

INDEPENDENCE FIRST . . . OR DON'T USE IT: A FIDUCIARY STANDARD FOR LAW-FIRM USE OF GENERAL AI

EDWARD S. ADAMS*

I. INTRODUCTION.....	463
II. THE EMERGING CONSENSUS: MAPPING THE NEW REGULATORY FLOOR.....	466
A. AMERICAN BAR ASSOCIATION FORMAL OPINION 512	466
B. STATE BAR GUIDANCE.....	468
C. THE JUDICIARY’S WARNING: HUMAN OVERSIGHT IS NON- NEGOTIABLE	469
III. A THEORY OF VENDOR INFLUENCE: HOW GPAI THREATENS PROFESSIONAL INDEPENDENCE UNDER MODEL RULE 5.4	470
A. THE ALGORITHMIC VECTOR: INFORMATION ASYMMETRY AND THE BLACK BOX PROBLEM	471
B. THE ECONOMIC VECTOR: PRICING MODELS AS INFLUENCE MECHANISMS	473
C. SOPHISTICATED PRICING MODELS: HIDDEN INFLUENCE PERSISTS.....	475
D. THE PRINCIPAL-AGENT PROBLEM REDEFINED.....	476
IV. A THEORY OF FIDUCIARY-GRADE CONFIDENTIALITY: PRIVILEGE, WAIVER, AND THE THIRD-PARTY AI DOCTRINE.....	477
A. PRIVILEGE: THE FRAGILITY OF ATTORNEY-CLIENT PROTECTION	478
B. WAIVER: HOW THIRD-PARTY AI DISCLOSURE DESTROYS PRIVILEGE IRREVOCABLY	479
C. THE THIRD-PARTY AI DOCTRINE: FIDUCIARY-GRADE SAFEGUARDS AS THE ONLY DEFENSE.....	483

* Howard E. Buhse Professor of Finance and Law, University of Minnesota; M.B.A. 1997, Carlson School of Management at the University of Minnesota; J.D. 1988, University of Chicago; B.A. 1985 Knox College.

V. THE FGIS STANDARD: AN AUDITABLE FRAMEWORK FOR FIDUCIARY-GRADE AND INDEPENDENCE SAFEGUARDS...	484
A. PILLAR A: INDEPENDENCE SAFEGUARDS	485
i. <i>Treatment of AI Vendor as Supervised Nonlawyer Assistant (Rule 5.3)</i>	485
ii. <i>Retain Economic Neutrality (Rule 5.4)</i>	486
iii. <i>Model Knowledge and Inquiry</i>	486
iv. <i>Requisite Human Verification Protocol</i>	487
B. PILLAR B: FIDUCIARY-GRADE CONFIDENTIALITY SAFEGUARDS.....	488
i. <i>Private Deployment</i>	488
ii. <i>Contractual “No-Self-Learning” Guarantee</i>	489
iii. <i>Data Tailoring—Minimization of Identifying Data</i>	489
iv. <i>Audit-Ready Logging</i>	489
v. <i>Informed Client Consent</i>	490
vi. <i>Organization of Firmwide AI Oversight Department</i>	490
VI. FROM THEORY TO PRACTICE: IMPLEMENTING FGIS WITH AUDITABLE GOVERNANCE ARTIFACTS	491
A. THE KIT OF GOVERNANCE ARTIFACTS	491
i. <i>AI Registry</i>	491
ii. <i>Use-Case Catalog</i>	492
iii. <i>FGIS Checklist (Per-Matter)</i>	494
iv. <i>Vendor Term Sheet</i>	494
v. <i>Incident Response Playbook</i>	495
VII. CURRENT FIRM ORGANIZATIONAL STRUCTURE FALLS SHORT	496
A. PROPOSAL: CREATION OF AI OVERSIGHT DEPARTMENT	496
i. <i>Structure & Staffing</i>	496
ii. <i>Core Responsibilities</i>	497
a. <i>Tool Vetting and Approval</i>	497
b. <i>Usage Protocols and Logging</i>	498

c. Compliance Audits	498
d. Training and Guidance	500
e. Client Disclosure and Consent	500
B. BENEFITS & RISK MITIGATION	500
C. CHALLENGES & IMPLEMENTATION CONSIDERATIONS	502
VIII. ALIGNING INCENTIVES: ECONOMICS OF ETHICAL AI	504
A. THE LAW AND ECONOMICS OF INFORMATION DISCLOSURE	506
B. MANDATORY, TRANSPARENT DISCLOSURE.....	508
C. THE PROHIBITION OF HOUR INFLATION.....	509
D. ALIGNMENT THROUGH ALTERNATIVE FEE ARRANGEMENTS	511
E. THE MEMORIALIZATION OF HUMAN VALUE	512
F. THE ECONOMIC TRANSITION: FROM BILLABLE HOURS TO VALUE-BASED PRICING	513
G. MARKET STRUCTURE CONSIDERATIONS	516
H. THE PATH FORWARD.....	518
IX. OBJECTIONS AND REBUTTALS	519
X. CONCLUSION.....	520
I. INTRODUCTION	

The legal profession stands at a technological precipice, confronted by the transformative power of general-purpose artificial intelligence (GPAI). These systems—capable of generating sophisticated text, analyzing complex documents, and synthesizing information—promise unprecedented gains in efficiency and productivity.¹ From a law and economics perspective, GPAI presents a classic case of technological disruption creating both efficiency gains and market failures.² The core economic challenge lies not in the technology itself, but in the misalignment of incentives it creates between lawyers, clients, and AI vendors.³ Despite these benefits, the debate over the

1. See, e.g., Jonathan H. Choi, Amy B. Monahan & Daniel Schwarcz, *Lawyering in the Age of Artificial Intelligence*, 109 MINN. L. REV. 147, 191–92 (2024) (noting in a blind graded study giving law students access to GPAI tools made them significantly more efficient at completing legal tasks without reducing the quality of their work).

2. See *infra* Sections II.A–C, VII.A–H.

3. See *infra* pp. 512–13.

ethical use of GPAI has focused almost entirely on AI “hallucinations,” meaning the generation of factually incorrect or entirely fabricated content.⁴

This narrow focus, while understandable in the wake of high-profile sanctions cases like *Mata v. Avianca, Inc.*,⁵ has created a dangerous “competence trap.”⁶ This trap is economically predictable: firms rationally invest resources in addressing the most visible, salient risk (hallucinations leading to sanctions) while systematically underinvesting in addressing structural, less visible risks (independence erosion and privilege waiver) that carry potentially greater expected costs.⁷ This represents a market failure driven by information asymmetry and bounded rationality.⁸ Law firms, believing that the primary ethical challenge is to verify the accuracy of AI-generated outputs, may invest heavily in review protocols while remaining blind to two structural threats that persist even when output is perfectly correct: (1) the erosion of a lawyer’s professional independence and (2) the systemic risk of waiving attorney-client privilege.

These twin threats reflect fundamental economic problems. The independence problem is a principal-agent issue compounded by third-party influence: clients (principals) hire lawyers (agents) to exercise independent judgment, but GPAI vendors introduce a third party whose economic interests systematically distort the agent’s incentives.⁹ The confidentiality problem reflects both information asymmetry (lawyers lack complete information about how vendors use data) and irreversible externalities (a privilege waiver creates costs borne by clients, not vendors, creating moral hazard).¹⁰

This Article argues that GPAI is not a neutral tool like a word processor or legal research database. Rather, it functions as an “opinionated co-author” governed by a commercial third-party vendor whose interests do not inherently align with those of the client.¹¹ In economic terms, GPAI vendors

4. See Matthew Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, 16 J. LEGAL ANALYSIS 64, 66 (2024) (explaining the “widespread occurrence of legal hallucinations”); see also Varun Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, 22 J. EMPIRICAL LEGAL STUDS. 216, 225 fig. 4 (2025) (finding that legal research companies’ GPAI systems hallucinate in seventeen percent to thirty-three percent of responses).

5. 678 F. Supp. 3d 443, 466 (S.D.N.Y. 2023) (ordering Rule 11 sanctions of \$5,000 after attorneys submitted non-existent judicial opinions with fake quotes and citations created by ChatGPT).

6. See *infra* p. 472.

7. See *infra* Section II.A.

8. See *infra* Part II.

9. See *infra* Section II.D.

10. See *infra* Sections II.A, II.B.

11. See *infra* Section II.A.

are profit-maximizing entities operating in a market where their revenue models create perverse incentives: they benefit from increased usage and data collection, even when such usage may compromise client interests.¹² This misalignment is not incidental but rather structural, embedded in the economic architecture of GPAI business models.¹³ This fundamental distinction gives rise to the twin crises of AI adoption. First, the design and commercial architecture of GPAI systems create subtle but powerful “influence vectors” that can steer a lawyer’s professional judgment.¹⁴ This dynamic presents a novel and direct challenge to the ethical mandate of professional independence, as codified by the American Bar Association (ABA) in Model Rule of Professional Conduct 5.4.¹⁵ Second, the structure of third-party GPAI systems, which routinely collect and reuse user inputs to train future models, creates a direct risk of disclosing confidential client information.¹⁶ This practice risks a catastrophic and likely irreversible waiver of attorney-client privilege under established third-party disclosure doctrines.

The *Mata* sanctions rightly clarified a baseline duty of verification under Model Rules of Professional Conduct 1.1 and 3.3.¹⁷ Early ethics opinions from the ABA,¹⁸ New York City Bar,¹⁹ D.C. Bar,²⁰ and Florida Bar²¹ have echoed this focus on competence and candor. While understandable, this emphasis has narrowed professional attention to output accuracy, obscuring deeper structural risks. A firm may diligently verify every AI-generated citation and assume its ethical obligations are met—even as its choice of vendor, driven by aggressive commercial practices, undermines strategic

12. See *infra* Section II.B.

13. See *infra* Sections II.A–C.

14. These include, for example, the algorithms that shape outputs, the training data that embed patterns and biases, and the pricing models that influence how lawyers use the system. See *infra* Sections II.A–C.

15. MODEL RULES OF PRO. CONDUCT r. 5.4 (A.B.A. 1983) (prohibiting lawyers from splitting legal fees with nonlawyers, forming partnerships with nonlawyers if any of the partnership’s activities involve the practice of law, and practicing law in a business that is owned at least in part by nonlawyers).

16. See *infra* Section III.A.

17. See *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443, 464 (S.D.N.Y. 2023) (“An inadequate or inattentive ‘inquiry’ may be unreasonable under the circumstances. But signing and filing that affirmation after making no ‘inquiry’ was an act of subjective bad faith.”); see also MODEL RULES OF PRO. CONDUCT r. 1.1 (A.B.A. 1983) (“A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.”); MODEL RULES OF PRO. CONDUCT r. 3.3 (A.B.A. 2025) (“A lawyer shall not knowingly make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer . . .”).

18. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512 (2024).

