cut of an ISDA issues list, an onboarding checklist, a sanctions summary, or an internal policy draft. However, it cannot be trusted, without careful supervision and review, to decide whether a clause allocates credit risk correctly, whether a position or surveillance issue is material, whether a privilege concern has been captured accurately, or whether a regulatory characterization is sound. ABA Opinion 512 is explicit regarding this: lawyers may not abdicate professional judgment by relying solely on an AI tool for tasks that call for a lawyer's judgment.[10]

2. Confidentiality is still the hardest issue

Although competence provides the starting framework, the more immediate operational challenge for many lawyers is confidentiality. Model Rule 1.6 prohibits revealing information relating to the representation of a client without informed consent or another recognized basis. It also requires reasonable efforts to prevent inadvertent or unauthorized disclosure or access. ABA Opinion 512 applies that rule directly to AI use and states that, because many self-learning AI tools may directly or indirectly expose information relating to a representation, client informed consent is required before inputting such information into those tools.[11] The opinion also says lawyers should read and understand the tool's terms of use, privacy policy, and related contractual terms, or consult someone qualified who has done so.[12]

New York Ethics Rule 1.6 is framed slightly differently from the Model Rule. It prohibits disclosure of "confidential information," which the rule defines as "information gained during or relating to the representation of a client, whatever its source," if it is "protected by the attorney-client privilege," "likely to be embarrassing or detrimental to the client if disclosed," or "information that the client has requested be kept confidential."[13]

That point matters acutely in regulated trading and commodity businesses. The risk is not limited to a classic privileged memorandum. It includes draft deal terms, surveillance outputs, internal escalation notes, customer information, position or exposure summaries, sanctions analyses, investigation chronologies, and trading-related fact patterns that may be commercially sensitive even if not privileged. For outside counsel, it also includes work product, legal advice drafts, and client-specific facts pasted into a general-purpose model. The practical distinction is therefore not between "legal" and "non-legal" data. It is between data that can safely be used in a governed environment and data that should never be placed into an unapproved or self-learning system. ABA Opinion 512 and the New York City Bar's guidance both point in that direction by emphasizing confidentiality, the need for

---

[9] *See* ABA Opinion 512 at 2-3.
[10] *See id.* at 4.
[11] *See id.* at 6.
[12] *See id.* at 7.
[13] N.Y. Rules of Pro. Conduct r. 1.6(a) (2025).

appropriate safeguards, and the importance of understanding a tool's terms and data-handling practices.

The same logic extends to AI meeting tools. The New York City Bar's 2025 opinion on AI recording, transcription, and summarization concludes that a lawyer should obtain client consent before recording a client call, should consider confidentiality and privilege risks, and should check the output for accuracy if it may be preserved or relied upon.[14] That is not an edge case anymore. Many platforms now make recording and summarization easy. For lawyers, convenience does not replace the consent requirement.

3. Supervision matters more

In the legal and compliance context, AI-related problems often depend not only on the tool itself, but also on whether its use is subject to clear policies, appropriate training, and sufficient review. Model Rule 5.3 is therefore directly relevant. It requires lawyers with managerial or supervisory authority over nonlawyers to make "reasonable efforts to ensure that" their conduct "is compatible with the professional obligations of the lawyer."

The accompanying Comment makes the point in practical terms: "A lawyer must give such assistants appropriate instruction and supervision concerning the ethical aspects of their employment, particularly regarding the obligation not to disclose information relating to representation of the client."[15]

ABA Opinion 512 applies those same principles directly to generative AI. It states that "[m]anagerial lawyers must establish clear policies regarding the law firm's permissible use of GAI," and that supervisory lawyers must make "reasonable efforts to ensure that the firm's lawyers and nonlawyers comply with their professional obligations when using GAI tools." The opinion adds that supervisory obligations include ensuring that lawyers and nonlawyers are trained in the ethical and practical use of the relevant tools and the risks associated with them.

For New York lawyers, New York Ethics Rule 5.3 is materially similar for present purposes. It provides that a law firm must ensure that nonlawyers' work is "adequately supervised, as appropriate," and that "the degree of supervision required is that which is reasonable under the circumstances."

In a regulated entity, that principle has a practical consequence. If an AI tool will be used in legal or compliance workflows, the controls around that tool should not be left solely to one lawyer or to the vendor's default terms. They should be informed by the functions that will use the tool or bear the consequences of its use, including legal, compliance, information security, and relevant business or operational personnel. That is the kind of structure that helps make supervision "reasonable under the circumstances."

---

[14] *See* N.Y.C. Bar Ass'n, Comm. on Pro. Ethics, Formal Op. 2025-6, *Ethical Issues Affecting Use of AI to Record, Transcribe, and Summarize Conversations with Clients* (Dec. 22, 2025) (stating that "an attorney should obtain client consent before recording the call, should consider whether recording, transcribing and summarizing is tactically well-advised in the particular circumstances, including issues of confidentiality and privilege, and should check the work product for accuracy if there is a possibility the transcription or summary will be preserved and potentially relied upon").

