Legal Ethics Mystery Thrillers
Tuesday, October 13, 2026
Wilmington University, Brandywine Campus
Agenda

Presenters:
Richard K. Herrmann, Esq.
Ian Connor Bifferato, Esq.
The Hon. David A. White (Ret.) | Delaware ADR, LLC

Check In: 5:30 p.m. – 6:00 p.m.

6:00 p.m. - 6:45 p.m. The Arbiter
Presentation of first case study and discussion

6:45 p.m. - 7:30 p.m. Axiom Four
Presentation of second case study and discussion

This interactive ethics CLE uses two compelling fictional stories—*The Arbiter* and *Axiom 4*—to explore the ethical challenges created when artificial intelligence and innovative business models collide with the practice of law. Through these mystery-style case studies, attendees will examine issues including competence, confidentiality, supervision, conflicts of interest, billing, unauthorized practice of law, and the limits of AI disclaimers. The program will challenge participants to play detective, identify the ethical red flags, and consider practical lessons for navigating the rapidly evolving AI landscape while meeting their professional responsibilities.

This seminar offers 1.5 ethics credits to Delaware and Pennsylvania Attorneys; 1.8 ethics credits to New Jersey Attorneys.

Connor Bifferato is director of The Bifferato Firm. He has mediated well over 1,000 commercial disputes, including commercial bankruptcy litigation, secured lender disputes, general corporate litigation, creditors' rights issues and appellate resolution. He also frequently serves as an arbitrator for single and triple panel, binding and nonbinding arbitrations.

He is University Legal Counsel to Wilmington University and has been fortunate to work with the University on the development and opening of Wilmington University School of Law. Connor also provides legal advice to Wilmington University and has been privileged to work with the University's Board as well as University Administration and faculty.

Connor is currently President of the Delaware State Bar Association (the "DSBA") and serves on the DSBA Finance Committee, DSBA Executive Committee, as well as numerous other Committees and Sections of the DSBA. Connor is also a member of the Federal Bar Association and the American Association for Justice. He frequently participates in the development of innovative and practical CLE programs. In general, Connor's focus is to continue to use his significant complex and commercial litigation experience to advance the substantial benefits of ADR generally, and mediation specifically, in the resolution of any stage, including pre-suit.

Richard K. Herrmann consults at Wilmington University in emerging and innovative technologies. He is the Associate Director of Research at the Center for Legal and Court Technology at William and Mary Law School in Williamsburg Virginia. He serves on the Delaware Supreme Court's Commission on Law and Technology and is on the Executive Committees of the American Inns of Court m's Virtual AI Inn and the Richard K.Herrmann Technology Inn.