19. See N.Y.C. Bar Ass’n Comm. on Pro. Ethics Comm., Formal Op. 2024-5 (2024).

20. See D.C. Bar Ethics Op. 388 (2024).

21. See Fla. Bar Ethics Op., Op. 24-1 (2024).

independence and triggers privilege waivers over the very documents being reviewed.

To counter these structural threats, this Article proposes a restrictive, fiduciary-based standard for the use of GPAI in substantive legal work. The proposed “Fiduciary-Grade & Independence Safeguards” (FGIS) standard posits that law firms should presumptively prohibit the use of GPAI for tasks involving legal analysis, strategy, or the synthesis of confidential facts and law.²² A firm may only depart from this default prohibition if it can certify, with auditable proof, the implementation of two pillars of protection: (1) robust safeguards that neutralize vendor influence and preserve the lawyer’s independent professional judgment, and (2) fiduciary-grade confidentiality controls, including private deployment and contractual “no-training” guarantees, that eliminate the risk of privilege waiver. Where either pillar is missing, the only ethically defensible posture is non-use. This standard translates the abstract duties of the Model Rules into an operational governance framework, providing a clear path for the legal profession to absorb the power of automation without surrendering its foundational values.

II. THE EMERGING CONSENSUS: MAPPING THE NEW REGULATORY FLOOR

The rapid escalation of GPAI has spurred a wave of regulatory and judicial responses that establish a clear and consistent floor for ethical conduct regarding GPAI use in the legal industry. While no single authority has yet to mandate a comprehensive framework like this Article posits, a detailed analysis of the leading ethics opinions and court rules reveals a strong convergence on the core principles that together logically compel its adoption. As a result of the ABA’s national guidance, trailblazing state opinions, and bellwether judicial policies, a clear consensus in the legal industry has emerged demanding rigorous human oversight, stringent confidentiality protections, full client transparency, and fair billing practices as the law becomes immersed in GPAI use. While the emerging regulatory floor is currently framed as separate sets of discrete duties, these separate duties provide the necessary elements to establish a holistic and complete approach to harnessing GPAI while mitigating its inherent risks.

A. AMERICAN BAR ASSOCIATION FORMAL OPINION 512

The most significant guidance at the national level stems from ABA Standing Committee on Ethics and Professional Responsibility Formal

22. *See infra* p. 488–89.

Opinion 512, issued on July 29, 2024.²³ The opinion serves as a model for the mapping of GPAI-specific ethical use with several existing Model Rules of Professional Conduct.²⁴ By emphasizing the duty of competence required under Model Rule 1.1, the opinion clarifies that lawyers need not be experts on AI but merely need a “reasonable understanding of the capabilities and limitations” of the specific tools used in the course of representation.²⁵

ABA Formal Opinion 512 also emphasizes the relevance of Model Rule 5.3, which governs a lawyer’s responsibilities regarding nonlawyer assistance. As GPAI increasingly contributes to work that informs lawyers’ final deliverables, “lawyers must make reasonable efforts to ensure that the” use of those tools is compatible with their own professional obligations.²⁶

On the central topic of client confidentiality under Model Rule 1.6, the opinion emphasizes that lawyers using AI must understand the nature of the AI model, such as its terms of use and privacy policies, before utilizing it for client matters.²⁷ The failure for a lawyer to undertake this additional step of security before inputting client matters into an AI model could directly result in a breach of client confidentiality if, like most AI models, the model stores inputs and uses the information for self-learning purposes.²⁸ The failure to obtain informed written consent from the client prior to inputting their information into the AI model could result in a direct violation of Model Rule 1.6 by that lawyer.²⁹

Finally, the opinion covers the ever-important topic of fees under Model Rule 1.5, highlighting the key impacts AI use has on firms’ billing to clients.³⁰ Evidently stated in the opinion, firms must adjust their billing in accordance with their use of AI as it increases efficiency in delivering work product to clients.³¹ Failure by lawyers to accurately account for AI efficiency as a reflection on their billing would be counter to Model Rule 1.5.³² Proper tracking of AI use by lawyers and accurate reflection on bills sent to clients will ensure that the efficiency gains are not solely benefiting firms, while negatively impacting clients by either inflating costs or through inappropriate billing practices.³³

23. A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 1.

24. *See id.* at 2–13.

25. *See id.* at 2–3.

26. *See id.* at 10 (citing MODEL RULES OF PRO. CONDUCT r. 5.1 (2023)).

27. *See id.* at 6.

28. *See id.* at 6–7.

29. *See id.*

30. *See id.* at 11–14.

31. *See id.* at 12.

32. *See id.*

33. *See id.*

B. STATE BAR GUIDANCE

At the state level, the State Bar of California issued one of the earliest and most comprehensive AI-ethics guidance documents in November 2023.³⁴ Recognizing the magnitude AI implementation will have on the legal practice, the Bar urges lawyers and firms to develop a working understanding of AI systems and their risks to ensure their continued compliance with California's ethical standards.³⁵ Unlike the ABA's Formal Opinion 512, California's guidance explicitly emphasizes that the type of AI vendor used directly shapes the level of scrutiny and governance required.³⁶

Florida's Ethics Opinion 24-1, approved in January 2024, places Florida alongside California as one of the first states to issue guidance on AI governance.³⁷ Echoing the ABA and California, the Florida Bar emphasizes the ethics concern surrounding client confidentiality.³⁸ Crucial to its focus, the Florida Bar underscores the diligence required of lawyers when using "self-learning" models, and the heightened risk of confidentiality breaches associated with use of such models.³⁹ The opinion mandates informed written client consent before any AI use, requiring lawyers to disclose the specific model involved, the nature of its data practices, and the role it will play in the representation.⁴⁰

To further mitigate the risks associated with client confidentiality and AI use, the Florida Bar suggests that the use of an in-house AI model could be better than an outside AI model, as it does not openly turn over client

34. *See generally* COMM. ON PRO. RESP. & CONDUCT, STATE BAR CAL., PRACTICAL GUIDANCE FOR THE USE OF GENERATIVE ARTIFICIAL INTELLIGENCE IN THE PRACTICE OF LAW 1, <https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf> [<https://perma.cc/A3U8-9V9E>] (last visited Apr. 10, 2026) (guidance document issued by California's attorney ethics body, updating its 2023 AI guidance at the request of the California Supreme Court, addressing how existing professional responsibility obligations of competence, confidentiality, communication, and candor apply to lawyers' use of AI tools in legal practice).

35. *See id.*

36. *See id.* An open AI system "refers to artificial intelligence systems that can be used, examined, altered, and distributed for any purpose, without having to request permission." *See* Rina Diane Caballar & Cole Stryker, *What is Open-Source AI?*, IBM, <https://www.ibm.com/think/topics/open-source-ai> [<https://perma.cc/QZ2G-UDBL>] (last visited Apr. 10, 2026).

37. *See generally* Fla. Bar Ethics Op., Op. 24-1, *supra* note 21 (advisory opinion concluding that Florida lawyers may ethically use generative AI in legal practice provided they safeguard client confidentiality by researching a program's data retention and self-learning policies, maintain supervisory responsibility over AI-generated work product, avoid duplicative or inflated billing practices, and ensure that client-facing AI chatbots clearly identify themselves as non-human and comply with lawyer advertising rules).

38. *See id.* at 2–3.

39. *See id.* (noting "[u]se of 'self-learning' generative AI raises the possibility that a client's information may be stored within the program and revealed in response to future inquiries by third parties").

40. *See id.* at 6.

information to a third party.⁴¹ The opinion further sets forth an oversight governance outline for AI that firms could implement into their practice entirely or in pieces.⁴² While not every state has the same ethics rules as Florida, the ideas touched upon in the Florida Bar's opinion overlap drastically with concerns in every jurisdiction and provide strong support for the development and implementation of FGIS standards at firms to ensure the deployment of AI in the legal industry not only flourishes, but does so in an ethically compliant manner.

C. THE JUDICIARY'S WARNING: HUMAN OVERSIGHT IS NON-NEGOTIABLE

This regulatory guidance is powerfully reinforced by the judiciary's response to AI misuse and ethical missteps. The landmark case of *Mata v. Avianca, Inc.*, in which the Southern District Court of New York handed down sanctions when attorneys submitted a brief containing fabricated case law generated by ChatGPT, serves as a stern warning.⁴³ The order found that the attorneys had acted in subjective "bad faith" by submitting fictitious cases and standing by those non-existent cases as the trial proceeded.⁴⁴ The case established a clear judicial expectation that lawyers are gatekeepers who have a non-delegable duty under Federal Rule of Civil Procedure 11 to ensure the accuracy of their filings.⁴⁵ While acknowledging that there is "nothing inherently improper about using a reliable artificial intelligence tool for assistance," the court made clear that uncritical reliance on its outputs constitutes an abdication of professional responsibility.⁴⁶

The *Mata* sanctions and a growing number of similar cases have rendered precise human verification of AI output a non-negotiable for legal practitioners. Currently, there are trends among administrative bodies that govern the court's codifying of these AI governance oversights into formal rules. For example, most recently, the Judicial Branch of California adopted

41. *See id.* at 3 ("It should be noted that confidentiality concerns may be mitigated by use of an inhouse generative AI rather than an outside generative AI where the data is hosted and stored by a third-party.").

42. *See id.* at 4 (explaining the lawyer "must: (1) make reasonable efforts to ensure [the] law firm has policies to reasonably assure the conduct of a nonlawyer assistant, [such as AI,] is compatible with the lawyer's own professional obligations"; (2) review the work product of AI, verifying it for legal accuracy and sufficiency of all research performed on the matter, and ultimately use their professional judgment for the final deliverable to the client; (3) not delegate duties to AI that could constitute as the practice of law or requires the lawyer's personal judgment and participation).

43. *See* 678 F. Supp. 3d 443, 458, 464 (S.D.N.Y. 2023).

44. *See id.*

45. *See id.* at 448.

46. *See id.*

Standard of Judicial Administration 10.80, effective September 1, 2025.⁴⁷ The standard explicitly sets forth that judicial officers “[s]hould not enter confidential, personal identifying, or other nonpublic information into a public [G]AI system” and must take reasonable steps to verify the output as accurate when AI is used.⁴⁸ This policy, coming from the nation’s largest court system and home of Silicon Valley, is likely to serve as a national blueprint—strengthening the consensus of keeping confidential information out of open AI and requiring human validation of all AI output as of the utmost importance as AI use takes root in the legal industry.⁴⁹

Collectively, this evolving body of law and practical guidance forms a coherent and consistent regulatory floor. It establishes the foundational concepts that lawyers must understand the tools they use, govern their application, verify their outputs, protect client confidences, communicate transparently, and bill fairly for their use.

III. A THEORY OF VENDOR INFLUENCE: HOW GPAI THREATENS PROFESSIONAL INDEPENDENCE UNDER MODEL RULE 5.4

While the regulatory consensus effectively addresses the duties of competence and confidentiality, it largely overlooks a more insidious threat posed by the integration of GPAI into substantive legal work: the erosion of a lawyer’s professional independence. From a law and economics perspective, Model Rule 5.4 serves as a structural safeguard against agency costs.⁵⁰ The Rule recognizes that when nonlawyers have financial interests in legal practice, they face incentives to maximize their own welfare at the expense of client interests.⁵¹ This creates classic principal-agent problems where monitoring costs are high, and information asymmetries favor the agent.

