



Episode 9 Transcript: “AI in the Boardroom and the Courtroom: Insights From the Tulane Corporate Law Institute”

Scott: Welcome back to the Business Court Benches. This is episode nine, AI in the Boardroom and the Courtroom: Insights From the Tulane Corporate Law Institute.

I’m your co-host Scott Barnard, a litigation partner in Akin’s Dallas office.

Stephanie: And I’m your co-host Stephanie Lindenmuth, a litigation partner in Akin’s New York office. On our recent episodes, we have talked a lot about how corporate law is evolving, through statutory reform, forum shifts, and courts grappling with increasingly sophisticated business disputes.

Scott: That’s right. Last month in March of 2026, Stephanie and I attended the Tulane Corporate Law Institute and across virtually every panel, whether it was M&A or fiduciary duties or litigation strategy, artificial intelligence was a recurring theme.

Stephanie: Yes, AI was not treated as a theoretical future development.

It came up repeatedly as something boards, executives, and litigators are dealing with in real time. So today we want to focus on takeaways from Tulane on what AI is already changing and what those changes mean from a legal and governance perspective.

Scott: One of the strongest messages from Tulane was that AI represents more than just an incremental efficiency.

The panelists were describing it as a structural shift that’s changing how companies think about value and long-term growth.

Stephanie: That was especially apparent in the M&A discussions. Speakers noted that the M&A market since early 2025 has been unusual, driven notably by the rapid integration of AI across industries.

The panelists also pointed out that because of AI, businesses are increasingly valued based on long-term expectations, what the company may look like a decade or more into the future, rather than short-term metrics.

Scott: And what we’re seeing is that that creates a specific tension that deal lawyers are already navigating today.

When a material portion of the company’s enterprise value rests on AI driven projections, you get harder questions about which DCF inputs are actually defensible in an appraisal proceeding or in a fairness opinion. Courts and experts are going to need to distinguish



between a reasonable forward-looking estimate and speculation dressed up as financial modeling.

That line is increasingly difficult to draw.

Stephanie: It also puts a spotlight on the diligence side. Acquirers are still developing the frameworks to properly diligence AI capabilities. There is often a real gap between a target's AI narrative, meaning what it says in the pitch deck, and what the actual infrastructure looks like—the data rights, the model governance, the third-party dependencies.

Buyers who rely on AI driven valuations without scrutinizing those underlying assets may find themselves in post-closing disputes that go well beyond a typical earnout disagreement.

Scott: It also means that decisions made today are more likely to be revisited later, particularly if the expectations don't materialize as planned, and that's where process and documentation and oversight really matter.

Stephanie: Several panelists emphasize that unlike earlier tech cycles, many companies tied to AI today already generate real cash flow. But even in those cases, a significant portion of the deal rationale turns on expectations about how durable that advantage will be. That shift in how companies are valued also shows up in a more practical way through how AI is already being used inside the deal process itself.

Another point that came up repeatedly was how integrated AI already is in transaction workflows. The panelists discussed its use in diligence, drafting, internal analyses, and even preparing materials for the boards.

Scott: Yeah. And I think the board materials point deserves particular attention.

If a board is approving a significant transaction based in part on AI generated analysis of synergies or risk factors, the quality and traceability of that analysis may become important later.

Stephanie: Mhmm.

Scott: Whether it's in post-closing litigation, books and records demand, or a challenge to the process, directors should be asking specifically how their briefing materials were prepared and what review occurred before those materials reached the board.

Stephanie: At the same time, there was a strong consensus at the conference that convenience comes with risk.

One clear warning was that AI generated work product is still treated like any other document from a legal standpoint.



Scott: That's right Stephanie, and that includes discovery. The speakers cautioned against situations where companies retained AI generated summaries or conclusions, but failed to preserve materials or inputs that would've shaped them.

Stephanie: And there's a related risk on the substantive side. AI tools used in contract diligence can miss atypical provisions or mischaracterize complex terms. If a deal team relies on AI generated contract summaries without attorney review of the flag provisions, especially in industries with heavily negotiated or non-standard documentation, that can create real exposure.

Scott: What about Stephanie from the litigation perspective?

Stephanie: From the litigation perspective, that can also create real problems. Producing outputs without context invites questions about what information decision makers actually considered and whether the record tells a complete story.