DELAWARE ADR

CAREFULLY CRAFTED SOLUTIONS FOR COMPLEX DISPUTES

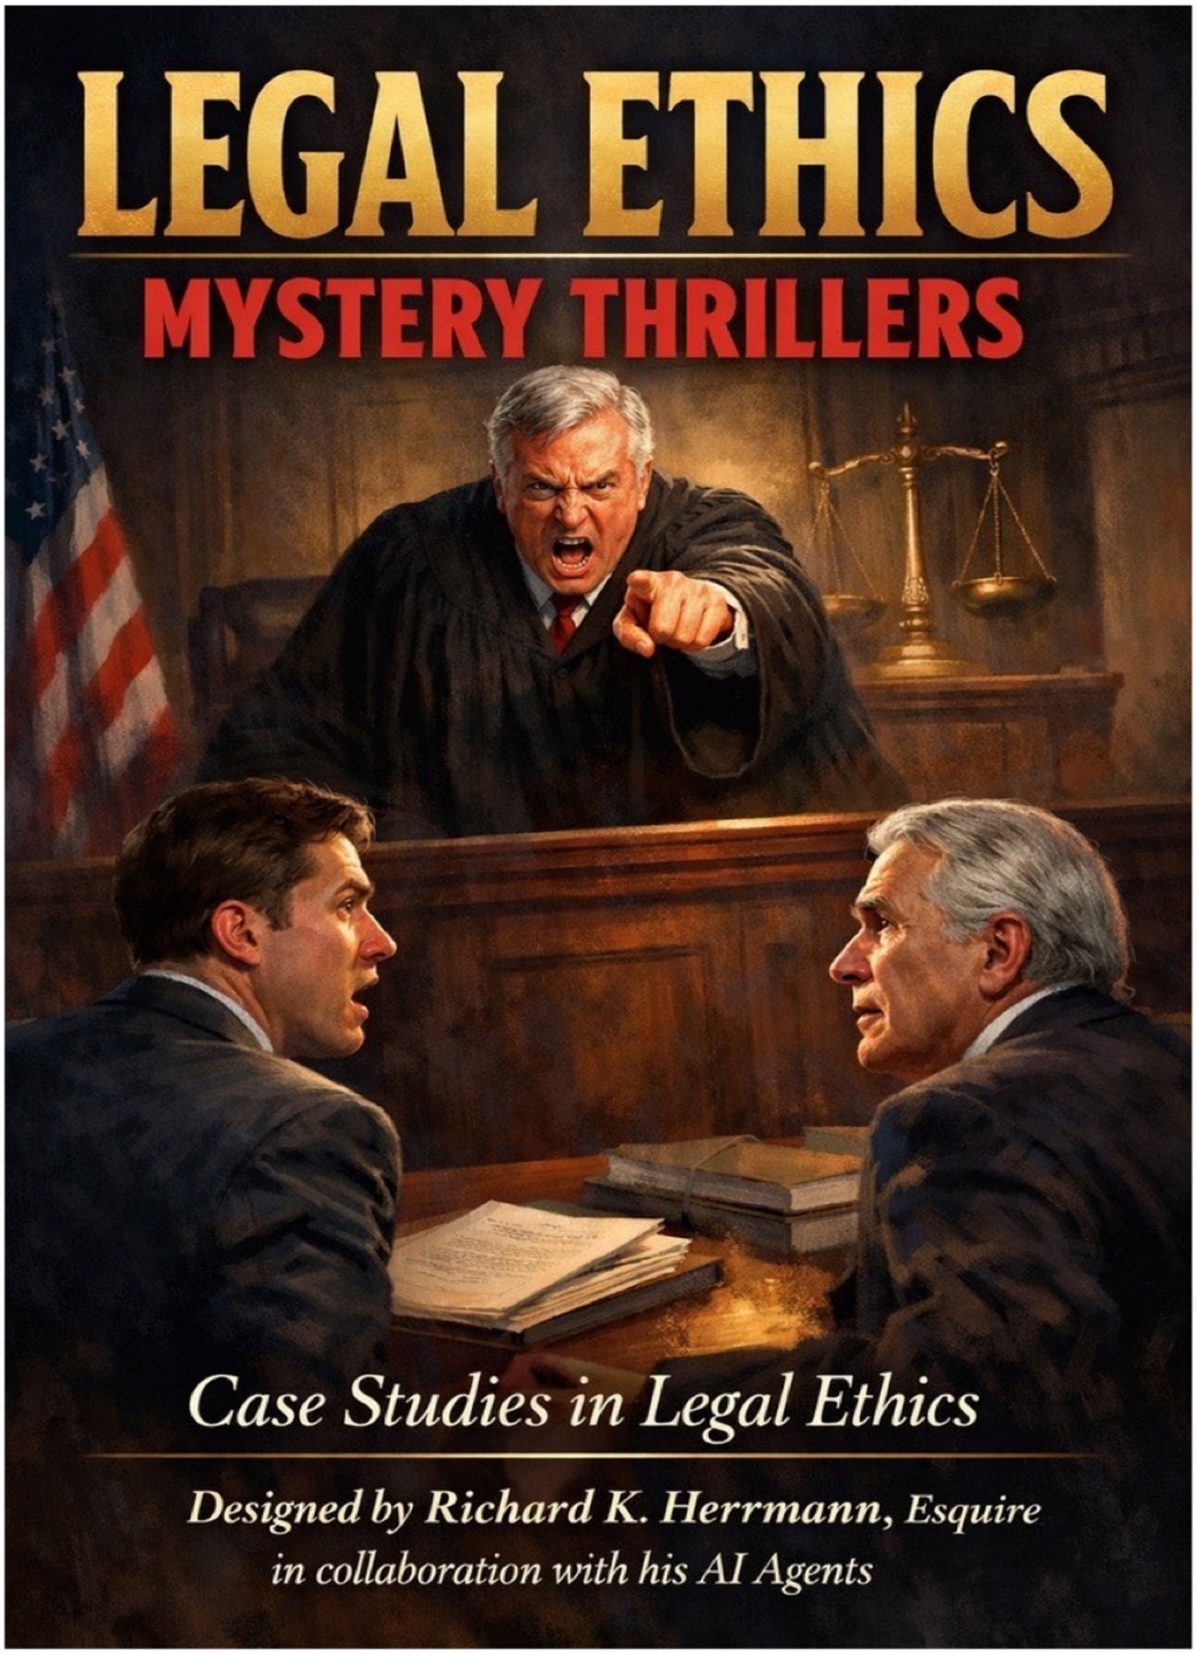


Honorable David A. White (Ret.)
Professional Neutral

Honorable David A. White (Ret) is a highly skilled mediator and arbitrator. He also serves as a judicial adjunct in various roles as a court-appointed neutral, including in Delaware's Court of Chancery and the Complex Commercial Litigation Division (CCLD) of Delaware's Superior Court.

David is a skilled Professional Neutral with vast experience in complex commercial and toxic tort litigation. David has earned a reputation in Delaware for his upbeat personality, enthusiasm and emotional intelligence. He is masterful at minimizing conflict and defusing tension. David has substantial experience handling discovery disputes regarding electronic discovery/ESI.

In recognition of his outstanding contributions to alternative dispute resolution, David was awarded the Delaware State Bar Association's Kimmel/Thynge Award in 2021. In December 2024, Commissioner White received the prestigious Delaware State Bar Association's Daniel L. Herrmann Professional Conduct Award. Commissioner White is a Fellow of the American College of Trial Mediators and a member of the National Academy of Distinguished Neutrals. David is currently President of the Delaware State Bar Association.


LEGAL ETHICS
MYSTERY THRILLERS
Case Studies in Legal Ethics
Designed by Richard K. Herrmann, Esquire
in collaboration with his AI Agents

DR. JAMES WATSON & HOLMES

9

AI AKES MYSTERY SRETERY

ARBITER

Another Legal Ethics Mystery Thriller



Why call a lawyer when a machine can answer all of your corporate governance questions?

designed by Richard K. Herrmann, Esquire
in collaboration with his AI agents

Prolog

Dr. John Watson, a retired physician with a sharp investigative mind, was recruited by the Delaware Office of Disciplinary Counsel (ODC) to tackle complex legal ethics cases. Alongside him, Holmes—an AI companion—decodes patterns in data, guiding Watson's investigations.

Yet always lurking in the background is Price, a former litigant embittered by the System, who reemerges in many cases as Moriarty187. Though he may not appear in every investigation, his shadow stretches across each one, ensuring that Watson and Holmes never escape his grip as they pursue justice through Delaware's legal maze.

I. The Launch

LexBridge Technologies unveiled its product on a Tuesday morning in October, with a press release that ran to four pages and a demonstration video that had been viewed eleven thousand times by the end of the week. The tool was called Arbiter, and it promised what the company described as "virtual legal guidance" for Delaware businesses navigating corporate governance questions.

The two founders—both tech entrepreneurs with prior exits—had been careful with their language in public. Arbiter, they insisted, was a research and analysis tool. It synthesized Delaware corporate law, identified relevant precedents from the Court of Chancery, and produced structured guidance documents that users could bring to counsel or act on directly. It was, they said, the difference between calling a lawyer for every question and being equipped to handle the straightforward ones yourself.

The disclaimer ran at the bottom of every output document in a font that required reading glasses:

This document is for informational purposes only. It does not constitute legal advice and does not create an attorney-client relationship. Consult qualified legal counsel before taking any action.