Model Rule of Professional Conduct 5.4, titled “Professional Independence of a Lawyer,” is designed to prevent a nonlawyer’s financial interests or control from interfering with a lawyer’s independent professional judgment.⁵² The rule traditionally has been applied to prohibit fee-sharing

47. Cal. R. Ct. 10.80.

48. *Id.*

49. Natasha Allen, *California Courts Announce New AI Regulations*, NAT’L L. REV. (Aug. 18, 2025), <https://natlawreview.com/article/california-courts-announce-new-ai-regulations> [https://perma.cc/N4CL-7SEQ].

50. See Edward S. Adams & John H. Matheson, *The Big Four Accounting Law Firms: It’s Past Time for the Practice of Law to Get into the Game*, 92 TENN. L. REV. 1147, 1156–58 (2025).

51. See *id.*

52. MODEL RULES OF PRO. CONDUCT r. 5.4 (A.B.A. 2025) (providing in part that: “(d) A lawyer shall not practice with or in the form of a professional corporation or association authorized to practice law for a profit, if: (1) a nonlawyer owns any interest therein, except that a fiduciary

with nonlawyers, forming partnerships with nonlawyers if any of the partnership's activities consist of the practice of law, and practicing law in a business that is owned at least in part by nonlawyers.⁵³ However, the unique nature of GPAI creates a potent vector for nonlawyer influence that operates through algorithmic and economic pressures embedded in the technology itself.⁵⁴ These vectors represent a novel form of agency cost—one that is continuous, algorithmic, and difficult to monitor, thereby maximizing information asymmetry while minimizing the principal's ability to detect misalignment.⁵⁵

When a law firm deeply integrates a third-party GPAI for core legal analysis, it is not merely licensing a tool; it is effectively entering into a relationship with a nonlawyer entity that has a direct and ongoing influence on the creation of legal work product, functionally circumventing the spirit, if not the letter, of Model Rule 5.4.

A. THE ALGORITHMIC VECTOR: INFORMATION ASYMMETRY AND THE BLACK BOX PROBLEM

The first vector of influence is algorithmic. Unlike a traditional database that retrieves information based on explicit queries, a GPAI generates new content based on patterns in its training data and the architecture of its system.⁵⁶ This process is inherently opinionated.⁵⁷ In economic terms, this opacity creates severe information asymmetry between the lawyer (who observes only outputs) and the vendor (who controls the algorithm, training data, and system architecture).⁵⁸ This asymmetry is not merely technical but economic: the vendor possesses private information about bias vectors,

representative of the estate of a lawyer may hold the stock or interest of the lawyer for a reasonable time during administration; (2) a nonlawyer is a corporate director or officer thereof or occupies the position of similar responsibility in any form of association other than a corporation; or (3) a nonlawyer has the right to direct or control the professional judgment of a lawyer").

53. *Id.* Professor John Matheson and I argue that Model Rule 5.4 should be repealed to improve innovation and access to legal service. *See* Adams & Matheson, *supra* note 50.

54. *See infra* Sections II.A–C.

55. *See infra* Section II.D.

56. Bill Tomlinson, Donald J. Patterson & Andrew W. Torrance, *Turning Fake Data into Fake News: The AI Training Set as a Trojan Horse of Misinformation*, 60 S.D. L. REV. 641, 644 (2023) ("Training sets are vast collections of data used to teach algorithms how to process and interpret information."); *see id.* at 646 ("They provide the basis for AI models to learn patterns, relationships, and associations, which enable them to make predictions, recognize objects, generate text, and perform various other tasks. The quality of an AI model's performance is heavily influenced by the accuracy, representativeness, and comprehensiveness of the training set it is built upon.").

57. *See generally id.* (noting that AI outputs reflect the biases and inaccuracies embedded by human annotators during training, rendering them inherently subjective).

58. *See* FRANCESCO PARISI, *THE LANGUAGE OF LAW AND ECONOMICS: A DICTIONARY* 15 (2013) (explaining the significance of asymmetric information).

training data composition, and algorithmic design choices that systematically influence outputs in ways the lawyer cannot observe or verify.⁵⁹

The biases baked into its vast training data can all “nudge” a lawyer’s analysis in a particular direction.⁶⁰ For example, a model trained predominantly on sources favoring a certain legal interpretation may consistently generate arguments that reflect that bias.⁶¹ The lawyer, presented with a coherent and well-written output, may be unaware of the algorithmic thumb on the scale that produced it.⁶² This represents what economists call “hidden action” in principal-agent theory: the vendor’s choices in model design and training remain hidden from the lawyer, yet systematically affect outcomes.⁶³ The lawyer cannot efficiently monitor these choices because (1) monitoring costs are prohibitively high, (2) the vendor treats this information as proprietary, and (3) even with access, the technical complexity creates cognitive barriers to meaningful oversight.⁶⁴

This problem is compounded by the “black box” nature of many advanced models, which operate through layers of statistical correlations and internal representations that are not readily interpretable by humans.⁶⁵ In many cases, even the developers cannot fully trace or explain why a particular

59. See generally Tomlinson, Patterson & Torrance, *supra* note 56 (explaining training data and algorithmic design choices).

60. Sonia Gipson Rankin, *Mitigating Algorithmic Bias: Strategies for Addressing Discrimination in Data*, 20 SCITECH LAW. 26, 27 (2024) (“[A]lgorithmic bias is . . . inherent in the design of large language models.”). Because AI systems learn to identify patterns and make predictions based on the content of their training data, any inaccuracies, gaps, or representational skews embedded in that data will be reflected in the model’s outputs. See Tomlinson, Patterson & Torrance, *supra* note 56.

61. See Rankin, *supra* note 60, at 28. AI promises “personalization and efficiency, tailoring experiences and decisions to individual patterns and precedent.” *Id.* Yet these benefits carry significant tradeoffs. See *id.* As AI systems rely on the personalization of user data to refine their recommendations, they can contribute to the “narrowing of our information horizons” and the “diminishing [of] our collective awareness and understanding.” *Id.*

62. See *id.* at 29. Consider the ways AI systems influence daily life. See *id.* AI outputs subtly steer attention toward certain options while filtering out others, often without our awareness. See *id.* In doing so, AI limits choice and optionality, shaping not only the decisions we make but also the set of decisions available to us in the first place. *Id.*

63. See ELMAR WOLFSTETTER, TOPICS IN MICROECONOMICS: INDUSTRIAL ORGANIZATION, AUCTIONS, AND INCENTIVES 278 (1999); Rankin, *supra* note 60, at 27 (“[S]pecific design decisions . . . can lead to biased outcomes [and] the challenges posed by AI’s ‘black box’ nature are particularly acute in legal settings, where the rationale behind decisions must be explainable and justifiable.”).

64. See Michael C. Jensen & William H. Meckling, *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*, 3 J. FIN. ECON. 305, 308 (1976) (defining agency costs to include monitoring expenditures by the principal); Rankin, *supra* note 60, at 27; Vikas Hassija et al., *Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence*, 16 COGNITIVE COMPUTATION 45, 47 (2024).

65. Rankin, *supra* note 60, at 27 (“The term ‘black box’ in AI refers to models where the internal workings and decision-making processes are difficult for humans to understand or even discern.”).

output was produced in response to a specific prompt.⁶⁶ This represents an extreme form of information asymmetry where even the vendor lacks transparency about its own system's decision-making process.⁶⁷ In economic terms, this creates a "double principal-agent problem": lawyers cannot monitor vendors, and vendors cannot fully monitor their own AI systems.⁶⁸ The result is a cascading failure of accountability where no party can effectively ensure alignment with client interests.

This opacity makes it difficult to scrutinize or verify the reasoning behind AI-generated recommendations, effectively requiring the lawyer to defer to the system's unseen logic.⁶⁹ As a result, the lawyer must place a meaningful degree of trust in the vendor's model, which can erode the exercise of truly independent professional judgment. In law and economics terminology, this constitutes a form of "bounded rationality": lawyers, facing cognitive and information constraints, must rely on heuristics and trust mechanisms that the market structure systematically exploits.⁷⁰ The vendor's business model depends on lawyers giving this trust, even when such trust may not be economically justified by verifiable performance metrics.

B. THE ECONOMIC VECTOR: PRICING MODELS AS INFLUENCE MECHANISMS

The second vector is economic. Many GPAI vendors employ a per-token or metered pricing model, where the cost of using the service is directly proportional to the amount of text processed and generated.⁷¹ From a microeconomic perspective, this pricing structure creates a direct marginal cost for each unit of legal work, fundamentally altering the lawyer's

66. See generally Hassija et al., *supra* note 64, at 47 ("These models make predictions based on input data, but the decision-making process and reasoning behind the predictions are not transparent to the user. This lack of transparency makes it strenuous for users to understand the model's behavior, detect potential biases or errors, or hold the model accountable for its decisions.").

67. See *id.*

68. See PARISI, *supra* note 58, at 5 (explaining "Agency Problems").

69. See Rankin, *supra* note 60, at 27 ("AI's 'black box' nature are particularly acute in legal settings, where the rationale behind decisions must be explainable and justifiable.").

70. See PARISI, *supra* note 58, at 31 (defining the term "Bounded Rationality" in the field of law and economics as the notion that humans have limited cognitive abilities, computational skills, and flawed memory).

71. See, e.g., *Pricing*, CLAUDE API DOCS, <https://docs.claude.com/en/docs/about-claude/pricing> [<https://perma.cc/UC9U-3TQT>] (last visited Apr. 11, 2026); *API Pricing*, OPENAI, <https://openai.com/api/pricing/> [<https://perma.cc/D839-2T3B>] (last visited Oct. 7, 2025) (A per-token pricing model charges users for each small unit of text processed or generated, while a metered pricing model bills users based on total consumption, meaning costs scale directly with how much they use the service.).

optimization problem.⁷² Instead of maximizing client welfare subject to professional duty constraints, the lawyer now faces an additional constraint: minimizing vendor costs. This transforms the objective function in ways that systematically bias decision-making away from thoroughness toward efficiency.

This structure creates constant pressure to minimize costs, which can subtly but surely influence legal strategy. A lawyer might be unconsciously incentivized to shorten a research query, limit the number of documents reviewed in a summarization task, or accept a “good enough” first draft of an argument to avoid incurring additional fees. These are not hypothetical concerns but predictable responses to price signals. Economic theory tells us that rational actors respond to marginal costs: when deeper research imposes marginal costs (additional tokens), lawyers face incentives to underinvest in research relative to the client-optimal level.⁷³ This creates a negative externality where cost savings accrue to the firm (or are partially shared with clients) while the costs of reduced thoroughness are borne by clients in the form of inferior legal work product.⁷⁴

This economic calculus, driven by the vendor’s profit motive, can directly conflict with the lawyer’s duty to provide thorough and zealous representation.⁷⁵ The decision of how deep to dig into research or how many drafts to refine should be governed by the needs of the client’s matter, not by the per-token cost of the firm’s GPAI vendor. This represents a classic case of third-party influence distorting agency relationships.⁷⁶ The vendor’s pricing model effectively becomes a “shadow principal” that influences the lawyer’s behavior without any fiduciary obligation to the client. The vendor maximizes its profits through usage-based pricing, while clients bear the agency costs of lawyers’ systematically skewed incentives.