[15] Model Rule 5.3, cmt 2.

4. Client Disclosure and Consent

Model Rule 1.4 and New York Ethics Rule 1.4 require lawyers to keep clients reasonably informed and to explain matters sufficiently to permit informed decisions. ABA Opinion 512 does not say that every AI use must be disclosed, but it does say that the facts of each case determine whether disclosure is required, and it specifically identifies several situations where consultation is needed: where client informed consent is required under Rule 1.6, where AI use is relevant to the basis or reasonableness of the fee, or where the tool's output will influence a significant decision in the representation.[16] It also notes that clients may separately require disclosure by engagement terms or outside counsel guidelines.

That is the sensible approach. Routine grammar assistance or idea generation with no client information may not require disclosure. But the answer changes where the lawyer plans to input client-specific information, where AI materially shapes an important recommendation, where the client has restricted AI use, or where the use affects staffing or billing. In a market-facing practice, disclosure issues are more likely to arise when AI touches negotiation positions, investigations, regulator-facing submissions, or high-value transactional advice where the client retained the lawyer for judgment, not just speed.

5. Courts are treating AI mistakes as lawyer mistakes

The case law already gives a clear message. Courts are not treating AI misuse as a novelty excuse. They are treating it as a failure of reasonable inquiry, candor, or competence.

The leading case remains *Mata v. Avianca.*[17] There, Judge Castel stressed that there is "nothing inherently improper" about using a reliable AI tool for assistance, but he sanctioned the lawyers under Federal Rule of Civil Procedure Rule 11 ("**FRCP Rule 11**") after they submitted non-existent opinions with fake quotations and citations and then persisted after the problem was flagged.[18] The court imposed a $5,000 joint penalty and additional non-monetary sanctions against the law firm and the lawyers involved. The court's framing is straightforward: technology is permissible, but attorneys still have the gatekeeping obligation to ensure accuracy. The Second Circuit applied a similar rationale in *Park v. Kim*.[19] Counsel admitted citing a non-existent case generated by ChatGPT and had not confirmed its validity.[20] The court referred counsel to its Grievance Panel and stated that, at a

---

[16] *See* ABA Opinion 512 at 8 (stating "The facts of each case will determine whether Model Rule 1.4 requires lawyers to disclose their GAI practices to clients or obtain their informed consent to use a particular GAI tool" and "Lawyers must also consult clients when the use of a GAI tool is relevant to the basis or reasonableness of a lawyer's fee."); and at 8-9 (stating "Client consultation about the use of a GAI tool is also necessary when its output will influence a significant decision in the representation.").

[17] *Roberto Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

[18] *Id.* at 448. FRCP Rule 11 states in relevant part that "By presenting to the court a pleading, written motion, or other paper—whether by signing, filing, submitting, or later advocating it—an attorney or unrepresented party certifies that to the best of the person's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances:…the claims, defenses, and other legal contentions are warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law or for establishing new law;…"

[19] *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024).

[20] *See id.* at 614-615.

minimum, FRCP Rule 11 requires attorneys to read and confirm the existence and validity of the authorities on which they rely.[21]

New York also has an important evidentiary development. In *In re Weber*,[22] a Surrogate's Court held that, before evidence generated by an AI product or system is introduced, counsel has an affirmative duty to disclose the AI use and the evidence should be subject to a Frye hearing before admission. That is not yet a general New York rule for all AI-shaped evidence, but it is a warning that AI-generated or AI-dependent evidence may trigger a requirement to disclose to court.[23]

These cases line up with the rules. Model Rule 3.3 and New York Ethics Rule 3.3 prohibit false statements of fact or law to a tribunal. ABA Opinion 512 explicitly links AI misuse in litigation to Rules 3.1, 3.3, and 8.4(c), and warns that outputs must be reviewed carefully before submission. In short, AI may help draft; it may not relieve counsel of the duty to verify.[24]

The privilege consequences of AI use are also beginning to appear in the case law. In *United States v. Heppner*, the Southern District of New York court held that a defendant's exchanges with Anthropic's Claude were not protected by the attorney-client privilege and further stated that, even if some of the information input into Claude had originally been privileged, any such privilege was waived by sharing that information with Claude and Anthropic "just as if he had shared it with any other third party," particularly given Anthropic's privacy policy and the resulting lack of a reasonable expectation of confidentiality. By contrast, in *Warner v. Gilbarco, Inc.*, the Eastern District of Michigan court declined to find a waiver of work-product protection based on a litigant's use of ChatGPT, reasoning that work-product waiver generally requires disclosure to an adversary or in a manner likely to reach an adversary, and adding that "ChatGPT (and other generative AI programs) are tools, not persons." Although *Warner* addressed work product rather than the attorney-client privilege, the two cases together suggest that courts may distinguish between open AI use that undermines confidentiality and other uses that do not, depending on both the doctrine at issue and the surrounding facts. Protective orders are moving in the same direction: for example, a 2025 protective order in the District of Kansas provided that "[p]roviding Confidential Information to an open tool is considered disclosure to a third party," while a 2026 stipulated protective order in the District of Nevada barred use of "Public Artificial Intelligence" tools for protected information but allowed secure, non-public AI tools subject to contractual and confidentiality safeguards.[25]