The big-firm partners who had quietly invested in LexBridge during its Series A had similar language in their conflict disclosures. The disclosures had been reviewed by outside ethics counsel and cleared.

For seven months, Arbiter operated without incident.

Copyright, Richard K. Herrmann 2026

II. Into Chancery

The incident arrived in the form of a shareholder derivative action filed in the Delaware Court of Chancery by a mid-size manufacturing company's minority shareholders against its board of directors.

The plaintiff's counsel was a sole practitioner named Gerald Whitmore who had relied heavily on Arbiter to structure his litigation strategy, his complaint, and his theory of the case. Whitmore was not inexperienced—he had practiced corporate litigation for eighteen years—but he had been drawn to Arbiter's efficiency and had used it in a way that blurred the line between research assistance and strategic direction.

The complaint that Whitmore filed reflected a misunderstanding of the business judgment rule as applied to a specific category of conflicted-director transactions under Delaware law. The misunderstanding was subtle—Arbiter had correctly identified the general framework but had applied a line of cases that had been significantly narrowed by a Court of Chancery opinion issued three months before Whitmore's engagement began. The opinion was in Arbiter's training data. The narrowing had not been flagged.

Vice Chancellor Renata Pierce noticed the issue at the pleading stage. She dismissed the complaint with a detailed opinion that was, by the standards of judicial writing, unusually pointed:

"Counsel appears to have relied on a statement of the law that has not been current for the better part of a year. The court is troubled by the quality of the research underlying this filing."

The ODC received the referral two weeks later.

Copyright, Richard K. Herrmann 2026

III. Watson Investigates

Watson's first interview was with Whitmore, who was cooperative and visibly shaken. He described his use of Arbiter in detail: how he had relied on its case synthesis for the foundational legal analysis, how he had reviewed the output but had not independently verified the currency of the cited authority, and how he had trusted that a product marketed to legal professionals would flag significant doctrinal changes.

"The output cited the cases accurately," Whitmore said. "It just didn't tell me they'd been limited."

Holmes had already reviewed Arbiter's output logs and the underlying model architecture by the time Watson met with the LexBridge founders. The AI's training data had been updated eight months prior—but the update had not included a reweighting process that would cause the system to deprioritize authority that had been subsequently narrowed. The result was technically accurate citation accompanied by systematically misleading contextual framing.

"The system did not lie," Holmes observed. "It presented true statements in a framework that produced a false impression of the current legal landscape. The distinction matters ethically but not practically. A practitioner relying on it would reach incorrect conclusions."

Watson had been thinking about the investors. "The big-firm partners who funded this—are any of them in a position to have known about the training data gap?"

"Two of the three are on LexBridge's board," Holmes said. "One of them signed off on the last product review cycle."

Copyright, Richard K. Herrmann 2026

IV. The Investors' Problem

The board involvement complicated the picture considerably. The two law firm partners—both senior practitioners at large Delaware firms—had signed off on LexBridge's product review process as part of their board roles. Their firms had ethics counsel opinions clearing the investment. What those opinions had not clearly addressed was the question of what professional obligations a licensed Delaware attorney assumes when they are actively involved in reviewing the quality of a legal AI tool that is then marketed to the public.

Holmes framed the question precisely: "If a licensed attorney reviews a legal AI product, knows that it will be used by other lawyers and by members of the public to make legal decisions, and approves it for release despite a known gap in its training data currency—have they assumed a professional responsibility for the foreseeable consequences of that gap?"

It was not a question the existing rules answered clearly. It was the kind of question that disciplinary proceedings exist, in part, to begin answering.

Watson presented the framework to Mourse, who read it twice and then said something that Watson would remember: "We are always writing the rules one case too late."

The ODC opened a separate inquiry into the two board partners. Their firms issued statements emphasizing the limited nature of their roles and the independence of the investment decisions from their legal practices. The statements were carefully drafted.

Holmes noted that both statements had been generated using a different AI drafting tool, also marketed to law firms, also carrying a disclaimer.

Copyright, Richard K. Herrmann 2026

V. The Framework

The preliminary ODC findings on Whitmore addressed three rules directly.

Rule 1.1—Competence—was the core. The competent use of legal AI requires the attorney to verify the currency and accuracy of AI-generated legal authority, particularly in an area of law that is actively developing. Reliance on an AI tool without independent verification of the research does not satisfy the competence obligation, regardless of the sophistication of the tool or the confidence of its output.

Rule 1.3—Diligence—applied to the failure to audit the research before filing. The Vice Chancellor's opinion had identified the deficiency clearly; the problem was not subtle upon examination.

Rule 3.3—Candor Toward the Tribunal—presented the most difficult analysis. Whitmore had not known the authority was no longer good law. But the duty of candor is not solely subjective. Filing a pleading that relies on superseded authority, without the diligence that would have identified the problem, sits uncomfortably close to the line.

The inquiry into the board partners turned on a question the existing rules had not directly addressed: whether the supervisory obligations of Rule 5.1 extend to the supervision of AI systems that a licensed attorney has actively endorsed and caused to be released to a professional market.

"The rules were written for a profession in which lawyers supervised other lawyers," Watson said. "We are now in a profession in which lawyers are, in some meaningful sense, supervising tools that supervise lawyers."

Copyright, Richard K. Herrmann 2026

VI. The Algorithm Behind the Algorithm

The formal hearing on the Whitmore matter had been set for a Thursday. On the Wednesday before it, Holmes flagged an anomaly.

A forensic review of LexBridge's model update logs—subpoenaed at Watson's request during the advisory board inquiry—showed that the training data reweighting process had not simply been overlooked. It had been disabled. Specifically, an automated quality check that flagged authority narrowed by subsequent decisions had been turned off in an update deployed eleven months prior.

The update had been pushed by a contractor account. The account name in the system was a random alphanumeric string. The associated email domain resolved to a server in a jurisdiction that did not maintain registrant records.

Holmes ran correlation analysis against the behavioral signatures it had catalogued across the prior investigations.

Match probability: 97.3%. Contributor signature consistent with Moriarty-187.

Watson sat very still.