72. See STEVEN A. GREENLAW ET AL., PRINCIPLES OF ECONOMICS 3e § 2.1, OPENSTAX (2022) (explaining that economic choices are constrained by costs, that “every choice has an opportunity cost,” and that marginal analysis examines the costs and benefits of choosing slightly more or less of a good or service); *Pricing*, *supra* note 71; *API Pricing*, *supra* note 71.

73. See GREENLAW ET AL., *supra* note 72, § 2.1 (explaining that marginal analysis requires comparing how costs and benefits change from one option to another and that a rational consumer purchases additional units only while marginal utility exceeds opportunity cost); see also Jensen & Meckling, *supra* note 64, at 308 (explaining that where an agent exercises delegated decision-making authority, the agent may not always act in the principal’s best interest).

74. See GREENLAW ET AL., *supra* note 72, § 12.1 (explaining that an externality occurs when an exchange affects a third-party outside the transaction and that externalities can produce market failure when the actor does not account for the broader social cost of its behavior).

75. See MODEL RULES OF PRO. CONDUCT r. 1.3 (A.B.A. 2025) (“A lawyer shall act with reasonable diligence and promptness in representing a client.”).

76. See generally RESTATEMENT (THIRD) OF AGENCY: AGENTS WITH MULTIPLE PRINCIPALS § 3.41 cmt. c (A.L.I. 2006) (explaining ambiguous agency relationships).

C. SOPHISTICATED PRICING MODELS: HIDDEN INFLUENCE PERSISTS

Some may argue that this framing overstates the immediacy of constraint. After all, most GPAI platforms abstract per-token pricing into fixed-fee, tiered, or enterprise licensing arrangements that shield lawyers from direct usage costs.⁷⁷ Yes—these models typically include generous caps or unlimited access, reducing the likelihood of overt economic constraint on daily workflows.⁷⁸ However, from a behavioral economics perspective, pricing architecture shapes decision-making even when direct costs are abstracted. The phenomenon, economists call it “mental accounting,” which means that lawyers internalize cost consciousness even under flat-rate pricing, particularly when firms emphasize cost control in performance evaluations or when firms choose vendors partially based on pricing structures.⁷⁹

Nevertheless, the influence of pricing architecture persists in subtler ways. Procurement teams may favor lower-cost vendors, internal budgeting may discourage excessive usage, and lawyers may internalize efficiency pressures that shape how deeply they research, how many documents they summarize, or how thoroughly they refine drafts. These dynamics reflect what economists call “organizational incentive structures”: even when individual lawyers face no direct usage costs, firm-level incentives cascade down through performance metrics, cultural norms, and resource allocation decisions. The economic pressure point shifts from individual transactions to organizational decision-making, but the fundamental misalignment persists.

When vendor pricing models or internal cost controls begin to influence the substance of representation, they raise the kind of structural concerns Rule 5.4 was designed to prevent: nonlawyer financial interests shaping the contours of legal service.⁸⁰ In law and economics terms, this represents regulatory capture at the micro level: a professional conduct rule designed to prevent third-party economic influence becomes ineffective because the mechanism of influence (algorithmic and pricing architecture) operates

77. See, e.g., *Pricing*, OPENAI, <https://openai.com/business/chatgpt-pricing/> [<https://perma.cc/BK74-MWLK>] (last visited Oct. 7, 2025) (monthly fee for individual, business, and enterprise plans); *Flexible Copilot Plans for Every Organization*, MICROSOFT 365 COPILOT, <https://www.microsoft.com/en-us/microsoft-365-copilot/pricing> [<https://perma.cc/79G5-TCDQ>] (last visited Apr. 11, 2026).

78. See generally *Pricing*, *supra* note 71; *Flexible Copilot Plans for Every Organization*, *supra* note 77.

79. See, e.g., Richard H. Thaler, *Mental Accounting Matters*, 12 J. BEHAV. DECISION MAKING 183, 184 (1999).

80. MODEL RULES OF PRO. CONDUCT r. 5.4 cmt. 2 (A.B.A. 1983).

below the threshold of traditional regulatory detection.⁸¹ The market has evolved a form of third-party influence that exploits gaps in existing regulatory frameworks.

The debate over reforming Rule 5.4 has historically focused on allowing nonlawyer investment in law firms to spur innovation and improve access to justice.⁸² Proponents argue that capital from outside the legal profession is necessary to build more efficient and technologically advanced service delivery models.⁸³ However, the deep integration of a single GPAI vendor for substantive work creates a different and arguably more dangerous form of nonlawyer entanglement. It is not an investment in the firm's equity, but a continuous operational influence on the firm's core intellectual product. The vendor's commercial goal is pursued through the design of its algorithms and the structure of its pricing. These mechanisms act as a constant, subtle force on the lawyer's day-to-day work. In this arrangement, the vendor becomes a *de facto* partner with a direct financial stake in how legal services are rendered, creating the very conflict between a nonlawyer's profit motive and a lawyer's professional duty that Rule 5.4 was designed to prohibit.

D. THE PRINCIPAL-AGENT PROBLEM REDEFINED

The GPAI context creates a complex, multi-layered principal-agent problem that law and economics scholarship has not fully addressed. In the traditional lawyer-client relationship, we have a simple principal-agent model: the client (principal) retains the lawyer (agent) to exercise professional judgment on the client's behalf.⁸⁴ Agency costs arise from: (1) information asymmetry (lawyers know more about legal strategy than clients), (2) divergent incentives (lawyers may prioritize their own interests), and (3) monitoring difficulties (clients cannot perfectly observe lawyer effort or quality).

GPAI introduces a third party (the vendor), creating a triangular relationship that fundamentally alters the incentive structure. The vendor is not merely a tool provider but an active participant whose economic interests shape outcomes. This creates three distinct agency relationships:

81. See PARISI, *supra* note 58, at 37 ("Regulatory Capture" or "Capture Theory" explains the phenomenon whereby regulation becomes ineffective at serving its intended public purpose because the regulated entity exerts influence over the regulatory process through mechanisms that escape the regulator's detection or enforcement capabilities).

82. See Adams & Matheson, *supra* note 50.

83. See, e.g., Edward S. Adams & John H. Matheson, *Law Firms on the Big Board?: A Proposal for Nonlawyer Investment in Law Firms*, 86 CAL. L. REV. 1, 30 (1998) (arguing that opening law firms to capital markets would enable investment in new technologies, increase competition, and improve access to legal services).

84. See PARISI, *supra* note 58, at 5 ("Agency Problems").

1. Client-Lawyer Agency: The primary fiduciary relationship where traditional agency costs already exist.
2. Lawyer-Vendor Relationship: Not a formal agency relationship but a contractual arrangement where the vendor's business model creates de facto influence over lawyer behavior.
3. Client-Vendor Indirect Relationship: Clients bear the costs of vendor influence on lawyers but have no direct contractual relationship with, or oversight over, the vendor.

This structure causes information asymmetries and incentive misalignments to compound across multiple relationships. The client cannot effectively monitor the lawyer's use of GPAI, the lawyer cannot effectively monitor the vendor's algorithmic choices, and the vendor has no fiduciary duty to either the lawyer or the client. Each layer of the relationship introduces additional agency costs that accumulate to create systematic misalignment with client interests.

Moreover, this structure creates moral hazard at multiple levels.⁸⁵ Vendors can profit from business models that compromise client interests (such as data monetization) because they bear no direct liability to clients. Lawyers may over-rely on GPAI to reduce their own effort costs, knowing that output verification is imperfect and errors may not be detected. Clients, lacking technical expertise and information about GPAI use, cannot efficiently contract around these problems *ex ante* or detect them *ex post*.

The economic solution to principal-agent problems typically involves some combination of monitoring, incentive alignment, and bonding mechanisms.⁸⁶ However, the GPAI context makes all three difficult: monitoring is costly and technically complex, vendor business models undermine incentive alignment, and bonding (such as warranties or insurance) is largely absent from the GPAI market. This suggests that market-based solutions are insufficient, and regulatory intervention—such as the FGIS framework proposed in this Article—becomes economically necessary to address systematic market failure.

IV. A THEORY OF FIDUCIARY-GRADE CONFIDENTIALITY: PRIVILEGE, WAIVER, AND THE THIRD-PARTY AI DOCTRINE

The second structural threat posed by GPAI, the risk to client confidentiality and attorney-client privilege, is widely recognized yet poorly understood in its full legal implications. This widespread but superficial recognition creates a dangerous complacency: lawyers believe they are

85. See PARISI, *supra* note 58, at 187–88 (“Moral Hazard”).

86. See *id.* at 5.

“being careful” with AI while inadvertently waiving privilege through structural vulnerabilities embedded in the technology itself.

Model Rule of Professional Conduct 1.6 requires lawyers to protect “information relating to the representation of a client” and to make “reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to,” such information.⁸⁷ Attorney-client privilege, though related, is a distinct evidentiary doctrine protecting confidential attorney-client communications from compelled disclosure in legal proceedings.⁸⁸ Both protections share a critical vulnerability: they can be waived through voluntary disclosure of confidential information to a third party.⁸⁹ This fundamental principle creates a direct collision with the operational architecture of most third-party GPAI systems.

This section demonstrates why standard GPAI use creates unacceptable privilege waiver risk, why protections developed for other technologies are insufficient, and why only “fiduciary-grade confidentiality,” private deployment, and contractual no-training guarantees can satisfy a lawyer’s duties. The confidentiality pillar of FGIS is not optional risk mitigation; it is the legal floor below which any AI deployment becomes professionally negligent.

A. PRIVILEGE: THE FRAGILITY OF ATTORNEY-CLIENT PROTECTION

Attorney-client privilege stands as one of the oldest protections in common law, encouraging full and frank communication by ensuring confidential communications made for the purpose of legal advice will not be disclosed in litigation.⁹⁰ The prevailing modern rationale is utilitarian.⁹¹ Absent privilege protection, clients would withhold critical information, impairing their attorneys’ ability to provide competent counsel.⁹² When properly maintained, the privilege is robust. However, the privilege is inherently fragile. It exists only so long as confidentiality is maintained from third parties.⁹³ Any voluntary disclosure of privileged communication to a

87. MODEL RULES OF PRO. CONDUCT r. 1.6(a), (c) (A.B.A. 2025).

88. *See, e.g.*, RESTATEMENT (THIRD) OF THE LAW GOVERNING LAWYERS § 68 (A.L.I. 2000).

