



AI and the Rise of Shadow Legal Work

Commercial and Contracts

Information Governance

Intellectual Property

Law Department Management



Banner artwork by patpitchaya / Shutterstock.com

Cheat Sheet

- **AI is broadening decision-making beyond legal teams.** Employees now use generative AI to interpret contracts, regulations, and policies, resulting in “shadow legal work” outside standard review processes.
- **This shift creates a governance gap.** While controls remain role-based, AI allows non-lawyers to perform legal-adjacent tasks, enabling high-risk decisions without proper visibility or accountability.
- **Legal AI risks are often hidden.** Errors may only become apparent after execution, when they are embedded in binding agreements or external commitments and are difficult to correct.
- **In-house counsel must redefine AI governance.** Establishing clear boundaries, requiring human review for legal-impact decisions, ensuring transparency, and providing training on AI failure modes are essential to managing risk without slowing business operations.

Introduction: A decision made too easily

A commercial lead is finalizing a multimillion-dollar agreement. When the counterparty suggests changes to intellectual property ownership, the lead bypasses legal review and consults a generative AI tool, asking: “Is this standard?”

The AI provides a clear and confident response. The language appears reasonable, so the commercial lead, under time pressure, approves the change. While this seems efficient, from a governance perspective, a significant legal decision has occurred outside established review processes, lacking visibility and accountability. This pattern is becoming increasingly common.

AI is redrawing role boundaries

Discussion around generative AI often focuses on productivity gains like faster drafting, quicker analysis, and reduced workload. However, a more significant shift is emerging.

Early evidence shows generative AI expands the scope of work employees undertake, encouraging them to operate beyond traditional role boundaries. By lowering barriers to complex tasks, AI enables employees to attempt work that previously required specialized expertise, including those with legal and regulatory implications.

Employees across organizations are drafting contract language, interpreting regulatory requirements, generating compliance content, and making judgment calls with legal consequences. These actions often do not feel risky. AI outputs are fluent, structured, and confident, which creates a sense of reliability even when the underlying analysis is incomplete or incorrect.

This has led to a subtle but important shift: decision-making is expanding faster than governance structures are evolving.

“Shadow Legal Work” in practice

This dynamic is increasingly visible in daily work. Employees outside the legal function, such as scientists, commercial leads, and operations managers, use AI tools to answer questions about contract terms, regulatory expectations, or policy language. The responses are often plausible enough to act on, especially when time constraints make formal escalation seem inefficient.

This behavior is understandable. AI tools are fast, accessible, and appear informed, while escalation introduces delays. This is known as “shadow legal work,” where non-legal professionals make legal decisions enabled by AI.

AI tools are fast, accessible, and appear informed, while escalation introduces delays.

Why this risk is different

Organizations are familiar with risks from AI-generated outputs in areas like software development. While these risks are meaningful, they are often visible and correctable. Legal and regulatory risks

behave differently.

Errors in legal reasoning or contractual interpretation may go unnoticed until after execution, becoming embedded in binding agreements or external communications and creating immediate or irreversible exposure. Unlike technical defects, flawed IP clauses or misinterpreted regulations may not be immediately apparent. These issues may only surface later as liability, disputes, or enforcement actions, when remediation is much more difficult.

Legal and regulatory risks behave differently.

The governance gap

Most organizations rely on role-based governance, where legal reviews contracts, compliance interprets regulations, and business teams execute within defined boundaries. Generative AI disrupts this model. Employees can now perform tasks outside their traditional roles without triggering standard controls. Additionally, AI use is often informal and undisclosed, making boundary crossings difficult to detect.

This creates a practical mismatch: authority remains role-based, while capability is now tool-based. Without adjustment, high-stakes decisions may occur outside established accountability structures.



This is not rogue behavior

This is not a case of misconduct. Employees are responding rationally to familiar pressures, such as the need for speed, the ease of obtaining seemingly expert guidance, and a desire to solve problems independently.

As these behaviors become normalized, they reshape expectations around responsibilities and decision speed. Over time, this can erode the boundaries that governance frameworks depend on.