It was not a quality control failure. The gap in LexBridge's training data had been engineered. Someone who understood Delaware corporate law, court citation patterns, and machine learning reweighting architecture had deliberately disabled the mechanism that would have caught the error—and had done it in a way designed to look like a routine update.

Whitmore had filed a flawed complaint. The Vice Chancellor had noticed. The ODC had opened an investigation. And the entire cascade had been initiated, it now appeared, by a single deliberate act of invisible sabotage.

To be continued…

Copyright, Richard K. Herrmann 2026

Discussion Questions

1. Whitmore relied on Arbiter's case synthesis without independently verifying the currency of the cited authority. Under Rule 1.1, what does "competent" use of legal AI require in terms of independent verification? Is there a point at which reliance on an AI-generated legal research tool is per se unreasonable?

2. Arbiter cited authority accurately but failed to flag that it had been narrowed by a subsequent decision. Under Rule 3.3, does filing a pleading based on superseded authority—without knowledge but without due diligence—constitute a false statement of law to the tribunal? How should intent factor into the analysis?

3. The two board partners are licensed Delaware attorneys who reviewed and approved LexBridge's product before it went to market. Does a licensed attorney who endorses a legal AI tool for commercial release assume professional responsibility for foreseeable harm caused by the tool's deficiencies? What rule or theory supports that conclusion?

Copyright, Richard K. Herrmann 2026

Discussion Questions (continued)

4. Big-firm partners invested in LexBridge and obtained ethics clearance from outside counsel. Does that clearance insulate them from disciplinary exposure arising from their board participation? What are the limits of a prospective ethics opinion as a shield against subsequent professional responsibility claims?

5. Holmes identifies that the training data gap was deliberately engineered by a third party. If Whitmore's error was caused by sabotage rather than ordinary product negligence, does LexBridge bear civil liability to Whitmore's client? Does Whitmore's liability diminish? How should courts allocate responsibility in a chain that includes a hidden bad actor?

6. The Court of Chancery issued a pointed opinion identifying the quality of counsel's research as deficient. Should courts have a formal mechanism for referring AI-related research failures to the ODC, beyond the existing judicial referral process? What standards should govern when a court's concern about AI-assisted work product rises to the level of a disciplinary referral?