89. *See* 2 CHRISTOPHER B. MUELLER & LAIRD C. KIRKPATRICK, FED. EVID. § 5.33 (4th ed. 2025).

90. *See id.* § 5.13.

91. *Id.*

92. *See id.* (“Without a privilege, the argument runs, clients would be deterred from seeking legal assistance or discouraged from telling their lawyers everything that might be important. Attorneys would not get all the information they need to represent their clients effectively.”).

93. *See id.* § 5.18.

third party destroys the privilege entirely.⁹⁴ This destruction is categorical, not a matter of degree.⁹⁵ Once the privilege is waived through disclosure, it cannot be selectively reasserted.⁹⁶ The information loses protected status and becomes subject to discovery and admissible as evidence.⁹⁷

This fragility creates special responsibilities for lawyers. Model Rule of Professional Conduct 1.6(c) imposes an affirmative obligation: lawyers must “make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.”⁹⁸ The “reasonable efforts” standard evolves with technology. Comment 18 to Rule 1.6 requires lawyers to “act competently to safeguard information relating to the representation of a client against unauthorized access by third parties.”⁹⁹

The emergence of GPAI systems that actively process, analyze, and potentially retain user inputs represents precisely the kind of technological development that triggers a lawyer’s duty to reassess confidentiality protections and implement new safeguards. For GPAI use, lawyers must deploy technology in a manner that: (1) makes “reasonable efforts” to prevent unauthorized disclosure, satisfying the ethical duty; and (2) avoids any disclosure to third parties that would trigger privilege waiver under evidence law. Meeting this dual imperative requires structural controls that eliminate the third-party disclosure problem at its root.

B. WAIVER: HOW THIRD-PARTY AI DISCLOSURE DESTROYS PRIVILEGE IRREVOCABLY

Privilege waiver doctrine is straightforward: voluntary disclosure of privileged information to a third party destroys the privilege. Courts consistently hold that once a party shares privileged communications with an outsider, someone outside the attorney-client relationship and not necessary to facilitate legal services, the confidential nature is compromised and privilege is lost.¹⁰⁰

94. *See id.* § 5.33 (“A client who voluntarily discloses the content of communications covered by the attorney-client privilege waives the privilege.”).

95. *See generally id.* (“[D]isclosure waives the lawyer-client privilege even if the client speaks without intentionally or purposefully relinquishing his privilege claim, so long as he intentionally and purposefully reveals the substance of a confidential communication.”).

96. *See id.*

97. *See id.*

98. MODEL RULES OF PRO. CONDUCT r. 1.6(c) (A.B.A. 2025).

99. MODEL RULES OF PRO. CONDUCT r. 1.6(c) cmt. 18 (A.B.A. 2025).

100. *See, e.g., In re Grand Jury Investigation of Ocean Transp.*, 604 F.2d 672, 675 (D.C. Cir. 1979) (noting a waiver may occur even without intent to waive the privilege (citing *Underwater Storage, Inc. v. U.S. Rubber Co.*, 314 F. Supp. 546, 549 (D.D.C. 1970)); *see also* 2 MUELLER & KIRKPATRICK, *supra* note 89, § 5.33.

When a lawyer uses third-party GPAI, they submit prompts containing confidential information, like document excerpts, case details, and strategy questions. This information is transmitted to vendor-owned servers where the system actively processes it: parsing, tokenizing, contextualizing, and running it through computational architecture.¹⁰¹ Depending on vendor terms, information may be stored for service improvement, quality control, safety monitoring, or model training.¹⁰² In self-learning systems, user inputs feed directly into training pipelines.¹⁰³ This is not passive data storage; instead, the AI vendor engages with confidential information to serve its own business objectives, including improving its product, monitoring liability, and extracting value from interactions.¹⁰⁴ Under these circumstances, the vendor is unambiguously a “third party” for privilege purposes, and submission constitutes “disclosure.”

Review of actual vendor terms confirms this. Many AI services explicitly state user inputs may be stored, logged, used for training, or shared with third-party service providers.¹⁰⁵ Even when vendors claim they do not use customer data for “general” model training, they reserve rights for service improvement, safety and security, or troubleshooting.¹⁰⁶ This retention transforms the vendor from a passive service provider into an interested third party with legitimate business purposes for accessing and using data.¹⁰⁷ Multiple bar associations have recognized this structural problem. Florida Bar Ethics Opinion 24-1 noted self-learning models may store a client’s information and may reveal it “in response to future inquiries by third parties,” explicitly suggesting “confidentiality concerns may be mitigated by use of an inhouse generative AI rather than an outside generative AI where the data is hosted and stored by a third-party.”¹⁰⁸ New York City Bar Formal

101. See generally Mindy Nunez Duffourc, Sara Gerke & Konrad Kollnig, *Privacy of Personal Data in the Generative AI Data Lifecycle*, 13 N.Y.U. J. INTELL. PROP. & ENT. L. 219, 244–49 (2024) (explaining how personal data is transmitted for processing, training, and model improvement through computational architectures).

102. See *id.* at 248–49.

103. See, e.g., *Privacy Policy*, OPENAI, <https://openai.com/policies/row-privacy-policy/> [<https://perma.cc/DZS5-DYKP>] (“[W]e may use Content you provide us to improve our Services, for example to train the models that power ChatGPT.”).

104. See *id.* (explaining the AI vendor uses personal data to (1) improve and develop its services and conduct research; (2) monitor any conduct on the platform to prevent fraud, illegal activity, or misuses; and (3) provide, analyze, and maintain its services to respond to questions in a personalized manner).

105. See *id.*

106. See *id.* (including “legitimate business purposes such as resolving legal disputes, safety and security reasons, or complying with our legal obligations” as reasons for retaining personal data).

107. See *id.*

108. Fla. Bar Ethics Op., Op. 24-1, *supra* note 21, at 2–3.

Opinion 2024-5 stated unequivocally: absent client consent, “a lawyer must not input confidential client information into any Generative AI system that will share the inputted confidential information with third parties.”¹⁰⁹ ABA Formal Opinion 512 requires lawyers to obtain informed client consent before using self-learning GPAI tools that might disclose confidential information.¹¹⁰

Federal Rule of Evidence 502(b) provides “clawback” for inadvertently disclosed privileged information if: (1) disclosure was inadvertent; (2) the privilege holder took reasonable measures to prevent disclosure; and (3) the holder promptly rectified the error.¹¹¹ This protective framework assumes disclosure was truly inadvertent; an unintended mistake despite reasonable care.

Inputting client data into public GPAI platforms is categorically different. It is not an accident. The lawyer decides to use AI, composes a prompt including confidential information, and clicks “submit.” Before this, the lawyer has agreed to vendor terms typically stating user inputs may be stored, processed, and used for training.¹¹²

The vendor’s subsequent use of client information in accordance with these terms cannot be characterized as “inadvertent” or “unauthorized.” It is the foreseeable, predictable, and contractually permitted consequence of the lawyer’s intentional act. The lawyer has knowingly placed confidential information into a system whose terms explicitly authorize the vendor to compromise confidentiality.¹¹³ This “knowing disclosure” problem fundamentally undermines Federal Rule of Evidence 502(b) protection. The first element cannot be satisfied when disclosure was the intended result. Moreover, if vendor terms disclosed inputs would be used for training and the lawyer failed to read those terms or proceeded anyway, the lawyer cannot credibly argue they “took reasonable steps to prevent disclosure.”¹¹⁴ Such negligence precludes relief under doctrines designed to protect parties who made reasonable efforts but still suffered accidental breaches.¹¹⁵

109. N.Y.C. Bar Ass’n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19, at 2.

110. A.B.A. Comm. on Ethics & Pro. Resp., Op. 512, *supra* note 18, at 6.

111. FED. R. EVID. 502(b).

112. *See, e.g., Privacy Policy*, *supra* note 103.

113. *See, e.g., D.C. Bar Ethics Op. 388*, *supra* note 20 (discussing a lawyer’s duty to prevent disclosure of client confidences to third-party GPAI providers that may access, store, or reuse the lawyer’s inputs or outputs in ways exposing client information).

114. *Id.* (“Lawyers . . . should understand that many GAI products currently on the market are specifically designed to collect and use information received from users—which may include client confidences and secrets—for the GAI’s own training and for transmission to future users of the technology.”); FED. R. EVID. 502(b).

115. *See generally* 2 MUELLER & KIRKPATRICK, *supra* note 89, § 5.34 (commenting on inadvertent or involuntary disclosure and waiver under Rule 502).

Once confidential information is input into a public GPAI system with self-learning capabilities, the resulting privilege waiver is likely permanent and irreversible. Unlike a misdirected email that can be recalled or a document clawed back through court order, data ingested into AI training datasets cannot be unseen or extracted.¹¹⁶

When confidential information trains a machine learning model, it becomes embedded in the model's neural network statistical parameters.¹¹⁷ The model does not store information as discrete, retrievable files.¹¹⁸ Instead, information influences learned patterns in a distributed manner pervading the entire system.¹¹⁹ While original input text might be stored separately and deleted, the influence that input had on model training cannot be precisely identified, isolated, or reversed.¹²⁰ Even if a vendor agrees to remove client data and retrain the model without it, these remedial measures are technically imperfect.¹²¹ The model may have been deployed in multiple versions or used as a foundation for derivative models, and there are complex technical issues about whether, if data can be located, it could be deleted "without requiring the destruction of the model itself."¹²² "Clawing back" data from a trained machine learning model will, technically speaking, be a challenge.¹²³ This irreversibility means confidentiality risks cannot be managed through ex post remediation—they must be prevented ex ante through structural controls, eliminating disclosure possibility in the first instance.¹²⁴

Model Rule 1.6(a) permits lawyers to disclose confidential information if "the client gives informed consent."¹²⁵ While such consent may satisfy the

116. JENNIFER KING & CAROLINE MEINHARDT, *Rethinking Privacy in the AI Era: Policy Provocations for a Data-Centric World*, STAN. INST. FOR HUM. CTR. AI 19 (2024), <https://hai.stanford.edu/assets/files/2024-02/White-Paper-Rethinking-Privacy-AI-Era.pdf> [https://perma.cc/39QS-HWA2] ("When personal data is included in the training dataset, research has demonstrated that these systems can memorize the data and then expose it to other users as part of the outputs.").

117. See Tomlinson, Patterson & Torrance, *supra* note 56, at 646–49.

118. See *id.* at 647 ("Once the data has been collected, it must be cleaned and preprocessed. This step involves removing duplicate or irrelevant entries, handling missing or incomplete data, and converting the data into a suitable format for the AI model to process.").

119. See *id.* at 646 ("[T]raining sets . . . serve as the foundational input for teaching machine learning algorithms.").

120. See Hassija et al., *supra* note 64, at 47.

121. See generally KING & MEINHARDT, *supra* note 116, at 39 ("There are also complex technical issues around (1) the question of whether one can request the deletion of data from a trained model without requiring the destruction of the model itself; and (2) determining whether one's personal data was used to train (or retrain) a model." (citation modified)).