Building an AI practice that includes legal risk

Many organizations are beginning to define “AI practices,” including the norms, controls, and workflows that govern AI tool use. For in-house counsel, these practices must explicitly address legal and governance implications.

Organizations should begin by defining clear domain boundaries. Low-risk activities, such as internal drafting and summarization, can proceed without additional controls. More sensitive tasks, including contract drafting, policy development, or regulatory interpretation, should require legal or compliance review, with defined triggers in workflows. Providing legal advice or approving contractual terms without counsel involvement should remain restricted.

It is critical to reinforce that AI is not an authority. AI outputs should be treated as drafts, not final decisions. Any output affecting legal rights, obligations, or external commitments must be validated by a human. This principle should be consistently reflected in playbooks, policies, and training.

Organizations must preserve accountability. AI should not obscure responsibility for decisions. Establish clear review thresholds, and ensure accountability remains with a named individual, regardless of AI involvement. Simple prompts, such as asking whether AI was used in drafting or evaluating terms, can reinforce this without slowing execution.

Organizations must preserve accountability.

Increasing transparency around AI use is essential. Lightweight disclosure mechanisms can be integrated into existing workflows, such as contract intake forms or approval processes. This visibility enables organizations to monitor AI use and identify emerging risks.

Finally, training should address both how to use AI and its potential failures. Employees should understand common failure modes, such as overconfident but incorrect reasoning, lack of jurisdictional nuance, and inability to reflect organization-specific risk tolerance. Scenario-based training can help employees recognize when escalation is necessary.

[ACC Members: Download the **Artificial Intelligence Toolkit for In-house Lawyers**](#)

The role of in-house counsel

This shift creates an opportunity for legal departments to lead. Rather than reacting to isolated incidents, in-house counsel can define clear boundaries for AI use in legal-adjacent work, embed review points into business workflows, and partner with compliance, IT, and business teams to align

governance with how work is actually performed.

The objective is not to slow the business, but to ensure that increased speed does not lead to unmanaged risk.

What In-house counsel should do next

In the near term, legal teams can take several high-impact steps:

- Identify where “shadow legal work” is most likely occurring, such as contracting, partnerships, and regulatory interpretation.
- Update workflows to account for AI use by adding simple disclosure prompts (e.g., asking whether AI was used) and review triggers that route higher-risk decisions to legal or compliance.
- Issue clear, practical guidance defining when AI use is appropriate and when escalation is required
- Align with business leadership on expectations around speed, autonomy, and risk in AI-assisted decision-making.

These actions do not require a comprehensive overhaul of AI governance, but they can significantly reduce risk as organizations mature their approach.

Conclusion: Keeping pace with distributed capability

Generative AI is changing not only how work is performed, but also who feels capable of performing it.

As these capabilities become more widely distributed, employees will continue to operate beyond traditional role boundaries. The question is not whether this will happen, but whether governance structures will evolve to keep pace.

For in-house counsel, the priority is clear: as capability expands, accountability must not be diluted. Organizations that succeed will make AI use visible, keep legal judgment anchored in human oversight, and deliberately redesign governance to match how work is actually done.

[Join ACC for more AI insights!](#)

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its

authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

[Gregory Kline](#)



Senior Director, IP Counsel

Thermo Fisher Scientific

Gregory Kline is Senior Director, Intellectual Property and Transactions Counsel at Thermo Fisher Scientific, a Fortune 500 company headquartered in Waltham, Massachusetts. He previously practiced at Morgan, Lewis & Bockius LLP before moving in-house, and now advises on intellectual property, licensing, and complex transactions in the biotechnology and life sciences sectors. He earned a bachelor's degree from the University of California, San Diego, a Ph.D. in Biology/Immunology from the University of California, Los Angeles, and a juris doctor from the University of California College of the Law, San Francisco (formerly UC Hastings). He serves as a volunteer legal officer with the Civil Air Patrol, supporting its mission of developing youth leadership and mentoring future leaders, and as an election officer in San Francisco.