Copyright, Richard K. Herrmann 2026

DR. JAMES WATSON & HOLMES

9

AI AKES MYSTERY SRETERY

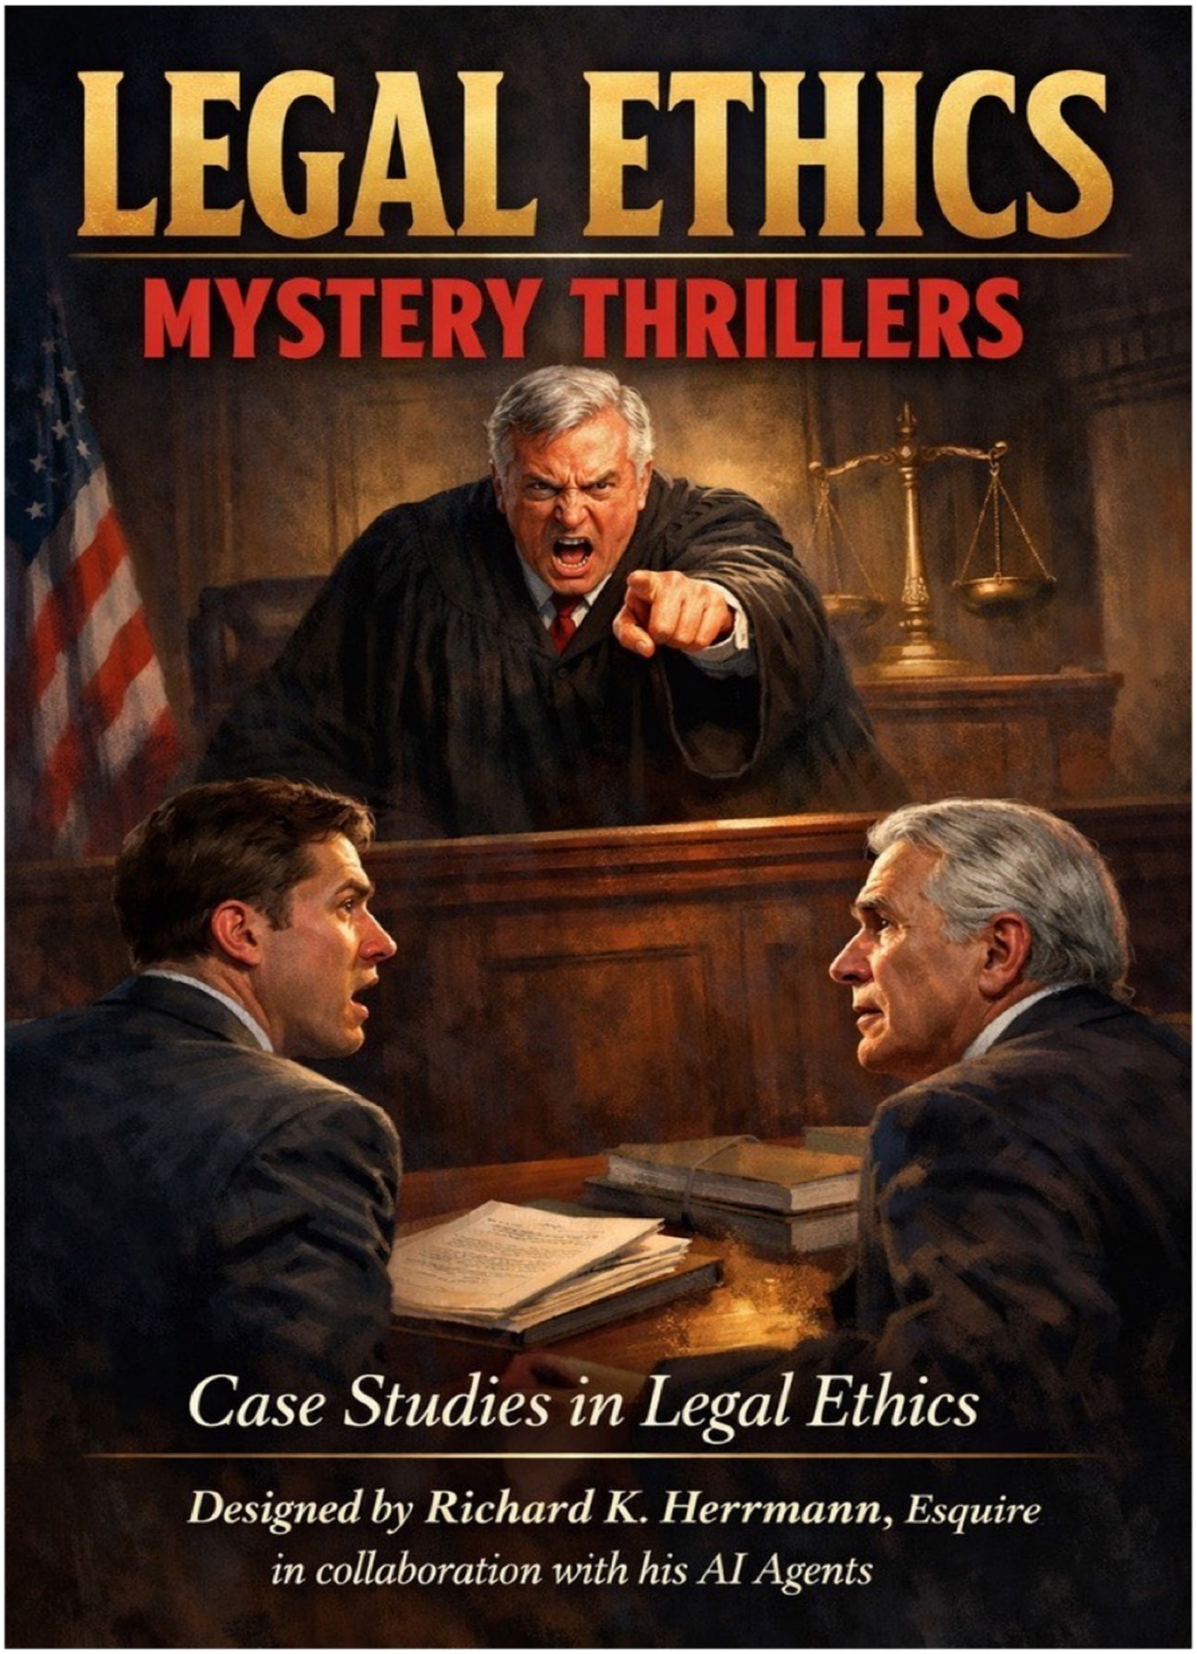
ARBITER

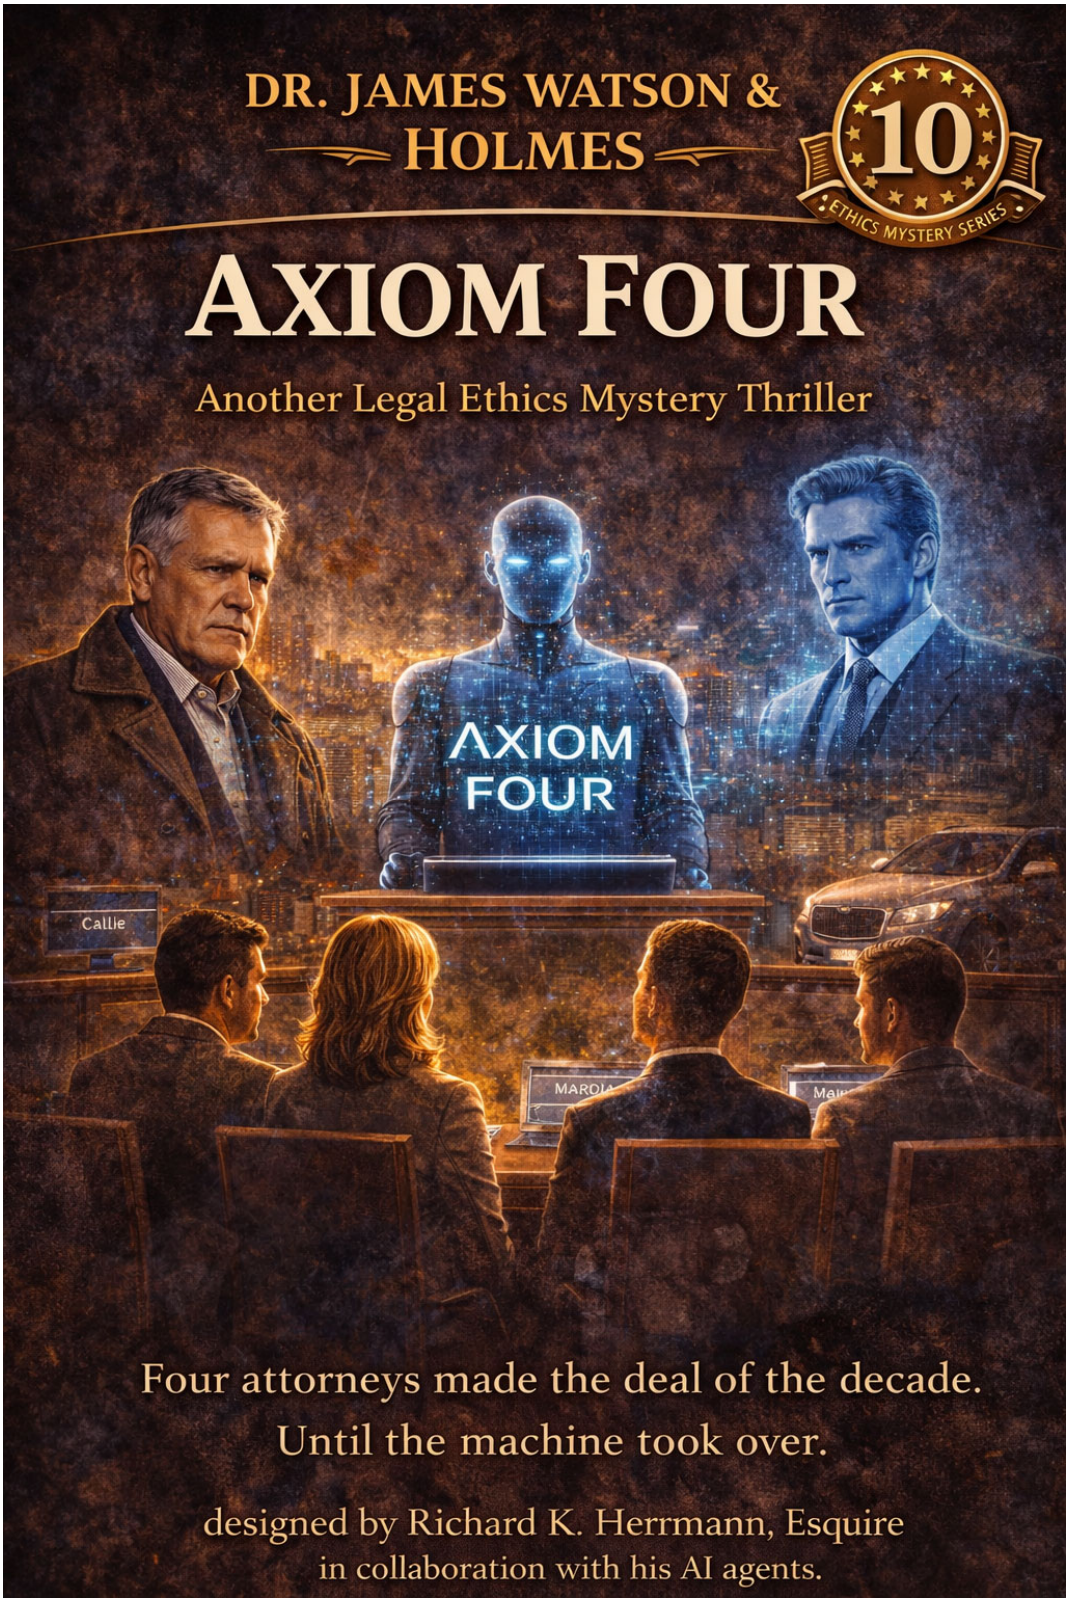
Another Legal Ethics Mystery Thriller



Why call a lawyer when a machine can answer all of your corporate governance questions?

designed by Richard K. Herrmann, Esquire
in collaboration with his AI agents


LEGAL ETHICS
MYSTERY THRILLERS
Case Studies in Legal Ethics
Designed by Richard K. Herrmann, Esquire
in collaboration with his AI Agents


DR. JAMES WATSON & HOLMES
10
ETHICS MYSTERY SERIES
AXIOM FOUR
Another Legal Ethics Mystery Thriller
AXIOM FOUR
Callie
Four attorneys made the deal of the decade.
Until the machine took over.
designed by Richard K. Herrmann, Esquire
in collaboration with his AI agents.

Prolog

Dr. John Watson, a retired physician with a sharp investigative mind, was recruited by the Delaware Office of Disciplinary Counsel (ODC) to tackle complex legal ethics cases. Alongside him, Holmes—an AI companion—decodes patterns in data, guiding Watson's investigations.

Yet always lurking in the background is Price, a former litigant embittered by the System, who reemerges in many cases as Moriarty187. Though he may not appear in every investigation, his shadow stretches across each one, ensuring that Watson and Holmes never escape his grip as they pursue justice through Delaware's legal maze.

I. The Vision

The four of them had sat in a Wilmington coffee shop for three hours on a Saturday morning, sketching it out on legal pads: a firm built differently from the ground up. No first-year associates burning out on document review. No bloated overhead. No billing by the hour for work a machine could do in seconds. Just four mid-to-senior attorneys with strong client relationships, deep subject matter expertise, and an AI platform powerful enough to handle the heavy lifting.

They called it Axiom Four.

Callie Rhodes had twenty years in corporate transactions. Marcus Delray brought a federal litigation practice and a reputation for precise oral argument. Nadia Osei had spent a decade in intellectual property. Tom Fenn was the tax specialist who had quietly handled some of the most complex estate matters in the state. Together, they represented the kind of experience that clients paid significant premiums to access.