122. See *id.*

123. See *id.*

124. See, e.g., *id.* at 40 (arguing for new AI data governance supply chain structures that prioritize greater transparency to the entire data development process used to train AI models).

125. MODEL RULES OF PRO. CONDUCT r. 1.6 (A.B.A. 2025).

ethical duty of confidentiality, it does not preserve attorney-client privilege. Privilege is designed to encourage full communication by ensuring such communications will not be disclosed in litigation.¹²⁶ Once privileged information is voluntarily disclosed to a third party, confidentiality is destroyed, and privilege is lost; regardless of whether the client consented.¹²⁷ Courts consistently hold that parties cannot selectively waive privilege by disclosing information to some third parties while maintaining confidentiality against others.¹²⁸ Client consent may shield the lawyer from ethics violations, but does not prevent adverse parties from compelling production of waived information or using it as evidence.¹²⁹

C. THE THIRD-PARTY AI DOCTRINE: FIDUCIARY-GRADE SAFEGUARDS AS THE ONLY DEFENSE

When law firms adopted cloud storage and Software as a Service (SaaS) platforms in the early 2000s, concerns arose about storing client data on third-party servers.¹³⁰ Consensus emerged that cloud services could be ethically permissible provided lawyers exercised reasonable care: investigating vendor security measures, reviewing terms of service, obtaining contractual confidentiality assurances, and ensuring regulatory compliance.¹³¹

Cloud computing's ethical acceptability rested on characterizing the vendor's role as a passive custodian.¹³² Cloud providers offer purely infrastructural services—secure storage and communication channels.¹³³ Data remained segregated, encrypted, and accessible only to authorized firm personnel.¹³⁴ The vendor's business model was based on charging for storage

126. See 2 MUELLER & KIRKPATRICK, *supra* note 89, § 5.13 (“By encouraging more open attorney-client communication, the privilege helps lawyers provide the kind of legal advice that is critical if clients are to assert their rights effectively, defend themselves properly, and avoid litigation in a complex regulated society.”).

127. See *id.* § 5.33.

128. See *id.* (“A client generally cannot make ‘selective’ or ‘limited’ disclosure to outsiders without waiving the privilege. The generally stated reason behind this principle is that the protection of the privilege already gives clients what they need in order to obtain legal representation, and allowing them to share with outsiders what is covered by the privilege would invite manipulation.”).

129. See *generally id.* (“Once the privilege is waived, the former holder generally cannot reclaim the privilege.”).

130. See Stuart Pardau & Blake Edwards, *The Ethical Implications of Cloud Computing for Lawyers*, 31 J. MARSHALL J. INFO. TECH. & PRIV. L. 69, 70, 74 (2014).

131. See *id.* at 71 n.10 (listing the state bar ethics opinions concluding that outsourcing storage of client files is permissible if the firm takes reasonable steps to ensure data is protected).

132. See *id.* at 73–74 (defining cloud computing as a storage system allowing files to be accessed from any computer and shared with others).

133. See *id.*

134. See *id.* at 78–79, 79 n.36 (listing thirteen broad factors to determine what constitutes “reasonable efforts” to ensure confidentiality: “(1) backing up data, (2) installing a firewall, (3) limiting information provided to others, (4) avoiding inadvertent disclosures, (5) verifying the

capacity, not extracting insights from customer data. Provided an attorney uses “reasonable care,” state bar ethics authorities concluded cloud services did not constitute privilege-waiving “disclosure.”¹³⁵

GPAI vendors are fundamentally different. A generative AI system’s entire value proposition lies precisely in active, substantive processing and analysis of input data.¹³⁶ When a lawyer submits a prompt, the system does not simply store information; instead, it immediately engages substantively: parsing input, contextualizing by relating it to learned patterns, and generating output based on probabilistic relationships.¹³⁷ In many systems, processing does not end with response generation. User inputs may be logged for histories, debugging, human review, and (most significantly) to serve as training data for future versions.¹³⁸ The system extracts meaning, makes predictions, and generates novel output by synthesizing input with learned knowledge.¹³⁹ This is analytical work performed on client information substance; precisely the third-party engagement the law treats as disclosure.

Moreover, the vendor’s business model fundamentally depends on this active processing, as continuous model improvement essential to competitive advantage depends heavily on exposure to diverse data, and user inputs are among the most valuable sources.¹⁴⁰ The vendor has direct commercial interest in retaining and analyzing user inputs.

This fundamental distinction, active processing versus passive custodianship, means the reasonable precautions framework developed for cloud computing, while still necessary, is not sufficient for GPAI.

V. THE FGIS STANDARD: AN AUDITABLE FRAMEWORK FOR FIDUCIARY-GRADE AND INDEPENDENCE SAFEGUARDS

With the theoretical risks associated with compromised independence and potential for waived client confidentiality, the creation and

identity of individuals to whom the attorney provides information, (6) refusing to disclose confidential information to unauthorized individuals without client permission, (7) encrypting confidential data, (8) ‘implementing electronic audit trail procedures to monitor who is accessing the data,’ (9) creating plans to address security breaches, (10) vetting service providers and service agreements, (11) training employees, (12) storing copy of digital data onsite, and (13) having an alternate way to connect to the internet” (citing Pa. Bar Ass’n Comm. on Legal Ethics & Pro. Resp., Op. 2011-200 (2011)).

135. *See id.* at 76.

136. *See* Tomlinson, Patterson & Torrance, *supra* note 56, at 646–47.

137. *See id.*

138. *See id.* at 647–48 (discussing data annotation and validation).

139. *See id.* at 647.

140. *See generally* Tal Roded & Pater Slattery, *What Drives Progress in AI? Trends in Data*, MIT FUTURETECH (Mar. 19, 2024), <https://futuretech.mit.edu/news/what-drives-progress-in-ai-trends-in-data> [<https://perma.cc/A36Y-33C6>] (explaining that the development of AI models requires increasingly large volumes of data).

implementation of a practical, prescriptive solution is a must. The Fiduciary-Grade & Independence Safeguards (FGIS) standard posited by this Article is an auditable framework firms can implement in an effort to turn these often-abstract duties imposed under the Model Rules into concrete operational oversights of GPAI in substantive legal work. Two pillars comprise the foundation of the FGIS framework; both presenting actionable requirements allowing firms to trace the iterations of GPAI use throughout their practice. While some may view the FGIS framework as a significant hinderance on AI usage, the reality is that FGIS is a rigorous checks and balance system put in place to ensure a firm is in full compliance with all traditional duties imposed on legal professionals when deploying GPAI in their substantive legal work.

A. PILLAR A: INDEPENDENCE SAFEGUARDS

This pillar is tailored to neutralize the “influence vectors” discussed in the previous sections and ensure that firms comply with Rules 5.3 and 5.4 of the Model Rules. Such treatment would require the firm to establish reasonable oversight mechanisms, ensuring that the nonlawyer’s conduct meets the professional standards expected of a licensed attorney.

i. Treatment of AI Vendor as Supervised Nonlawyer Assistant (Rule 5.3)

To start, the AI vendor must be treated as a nonlawyer assistant and held to the same standards as those outlined in Model Rule 5.3.¹⁴¹ Such treatment would obligate the firm to establish reasonable oversight mechanisms ensuring that the nonlawyer assistant’s conduct aligns with the professional standards expected of a licensed attorney.¹⁴² As set forth in ABA Formal Opinion 512, these duties are critical for the successful integration of AI in law firms.¹⁴³ In effect, law firms must impose on vendors a strict duty of confidentiality and adherence to firm policies, as well as prompt correspondence in the case of a breach.¹⁴⁴ The simplest way for firms to achieve this is by including it as a mandatory clause in their vendor contract, which cannot be omitted or drafted around. Furthermore, strict adherence to firm policies will ensure quality control of the output from the GPAI. Closely monitoring the AI vendor as if it were a nonlawyer assistant mitigates the opportunity for hallucinations or other AI tics to rear their ugly head and be of detriment to the firm’s clients’ matters. In the end, similar to nonlawyer

141. MODEL RULES OF PRO. CONDUCT r. 5.3 (A.B.A 2025).

142. *See id.*

143. A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18.

144. *See id.*

assistants, the lawyer must use their own independent professional judgment when finalizing the deliverable.¹⁴⁵

ii. Retain Economic Neutrality (Rule 5.4)

Firms must avoid using the vendor's pricing system, in nearly all cases, as it would likely generate a degree of influence on the firm's legal strategy, which violates the firm's professional responsibility. Maintaining economic neutrality is crucial to attorneys, as even the smallest amount of influence cuts counter to their fiduciary duties taken when sworn in under the bar. Per-token or other use-based fees should be avoided at all costs by firms, as those models incentivize the users to limit their research in the attempt to save money. The result is a less inclusive, exhaustive degree of research and ultimately an ill-informed opinion given to the client. To eliminate this risk, firms should seek out fee systems structured as flat-fee subscriptions, similar to that of Westlaw, LexisNexis, or Microsoft Office Suite. Doing so should result in lower overall cost to the firm, and again, reduce the risk of less comprehensive research for the final work product.

The argument posited here is maintaining economic autonomy by reducing the influence exerted on the firm and its lawyers, which a flat-fee subscription accomplishes as it neutralizes any lingering concerns of limits imposed on access for each respective matter. If firms wish to pass some or all of the costs onto their clients, they must disclose this as direct, non-marked-up expenses in accordance with the fundamental principle of reasonable fees imposed on lawyers under Model Rule 1.5.¹⁴⁶ The end goal is to disconnect the marginal cost of an AI query and the lawyer's decision-making process.

iii. Model Knowledge and Inquiry

In order to exercise independent judgment, the lawyer must have a fundamental understanding of the tool he is using. Blind use of a tool to inform a lawyer's work product cuts counter to their obligation to render the final verdict on a work product through professional judgment. As such, it is on the firm to instill methods which provide sufficient due diligence to aid in their lawyers' fundamental understanding of the GPAI tool being used.¹⁴⁷

145. *See id.* The firm cannot delegate the responsibility of final work product to the vendor. This remains the sole responsibility of the lawyer using the GPAI, or even the overseeing partner of the lawyer using the GPAI. A delegation of final judgment of the work product to the vendor would be in direct conflict with the firm's, and its lawyers', duties pursuant to Model Rule 5.3.

146. *Id.* (citing MODEL RULES OF PRO. CONDUCT r. 1.5 (2025)).