Scott: And those risks don't stop at the transaction table.

They're already showing up in court. AI came up repeatedly in the litigation panels at Tulane, not as a hypothetical concern, but as something that courts are already starting to confront in active disputes.

Stephanie: Delaware Chief Justice Seitz addressed how often AI is now showing up in filings and proceedings, particularly in cases he noted involving self-represented parties.

But he also noted that Delaware judges themselves have access to AI tools and are thinking actively about how they should and should not be used.

Scott: One Delaware case discussed in detail at the conference really pulled those issues together. It was Fortis Advisors v. Krafton, a Delaware Court of Chancery decision issued last month arising out of a post-closing earnout dispute.

Stephanie: Yes. The buyer in that case had agreed to a deal structure that left the founders and CEO of the target in operational control with an earnout tied to the company's future performance. As the business began to perform well, the buyer became increasingly concerned that the earnout could end up costing far more than anticipated.

Scott: Yeah, and faced with that risk, the buyer CEO did not limit his analysis to discussions with counsel. He personally used ChatGPT, entering prompts, asking how the earnout might be avoided, and what strategic options were available if the buyer wanted to prevent it from being paid.

Stephanie: The AI responses became important evidence. The prompts and outputs showed that the idea of removing the executives and taking control of the company was being



explored specifically as a way to disrupt the contractual earnout, not as a response to misconduct or performance failures.

Scott: And when the buyer later terminated the founders and CEO and asserted that the terminations were for cause, the court was able to compare those justifications against what was reflected in the ChatGPT interactions.

Stephanie: And that comparison mattered. The court concluded that the explanations offered after the fact did not align with what the decision makers were actually thinking at the time.

The AI prompts helped demonstrate that the terminations were pretextual and driven by financial incentives tied to the earnout.

Scott: In other words, the AI usage didn't just sit in the background. It became a part of the factual record that the court relied on in determining and evaluating intent, good faith, and compliance with the contract.

Stephanie: Yes. And another key aspect of this case is privilege. Had the CEO sought guidance through counsel, the analysis and strategy discussions would likely have been protected by attorney-client privilege.

Scott: Instead, because the analysis occurred through a publicly available AI tool, the communications were discoverable, readable by the court, and ultimately harmful to the buyer's litigation position.

Stephanie: This is one of the most important practical lessons from the case, and it's worth being direct about it. The lesson is not "don't use AI." Instead, the lesson is understand where the privilege boundary is. When an executive is contemplating a course of action that could generate litigation, particularly something as consequential as terminating founders in the middle of an earnout period, that analysis should run through outside counsel. The moment it goes through a consumer AI tool on a personal account, you have likely created a discoverable record and you have broken privilege.

Scott: And the flip side is worth noting too. Companies need to be thinking carefully about their AI retention policies.

Not to avoid preserving things they shouldn't, but because what gets deleted or saved itself can become an issue. Litigation holds need to be drafted while thinking about AI tool usage. The question of whether a ChatGPT account history was preserved or deleted after a dispute arose is exactly the kind of issue that comes up in discovery disputes.

Stephanie: So in this case, the Court of Chancery ordered equitable relief, including reinstating the CEO and extending the earnout period to account for the disruption caused by



the improper terminations. The remedy itself underscores how seriously the court viewed the conduct.

Scott: And I think that the broader lesson from Fortis Advisors is that AI use can directly affect governance outcomes, privilege determinations, and the evidentiary record. It's no longer just a tool for efficiency. It can shape how courts understand corporate decision making.

Stephanie: And as AI becomes more common in the boardroom, companies should expect judges and opposing counsel to scrutinize not just what decisions were made, but how those decisions were reached and whether AI was used in that process.

Scott: And that's what makes Fortis Advisors particularly instructive. It's that it's not an outlier.

Courts aren't just seeing AI issues in the headline-grabbing deal disputes. They're encountering them now in much more routine litigation as well.

Stephanie: And it's worth framing that contrast explicitly. Fortis Advisors involves a sophisticated corporate actor using AI to interpret contractual obligations. What we are about to discuss in another case, *Gupta v. Safko*, involves an unsophisticated litigant using AI to generate court filings in a books and records action.