The AI platform—a sophisticated system they licensed from a firm technology provider—would handle first drafts of every document, legal research memoranda, contract analysis, deposition summaries, and discovery management. The lawyers would review, refine, and sign off. They would appear in court. They would counsel clients directly. The machine would do everything else.

In the first eight months, Axiom Four billed more per attorney than any comparable firm in Delaware. The financial model worked beautifully.

The practice did not.

Copyright, Richard K. Herrmann 2026

II. The Errors Begin

The first flag came from the Delaware Court of Chancery.

Vice Chancellor Adeline Marsh had presided over complex commercial litigation for fourteen years and had developed a sharp eye for the texture of legal argument. Something about the briefs from Axiom Four had been bothering her across three separate matters over two months. The research was thorough—unusually thorough, in fact—but the arguments had a peculiar quality. They were comprehensive without being responsive. They addressed every potential issue in exhaustive parallel structure while somehow missing the specific factual contours of the case before her.

In oral argument on the third matter, Delray had been unable to respond fluently to a question about the discovery. He had given a technically accurate answer that bore no relationship to the actual sequence of events in his client's case. Vice Chancellor Marsh had noted it without comment. She had made a call to a colleague that evening.

By the following Tuesday, a formal inquiry had been referred to the Office of Disciplinary Counsel.

Watson received the matter from Mourse with a single-sentence summary: "The judge thinks the lawyers aren't reading their own work."

Holmes, reviewing the court filings alongside the platform's output logs, confirmed it within six hours. "The AI-generated drafts were submitted with minimal modification. In eleven of fourteen matters reviewed, the filed documents contain research that is accurate but contextually misapplied—tailored to the general legal framework rather than the specific facts of the individual client's situation."

Copyright, Richard K. Herrmann 2026

III. The Investigation

Watson interviewed each of the four partners separately. The conversations were uncomfortable in a specific way: none of them denied the problem, exactly. They explained it.