147. *See* COMM. IN PRO. RESP. & CONDUCT, *supra* note 34, at 2. The lawyers need not have a detailed knowledge of the system used, but rather reasonably understand the model's capabilities,

Arming lawyers with a general understanding of the AI model used in their practice ensures the lawyer is in compliance with a multitude of ABA rules, such as Rules 1.1 (Competence), 1.3 (Diligence), and 3.3 (Candor Toward the Tribunal).¹⁴⁸ More recently, *ABA Business Law Today* published an article discussing the current developments within AI, emphasizing that “lawyers have an ethical obligation to learn, understand and keep pace with how AI can be used in the daily practice of law.”¹⁴⁹ While the duty to inform oneself of the certain capabilities and limits of models used rests on the firms and their lawyers, the vendor must cooperate by providing sufficient transparency to allow for such continuous learning by legal professionals. If a vendor falls below this minimum threshold of transparency regarding recent developments in their model, then firms must use their professional judgment and disqualify that tool for substantive use firmwide.

iv. Requisite Human Verification Protocol

Evidently established in *Mata v. Avianca* by the sanctions handed down and coupled with the subsequent opinions from the ABA and state bar associations, no generative output used for substantive legal purposes should be relied upon or delivered to a court or client absent rigorous human verification. A firmwide formal human verification protocol is foundational to a compliant FGIS system implemented within a firm. This protocol should require a qualified lawyer to independently check every factual assertion and legal citation against its primary source. The process of verification, including any corrections or modifications made to the AI’s output, must be documented in the firm’s workpapers or matter management system, creating an auditable record of the lawyer’s supervisory role and intellectual contribution. The creation of an AI Oversight Department within firms, similar to that of the conflicts department, is a surefire method to ensure full compliance firmwide, while reaping the benefits of the improved efficiency generated by the developments and immersion of AI in the legal industry.

limitations, and any potential biases that might surface. *See id.* Essentially, firms need only ensure their lawyers have an understanding of the key aspects of the model informing their practice. *See id.*

148. *See id.*; MODEL RULES OF PRO. CONDUCT r. 1.1, 1.3, 3.3 (2025).

149. Bradford Newman & Adam Aft, *Recent Developments in Artificial Intelligence Cases and Legislation 2025*, A.B.A. BUS. L. TODAY (Aug. 5, 2025), https://www.americanbar.org/groups/business_law/resources/business-law-today/2025-august/recent-developments-artificial-intelligence-cases-legislation/ [https://perma.cc/4PUE-596T] (explaining that lawyers must take on the burden of staying abreast to advancements in the field of AI so that they can continue to use such tools to maximize efficiency for clients, while upholding the professional standard of legal practice).

B. PILLAR B: FIDUCIARY-GRADE CONFIDENTIALITY SAFEGUARDS

The purpose of this pillar is to thwart the risk of waiving privilege and ensure strict adherence to lawyers' duty of confidentiality pursuant to Model Rule 1.6.¹⁵⁰ Its thrust is the creation of clear contractual provisions against client data misuse when utilizing GPAI tools by generating a secure and isolated environment to work within. Creating such an environment is conducive to the continued adoption of GPAI throughout the legal industry as AI progresses.

i. Private Deployment

The ideal default environment for deploying GPAI in substantive legal work is a private, segregated network that prevents client data from being exposed to public-facing systems. This can be achieved through on-premise hosting or a dedicated Virtual Private Cloud (VPC) instance provided by a major cloud provider.¹⁵¹ The central idea here is that firms must ensure that client data is separated from that of the general public.¹⁵² Firms exerting this control over the GPAI models is recommended in a vast number of recent bar opinions, such as ABA Formal Opinion 512, Florida Ethics Opinion 24-1, California Guidance, and New York City Bar Formal Opinion 2024-5.¹⁵³ The emphasis on data isolation by four separate bar opinions makes evident the power this method has on mitigating the commingling of privileged, confidential client data with that of the general public. Taking this step is arguably one of the most effective technical means of ensuring continued client confidentiality for firms as GPAI permeates the legal industry.

150. MODEL RULES OF PRO. CONDUCT r. 1.6(a), (c) (A.B.A. 2025) (stating a lawyer must protect "information relating to the representation of a client" and make "reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to," such information).

151. See Jeff Johnson, *Best Practices for Confidentiality in the Age of AI-Powered Legal Tools*, JD SUPRA (Sep. 8, 2025), <https://www.jdsupra.com/legalnews/best-practices-for-confidentiality-in-1379976/> [<https://perma.cc/4U6D-CY4S>].

152. See *id.* By using on-premise or VPC systems, the firm effectively segregates sensitive, privileged client data from that of the general public by not prompting a publicly accessible model. Publicly accessible models, such as ChatGPT, are self-teaching, meaning that the model utilizes its users' input to inform itself and learn. After the self-learning model internalizes this new information from the input, it will spit the information out as "new" output moving forward. Thus, if a firm were to use a publicly accessible GPAI model, there is a high likelihood that privileged client data could be leaked, compromising the firm's duty of confidentiality pursuant to Model Rule 1.6.

153. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18; Fla. Bar Ethics Op., Op. 24-1, *supra* note 21; COMM. IN PRO. RESP. & CONDUCT, *supra* note 34; N.Y.C. Bar Ass'n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19.

ii. Contractual “No-Self-Learning” Guarantee

To further bolster efforts to retain client confidentiality during the deployment of GPAI in substantive work, firms must champion contractual negotiations that prohibit models from using input for self-learning purposes. These negotiations should include at least a provision in the contract with the vendor making it explicit and unambiguous that no firm data, client data, or user prompts shall be used to train, retrain, or otherwise improve the vendor’s model.¹⁵⁴ No-learning guarantees laid out in explicit, unambiguous language are non-negotiable if firms are to successfully integrate GPAI into their substantive work at scale.¹⁵⁵ These contractual guarantees further insulate client confidentiality, while still allowing the firm to streamline substantive work to optimize outcomes and efficiency in services rendered to their clients.

iii. Data Tailoring—Minimization of Identifying Data

Firms must adopt and strictly enforce a policy of data minimization amongst their lawyers when prompting GPAI. Extensive, thorough vetting of client information must be a priority firmwide before inputting it into the AI model to ensure only the most critical data is utilized to inform the prompt. To accomplish such a task, firms should consider implementing a FGIS framework like the one posited in this Article. By doing so, firms would establish a firmwide safety net ensuring that the minimum amount of data necessary to accomplish the task is used. Furthermore, all personally identifiable information and other sensitive identifiers should be redacted unless inclusion of such information is strictly necessary for proper prompting and output analysis. In all cases where GPAI is informed with confidential data, a concerted, conscious decision must be made and documented regarding the scope of the data used by at least the using lawyer.

iv. Audit-Ready Logging

To ensure the implementation of these verification duties and maintenance of confidentiality standards, firms must demand that their lawyers follow a standardized logging system which provides comprehensive, immutable logs of the lawyer’s use for each particular case. The breadth of these logs must include every prompt entered, all output received (whether used or not), the identity of the user, a timestamp, and

154. See Jordyn C. Tye, *Exploring the Intersections of Privacy and Generative AI: A Dive into Attorney-Client Privilege and ChatGPT*, 64 JURIMETRICS J. 309, 329 (2024).

155. See *id.*

record of the human verification process of the output ultimately used to inform the work product delivered to the court or client.¹⁵⁶ By harnessing the power of GPAI as a tool in the course of legal representation, firms are opening themselves up to additional liability, which these logs should minimize if a claim is brought by a client, court, agency, etc. against the firm. The use of these comprehensive logs serves as a shield for firms in the event a claim is brought. This system merely provides audit trails serving as evidence the firm can use to demonstrate to clients, courts, or disciplinary authorities that it has acted with reasonable care. They are in no way a silver bullet for firms.

v. *Informed Client Consent*

When it comes to client consent regarding firm usage of GPAI, the generic clauses which are standard in engagement letters are insufficient. Instead, due to the heightened possibility of breaches and compromised confidential data, informed consent by the client is required.¹⁵⁷ The ABA's Formal Opinion 512 underscores that informed consent in these cases requires a specific, detailed discussion with the client.¹⁵⁸ Such discussion, to be FGIS-compliant, must be matter-specific and clearly explain the nature of the GPAI tool(s) to be used and the scope of application to the client's matter. Furthermore, the discussion should include full transparency of the firm's internal controls, such as the FGIS framework posited in this Article if implemented by the firm. A comprehensive discussion hitting on the model, scope of use, internal controls, and the cost-benefit analysis is the minimum.

vi. *Organization of Firmwide AI Oversight Department*

Although law firms' existing structures can support FGIS adoption, implementation without modification risks creating inefficiencies in client representation. It will always be a mainstay that the lawyers who use the GPAI must always use their professional judgment in what to use from the output in their work product. However, the creation of additional duties for supervising lawyers within the firm to review that work product before delivering it to the court or client could detract from their capacity to handle

156. Recording of the verification process for the output actually used for informing the deliverable is of the utmost importance, rather than a record verifying all the output provided. Since the duty ultimately rests upon the lawyer to use their professional judgment to decide what should be used in the course of representation, then it implicitly follows that there was a vetting of *all* the output provided to the user. Their process of evaluating the output will inherently walk through the steps of why certain information was not used in the representation, so a separate duty to emphasize that would be redundant.

157. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18.

158. See *id.*

their current workload or take on more work, which could cut into firm profitability. Thus, the creation and implementation of an AI Oversight Department to handle the ultimate verification of the work product could drive greater efficiency and higher firm profitability. The structure of this department would resemble that of the conflicts departments most firms already deploy.

VI. FROM THEORY TO PRACTICE: IMPLEMENTING FGIS WITH AUDITABLE GOVERNANCE ARTIFACTS

When something is merely in the form of an abstract principle, it is of little use to those to whom the theory is meant to apply. The theory behind FGIS has been laid out, yet to be effective, it must be translated into a practical, operational, and auditable governance framework. In terms of an ethical theory, anyone can say they comply—but those who can show they comply through tangible “governance artifacts” like policies, checklists, and records that document the process will be given higher credence. FGIS demands the implementation of such governance artifacts in the form of a “governance kit,” which allows for seamless implementation of a standardized ethical approach across the firm in the ever-evolving world of GPAI and interconnectivity with remote work and industry globalization. While this governance kit surely benefits the firm by providing for clear oversight for all firm employees, it also provides a tangible form of compliance for any stakeholder, including clients, auditors, courts, agencies, etc., who might require such review. These artifacts are the natural, practical outgrowth of the duties of competence, supervision, and communication demanded from the developments in GPAI and articulated in the recent wave of ethics opinions.¹⁵⁹

A. THE KIT OF GOVERNANCE ARTIFACTS

i. AI Registry

Foundational to a firm’s AI governance is a centralized and exhaustive inventory of all firm-approved AI tools and models for attorney use.¹⁶⁰ By

159. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18; Fla. Bar Ethics Op., Op. 24-1, *supra* note 21; COMM. IN PRO. RESP. & CONDUCT, *supra* note 34; N.Y.C. Bar Ass’n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19.