Scott: That's right, Stephanie. The AI use in *Gupta v. Safko* came out of a 220 demand, Section 220 demand, where the plaintiff who was self-represented sought inspection of corporate records from officers rather than the corporation itself. At the pleading stage, the defendants moved to dismiss and requested sanctions.

Stephanie: In reviewing the record, the court noted that portions of the plaintiff's submissions were likely generated using AI tools. While the court ultimately dismissed the action on pleading grounds, it took the opportunity to address the implications of generative AI use in litigation filings.

Scott: The court emphasized that using AI does not lessen a party's obligations under Delaware rules. Filing still must be accurate—they have to be grounded in existing law, and they have to be supported by legitimate legal authority regardless of whether they're drafted manually or with the assistance of AI.

Stephanie: Importantly, the court declined to impose sanctions in that case, but it did so largely because the plaintiff was acting pro se, and the deficiencies did not rise to the level of what the court called "glaring legal or factual misrepresentations."

The court also noted there was no clear showing of AI hallucinations, meaning the AI system did not generate information that was entirely made up.



Scott: And that leniency came with a clear warning. The court required the plaintiff to certify any future use of generative AI, including identifying the tool used, identifying where it was used, and confirming that all AI generated material had been reviewed for accuracy.

Stephanie: The message was unmistakable. Generative AI may be a tool, but it does not change who bears responsibility for what is filed with the court. Human judgment and verification remain nondelegable.

Scott: It's also important to note that the Delaware Court's approach is part of a growing but uneven patchwork across different jurisdictions.

Several federal courts have adopted standing orders requiring AI disclosure, for example, in the Northern District of Texas, while others have not. Some required disclosure only if AI was used for legal research, while others are applying it more broadly. For litigators handling multi-district or multi-form disputes, that lack of uniformity creates real practical challenges.

Stephanie: Yes, and the lenience of the court here should not be generalized. Other jurisdictions have taken a much harder line.

In the Southern District of New York, for example, the *Mata v. Avianca* line of cases resulted in sanctions for attorneys who submitted AI generated filings containing fabricated case citations.

That contrast is instructive. Pro se leniency is not a universal standard, and the risk calculus for represented parties can be materially different.

Scott: Seen together, these cases raise a broader question about where AI fits in litigation and decision making, and where its limits need to be clearly defined.

Stephanie: So what should boards and companies take away from all of this?

Scott: Well, first, I think AI is already influencing core business decisions—from valuation to transaction structure to litigation structure.

Stephanie: Yes. And second, governance discipline is critical. Poor documentation or careless use of AI can increase rather than reduce legal exposure.

Scott: And third, lawyers remain central to this process.

As one of the panelists at the Tulane Conference noted, clients are using these tools, whether lawyers like it or not. So effective counsel requires understanding both the technology and the risk it creates.



Stephanie: Yes. And fourth, written AI governance policies are no longer optional. Companies that do not have them are already behind.

At a minimum, those policies should address which tools are approved for use, what categories of information may or may not be entered into external tools, how AI outputs must be reviewed before anyone relies on them, and how retention and litigation hold obligations apply to AI generated content.

Fortis Advisors is a compelling illustration of what can happen in the absence of that kind of structure.

Scott: So, in closing, AI is reshaping how companies operate and how courts evaluate decision making. That makes this an area that deserves close attention. And the right approach is going to depend on the facts, the transaction, and the risk profile.

But treating AI as purely an operational shortcut and not a governance issue, that would be a mistake.

Stephanie: And perhaps the most important thing we'd leave listeners with is this: AI tools move faster than governance frameworks, court rules, and legal precedent. That gap between what the technology can do and what the institutional structures around it have caught up to is where most of the current risk lives.

The role of counsel right now is not just to understand the technology; it is to help clients build the structures that keep pace with it. That is a different kind of legal advice than most clients have needed before, and it is work that is already underway.

Scott: Thanks for listening to this episode of the Business Court Benches - AI and the Boardroom, and the Courtroom Insights From the Tulane Corporate Law Institute. I'm your co-host Scott Barnard, a litigation partner in Akin's Dallas office.

Stephanie: And I'm your co-host Stephanie Lindemuth, a litigation partner in Akin's New York office. Thanks for listening, everyone. We'll catch you next time.