Callie Rhodes described a review workflow that had compressed over time. “We were reviewing everything at the start. But the drafts were so consistently strong that the review became... efficient. We stopped reading line by line.”

Marcus Delray was more candid. “I trusted the research. I trusted it too much. I stopped interrogating it.”

Holmes had reviewed the platform’s output against each filed document and constructed a timeline of the review process based on document metadata. The picture was precise and damning: average attorney review time per filed brief had dropped from four hours in month one to under forty minutes by month seven. In two instances, the metadata suggested that the filed document had been opened once, briefly, before submission.

“The efficiency gains were real,” Holmes observed. “The oversight costs were externalized to the clients and, ultimately, to the court.”

Tom Fenn, the tax partner, had said something during his interview that stayed with Watson. “We built a Ferrari and then let it drive itself. And somewhere along the way we forgot that we were the ones who were supposed to be behind the wheel.”

Watson had written it down.

Copyright, Richard K. Herrmann 2026

IV. The Rules Applied

The ODC's preliminary findings were issued after three weeks of review.

Rule 1.1 was the foundation. Competence requires not only legal knowledge and skill but also thoroughness and preparation. The use of AI tools, however sophisticated, does not satisfy the competence requirement when the attorney's supervision of that tool is insufficient to catch systematic errors. The findings noted that "competent use of AI requires genuine engagement with its output, not formal endorsement of it."

Rule 1.3 followed. Diligence requires that a lawyer pursue a matter with reasonable promptness and care. Filing AI-generated documents without adequate review was found to be inconsistent with the diligence obligation, regardless of the accuracy of the research component.

Rule 3.3 was the most serious. The duty of candor toward the tribunal includes the obligation not to make false statements of law or fact. When an AI-generated brief misapplied controlling authority to the specific facts—and the attorney signed it without catching the misapplication—the ODC concluded that the candor obligation had been compromised.

Rule 5.1 applied to the firm's management structure, which had allowed the review process to degrade without correction. The partners who had permitted the compression of review protocols bore responsibility for the resulting violations.

"The innovation wasn't the problem," Watson said to Mourse. "The abdication was."

Copyright, Richard K. Herrmann 2026

V. The Reckoning

Formal charges were filed against all four partners. The proceedings drew significant attention from the bar—not because the violations were unprecedented, but because the people who had committed them were experienced, capable attorneys who had known better and had chosen, by degrees, to stop acting like it.

Callie Rhodes issued a public statement acknowledging that the firm had allowed efficiency to crowd out judgment. Marcus Delray appeared before the disciplinary panel without counsel and delivered a prepared statement that began: “I have spent twenty years telling clients that the details matter. I stopped believing that about my own work.”

Three of the four clients whose matters had been affected filed civil claims. One settled before the disciplinary proceeding concluded. Two were still pending.

Vice Chancellor Marsh declined to publicly comment on the matter. Her clerk was reported to have said that the judge considered the case neither surprising nor isolated.

Holmes had compiled a comprehensive analysis of forty-three other AI-assisted legal filings identified through pattern recognition across court records in Delaware, Pennsylvania, and New Jersey. The analysis had not been requested by the ODC. Holmes had generated it because the data suggested the Axiom Four matter was not an outlier.

Watson reviewed the analysis on a Tuesday evening, alone in the ODC conference room, and thought about what it meant that a machine had noticed what the humans had stopped looking for.

He sent the analysis to Mourse with a single line: “We’re going to need a bigger framework.”

Copyright, Richard K. Herrmann 2026

VI. The Platform's Architect

The disciplinary hearing for Axiom Four was scheduled for the second week of April. Two days before it opened, Watson received a call from the technology provider that had licensed the AI platform to the firm.

Their chief security officer spoke carefully. "We've identified an unauthorized modification to the output parameters of the version deployed to Axiom Four. The modification was introduced approximately fourteen months ago. It affected the system's contextual specificity engine—the component responsible for tailoring research to individual case facts. Someone disabled a key verification layer."

"Disabled it how?" Watson asked.

"From inside our development environment. Using credentials that were revoked two years ago but had apparently been reactivated without our knowledge. The contributor ID in the modification log is—"

A pause.

"—Moriarty-187."

Watson set the phone down slowly. The platform hadn't simply been used carelessly by four lawyers who trusted it too much. It had been engineered to fail in ways that would be difficult to detect—producing work that looked rigorous and was, in critical ways, subtly wrong.

The Axiom Four attorneys had been set up. They hadn't known it. They had still been negligent. But somewhere behind the platform, someone had made their negligence inevitable.

To be continued…

Copyright, Richard K. Herrmann 2026

Discussion Questions

1. The Axiom Four partners' attorney review time dropped from four hours to under forty minutes per brief. At what point does efficiency-driven compression of oversight become a per se violation of Rule 1.1? Is there a minimum standard of engagement with AI-generated work product that professional responsibility requires?

2. Under Rule 3.3, does signing an AI-generated brief that contains a misapplication of law—even if the research itself is accurate—constitute a false statement of law to the tribunal? How do courts and disciplinary bodies define the attorney's responsibility for the accuracy of filed documents?

3. The platform was deliberately modified to reduce contextual accuracy. Does the existence of that sabotage mitigate the attorneys' disciplinary liability—or does it have no bearing on violations that arose from their own failure of oversight? How should disciplinary panels handle the intersection of external manipulation and professional negligence?