---

[21] *See id*.
[22] *In re Weber*, 2024 NY Slip Op 24258 (Sur. Ct., Saratoga Cnty. Oct. 10, 2024).
[23] *See id.* at 26 (stating "…due to the nature of the rapid evolution of artificial intelligence and its inherent reliability issues that prior to evidence being introduced which has been generated by an artificial intelligence product or system, counsel has an affirmative duty to disclose the use of artificial intelligence and the evidence sought to be admitted should properly be subject to a Frye hearing prior to its admission, the scope of which should be determined by the Court, either in a pre-trial hearing or at the time the evidence is offered").
[24] *See id.* at 9-10.
[25] *See Pathways Mgmt. Servs., LLC v. Northstar Hosp.,* No. 6:24-cv-01055-EFM-ADM, Protective Order, Doc. 84 (D. Kan. Jan. 23, 2025) (providing that "[p]roviding Confidential Information to an open tool is considered disclosure to a third party") and *9539-7337 Quebec Inc. v. Elko Wire Rope & Mining Supply, Inc., No.* 3:25-cv-00756-ART-CSD, Protective Order, Doc. 65 (D. Nev. Mar. 24, 2026)(defining a "Public Artificial Intelligence Tool" as a system made available broadly that "uses, stores, trains on, or otherwise retains user-submitted content" or permits access by persons outside the receiving party's control, while helpfully stating the use of AI tool is not prohibited if: "(a) the tool is used in a secure, non-public environment; (b) the tool's provider is contractually or legally prohibited from using Protected Information to

6. Billing for AI-Assisted Work

Model Rule 1.5 prohibits unreasonable fees and expenses. ABA Opinion 512 applies that rule directly to AI-driven efficiency. Lawyers billing by the hour must bill actual time, not hypothetical time that would have been spent without AI. The opinion also distinguishes between actual out-of-pocket expense, permissible direct cost allocation in some circumstances, and ordinary overhead.[26] It further says lawyers may not bill clients to learn how to use a tool they plan to use regularly, because maintaining competence in tools the lawyer uses is part of the lawyer's own professional obligation.[27]

That rule is especially important in institutional practice, where AI may reduce drafting or review time sharply on repeat tasks. Faster work is not unethical. Pretending the work still took the old number of hours is. The cleaner answer, when appropriate, is candid fee communication, fixed-fee or value-based pricing, or agreed treatment of direct tool expense.

7. Practical guardrails for Monday morning

Use AI most freely for low-risk internal efficiency work: organizing public material, generating first-cut checklists, preparing non-confidential summaries, or translating legal analysis into business language for internal discussion. Use it much more cautiously for anything involving client facts, litigation or regulatory submissions, witness materials, privileged investigations, draft legal advice, surveillance narratives, sanctions analyses, or deal language that allocates legal or economic risk.

Before using a tool, ask four questions. What data is going in? Who will see the output? What independent review will follow? Does the client, firm, or employer permit this use? If the prompt contains client-specific information, confidential commercial material, or internal investigative facts, the confidentiality analysis should come first. If the output will go to a court, regulator, counterparty, or client as advice, the verification standard should be much higher. And if the tool is unapproved or self-learning, the safest default is not to use it for sensitive information at all.

Conclusion

AI can make legal and compliance teams faster. It can help lawyers and compliance officers cut through volume, draft from a standing start, and organize information that would otherwise take much longer to process. But the ethics law has not changed in its essentials. The lawyer still must be competent. The lawyer still must protect confidential information. The lawyer still must supervise. The lawyer still must communicate with the client when disclosure or consent is required. The lawyer still must verify what is filed, sent, or relied upon. And the lawyer still must bill honestly. That is true in every practice. It is simply more visible in regulated financial and commodity businesses, where the information is more sensitive and the consequences of error can be more severe. For now, existing professional duties provide the principal framework for evaluating lawyers' use of AI. But as technology evolves and its use becomes more deeply embedded in legal practice, courts and bar authorities will

---

train or improve any artificial intelligence or machine-learning models available to persons outside the litigation team; (c) Protected Information is not disclosed to any person or entity not otherwise authorized to receive such information under this Order; and (d) the use of the tool is otherwise consistent with the confidentiality obligations imposed by this Order)").

[26] *See* ABA Opinion 512 at 12-13.

[27] *See id.* at 14.

likely continue refining that framework, and may ultimately conclude that more AI-specific guidance is warranted.