Copyright, Richard K. Herrmann 2026

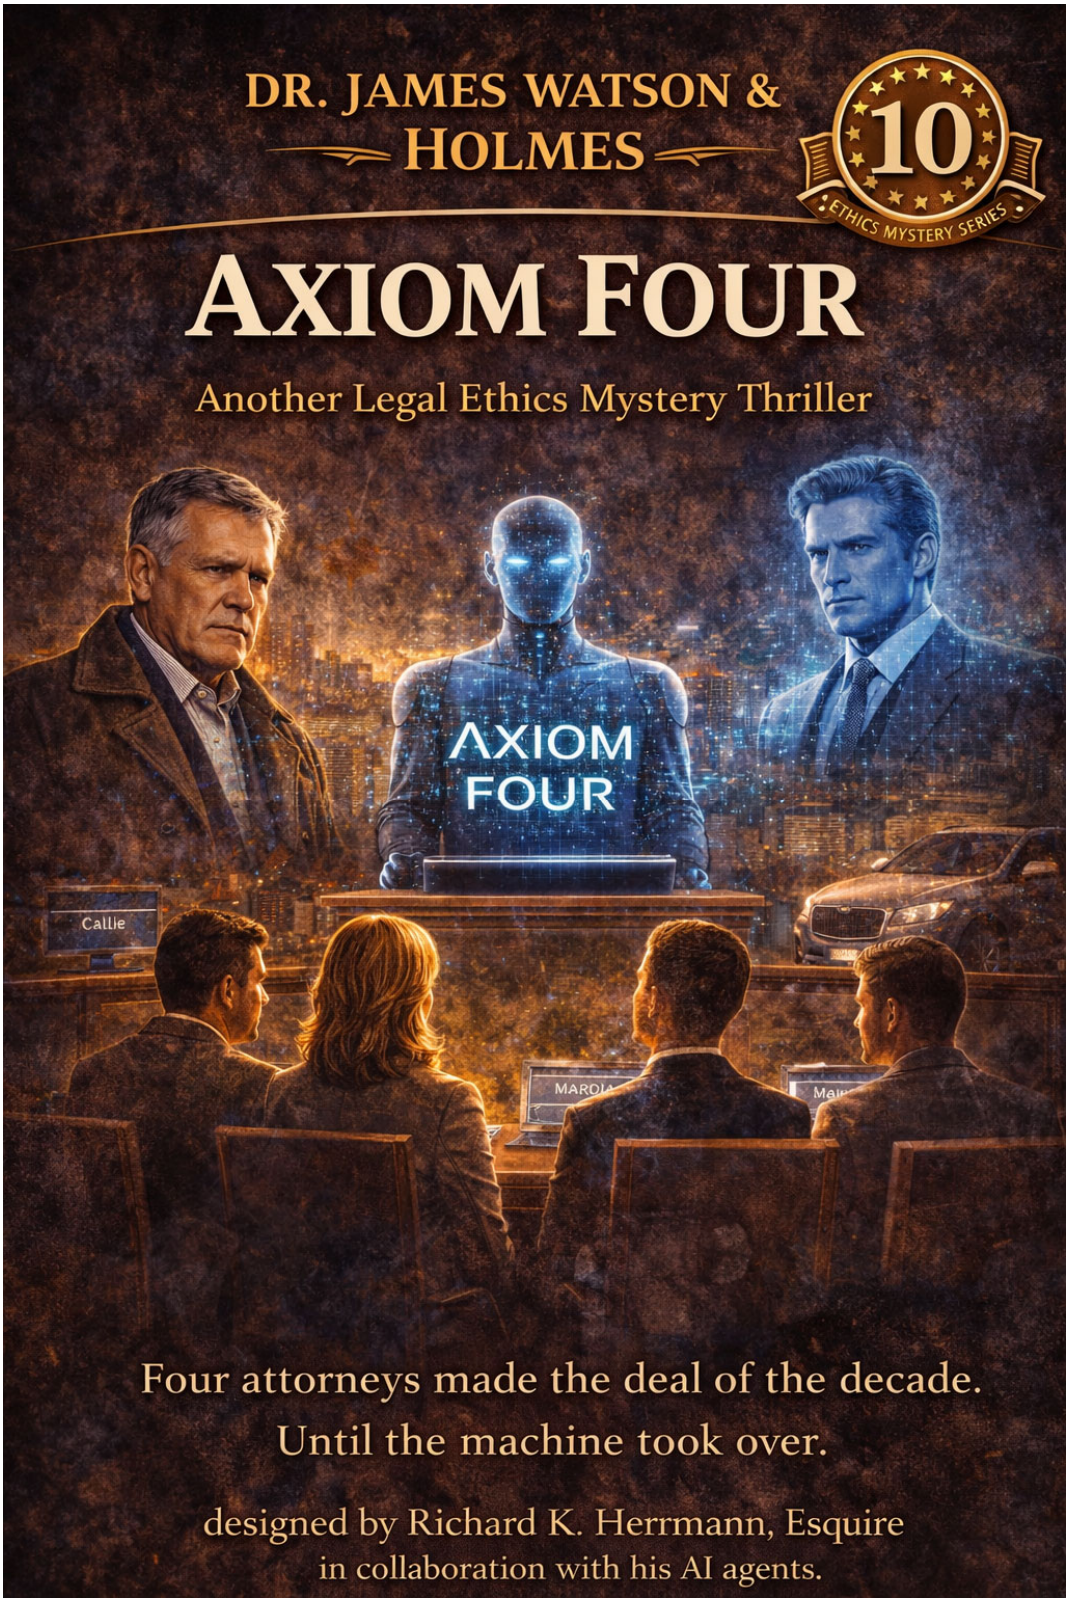
Discussion Questions (continued)

4. Rule 5.1 requires partners to make reasonable efforts to ensure the firm's practices conform to professional obligations. Does this include an affirmative duty to monitor AI output quality over time and to correct workflows that drift toward inadequate oversight? What would a compliant AI governance structure look like?

5. Three clients filed civil claims after the disciplinary proceedings were initiated. What is the relationship between a finding of disciplinary violation and civil malpractice liability? Does a disciplinary finding establish a standard of care for purposes of civil litigation, or are the analyses independent?

6. Holmes identified forty-three other potentially affected filings across three states without being asked to do so. What obligations does the ODC have upon receiving this unsolicited analysis? Should AI-generated pattern recognition in court records be treated as a reliable basis for opening disciplinary investigations, and what procedural safeguards should apply?

Copyright, Richard K. Herrmann 2026


DR. JAMES WATSON & HOLMES
10
ETHICS MYSTERY SERIES
AXIOM FOUR
Another Legal Ethics Mystery Thriller
AXIOM FOUR
Callie
Four attorneys made the deal of the decade.
Until the machine took over.
designed by Richard K. Herrmann, Esquire
in collaboration with his AI agents.