160. See Mark Doble & Alexi Al, *Most Law Firms Are Building on Leaky Foundations, and AI Will Expose It*, NAT’L L. REV. (July 8, 2025), <https://natlawreview.com/article/most-law-firms-are-building-leaky-foundations-and-ai-will-expose-it> [<https://perma.cc/5A32-KMBB>]. AI governance should address tool approval processes that assures all software is approved for compliance and security before use. *Id.* This is necessary as firms that are building from publicly

circulating and enforcing the notion of the approved AI tools and models, firms mitigate the introduction of “shadow AI,” which is an internal threat posed by unsavvy attorneys using unapproved models.¹⁶¹ Strict enforcement of the governance kit drastically drives down the occurrence of shadow AI usage on client matters, which, in turn, reduces the chances of breaches of confidential client information.¹⁶² Firm certification of AI tools and models should not be a one-off, but rather a continuous, ongoing review process to ensure that those certified tools and models meet the minimum threshold demanded to meet ethical obligations assigned to lawyers.¹⁶³ For each certified tool or model, the registry should document critical information, including: the specific model name and version (e.g., “GPT-4 Turbo via Azure OpenAI Service”); the deployment type (e.g., “Firm’s Private Azure Tenant”); key contractual terms, with a simple yes/no indicator for the presence of a “no-training” guarantee; a link to the vendor’s security and privacy documentation; a list of approved use cases; and the name of the internal business or IT owner responsible for managing the vendor relationship and overseeing the technology. This comprehensive registry system is a practical model firms can adopt to prevent shadow AI use and ensure that only technologies that meet the firm’s FGIS standards are accessible to its lawyers.

ii. Use-Case Catalog

Development of a use-catalog will provide the firm’s lawyers with clear, actionable guidance on navigating the use of GPAI. The best way to accomplish this is by listing out a color-coded framework with examples of what activity falls under each varying level of human verification rigor. The following tasks involving generative AI in legal workflows are categorized by risk level and require varying degrees of control and approval:

accessible AI models, like ChatGPT, exponentially increase their chances of having a breach to their system. *See id.*

161. *See generally id.* (“Shadow AI refers to AI tools introduced into a workplace without the awareness or approval of IT, legal operations, or firm leadership. These tools are widely available and easy to access, which has led to their rapid adoption across all types of organizations. A recent study found that 99 percent of enterprises are exposed to Shadow AI, and fewer than one in seven have implemented any AI-specific security or governance controls.”).

162. Firms may have to go as far as locking down accessibility of certain site public models and tools when devices are connected to the firm’s VPN (or Wi-Fi if on firm premises) to ensure full compliance and eliminate all risk of shadow AI use by otherwise noncompliant lawyers. While this drums up thoughts of Big Brother if viewed from a broader scope, it is not out of the ordinary for businesses to limit accessibility of sites for their employees when devices are connected to the firm’s VPN (or Wi-Fi if on firm premises).

163. *See id.*

- i. Level 1 – Green: Low-risk tasks, such as formatting documents or checking grammar and style in firm-drafted briefs, are classified as “Green (Utility)” and may be performed at the user’s discretion, provided client or matter identifiers are redacted. Similarly, translating publicly available court opinions or news articles falls under the same category, assuming no confidential information is input.
- ii. Level 2 – Amber: Moderate-risk tasks, labeled “Amber,” require more oversight. For example, generating a first-draft summary of a non-confidential court filing or deposition transcript is considered “Amber (Supervised Utility)” and necessitates initiation of the FGIS Checklist, along with documented human verification of key facts and conclusions by a responsible attorney. Drafting initial versions of standard NDAs or engagement letters using firm-approved templates and client-provided terms is “Amber (Supervised Substance)” and demands completion of the FGIS Checklist, private deployment with a no-training guarantee, full human review and validation, and documented client consent, all under the supervision of an AI Oversight Department lawyer. Legal research to identify relevant case law also falls under “Amber (Supervised Substance)” and requires private deployment, checklist completion, and independent verification of citations and summaries by a qualified lawyer, with documentation maintained in workpapers submitted to the AI Oversight Department for compliance review.
- iii. Level 3 – Red: High-risk tasks, marked “Red (High-Risk Substance),” include drafting core analytical sections of legal briefs or client advisories, which require all Amber-level controls plus the matter partner’s review and AI Oversight Department approval of the human verification process. Modeling litigation outcomes or predicting judicial rulings based on confidential facts also falls into this category and must be explicitly approved by the head of the AI Oversight Department, with clear communication to the client about the probabilistic and non-binding nature of the output.
- iv. Level 4 – Black: The most restricted category is “Prohibited,” which includes finalizing dispositive legal

advice based primarily on generative AI analysis; such use is not permitted under any circumstances.

In sum, this catalog operationalizes the crucial distinction between low-risk “utility” uses of GPAI and high-risk “substantive” uses. By creating a color-coded system (e.g., Green, Amber, Red, Black), the firm can provide intuitive, at-a-glance rules that guide daily practice and reduce ambiguity. This artifact directly addresses the need for clear firm policies as mandated by the supervisory duties under Model Rules 5.1 and 5.3.

iii. FGIS Checklist (Per-Matter)

For the substantive matters approved for GPAI use, the lawyers must complete a mandatory FGIS Checklist which will be submitted and kept on record. This will be a one-page document that serves as not only a reminder for the lawyer to comply, but also a formal attestation that they did.¹⁶⁴ The checklist should also require AI Oversight Department review of both the work product and the process by which the GPAI output was used, ensuring full compliance with firm standards.¹⁶⁵ Additionally, for “Red” matters, the checklist must also be reviewed by the matter partner before it gets the stamp of approval to submit to the court or send to the client.¹⁶⁶ Details to be documented on the checklist include: the specific model from the registry being used; confirmation of the private deployment and no-training guarantee; the name of the lawyer assigned as the “verification lead” responsible for final review; confirmation that audit-ready logging is enabled; the billing method being used (e.g., fixed fee, overhead); and the date and method by which matter-specific, informed client consent was obtained. Again, this checklist serves a dual purpose of reminding the lawyer of compliance requirements, while also serving as another form of auditable material for the firm.

iv. Vendor Term Sheet

To ensure consistency and rigor in contracting, the firm’s legal or procurement department should collaborate with the AI Oversight Department to develop a stock Vendor Term Sheet or a standard contract addendum for all AI vendors. In this document, the firm should clearly put forth all non-negotiable clauses which ensure the firm remains in full

164. The checklist is yet another auditable form for the firm to use in instances when ethics are called into question regarding the use of GPAI. It serves as a form of evidence for the firm to point to when necessary to protect itself when faced with a challenge.

165. See *infra* Section VI.2.b (discussing AI Oversight Department as a posited creation within law firms moving forward).

166. See *supra* Section V.A.2 (aligning with the use-catalog framework).

compliance with the FGIS standard. While the terms included by each firm may vary, the key terms that must be included in all FGIS-compliant Vendor Term Sheets are an explicit “no-training” clause;¹⁶⁷ robust data security and encryption requirements;¹⁶⁸ quick response notification system for instances when a breach occurs;¹⁶⁹ the right for the firm to conduct periodic audits or tests on the vendor’s security controls;¹⁷⁰ and strong indemnification for any breaches of confidentiality or IP infringement claims that flow from the firm’s use of the vendor’s tool(s).

v. *Incident Response Playbook*

Considering that no system is foolproof, a firm must have a pre-planned response playbook in place specifically tailored to incidents of AI-related issues. Failure to create such a response playbook would make the firm hard-pressed to ensure an enduring FGIS-compliant practice. The playbook should be clear and easy to apply on the fly, as incident response time is traditionally of the essence. All triggering events should be clearly defined for their respective threshold.¹⁷¹ For each class, the playbook should adjust accordingly in addressing the proper steps to be taken by the firm for matters such as internal escalation paths, client notification procedures, and protocols for initiating a hold on further submissions for the matter until the situation is under control.¹⁷²

167. In this clause, firms should consider adding strict language for confidentiality purposes stating that all inputs and outputs are the sole property of the firm, and that the vendor cannot retain a copy, unless required to do so for legal purposes. But the former exception is only applicable for the output, and any kept outputs must have all potentially identifying language redacted when kept on file for the vendor. The exception should not extend to inputs because there is a no-training provision which renders retaining the inputs pointless for the vendor in the first place, but also from the firm’s perspective, this further insulates the protection of confidential client information by keeping it out of the hands of third parties.

168. It might be necessary to do a deep-dive analysis on the vendor’s current state of affairs on this to determine if they will be adequate or if more is required for the firm to trust the vendor with their business and ultimately their client’s potentially confidential information.

169. While there is no one perfect time, a reasonable and general time limit would be within seventy-two hours. Three days is more than enough time for a vendor to identify a breach and analyze the extent and severity of it before reporting to the firm. Even if the breach is resolved with no casualties to the firm’s data, the vendor must still notify the firm of the breach, how it was resolved, and how it will be prevented from occurring again in the future.

170. This could be included in the language under the data security and encryption clause as well, or even a “data security” section of the agreement might be better as it would be all encompassing and put it all together for ease of review at later times.

171. Instances such as a vendor security breach, the discovery of a significant hallucination in a document already filed with a court or sent to a client, or a material change in a vendor’s terms of service will likely all require a different response from the firm when compared with one another, or with other triggering events the firm lists.

172. Additionally, the firm should consider creating a separate playbook for the instances when a court filing was impacted by AI use. The response needed to address court-related issues

VII. CURRENT FIRM ORGANIZATIONAL STRUCTURE FALLS SHORT

Most law firms already have departments dedicated to conflicts of interest, risk management, and professional responsibility. However, none of these are designed to handle the unique challenges posed by the integration of AI into legal practice. Unlike traditional legal tools, AI models operate probabilistically rather than deterministically; may hallucinate or fabricate information; and continuously evolve through updates and machine learning. These characteristics demand a new kind of oversight in addition to the FGIS standard posited in this Article. An oversight that is proactive, technically informed, and ethically grounded. Delegating the duties under the FGIS standard to the lawyers while doing their work should be relatively seamless; however, assigning additional duties to the partners of the firm could ultimately cut into profitability as it draws the partners' attention away from generating business. The creation and deployment of an AI Oversight Department within the firm is a feasible solution to any issues posed by the current firm structure and successful FGIS implementation. This department would function analogously to a conflicts department but with a distinct mandate: to ensure that all AI usage within the firm complies with the FGIS framework, and thus with the ABA Model Rules and applicable state ethics opinions.

This department would be staffed entirely by licensed attorneys, which avoids additional precautions for nonlawyer issues under ABA Model Rule 5.3. The AI Oversight Department would operationalize the FGIS framework set forth in Section IV of this Article, including monitoring AI tools and models, reviewing case law for hallucinations, auditing lawyer entries in the use-catalog, negotiating vendor terms, acting as point of contact for security concerns, and other duties firms may deem proper.

A. PROPOSAL: CREATION OF AI OVERSIGHT DEPARTMENT

i. Structure & Staffing

The AI Oversight Department is handling the enforcement of arguably the most crucial aspect of the firm's future success. This demands that only the most qualified legal professionals be appointed to the Department to ensure rigorous oversight and full compliance with ABA standards. As such, the Department will be staffed exclusively with licensed lawyers. This requirement ensures that each person vetting the AI models, AI output, etc.

will likely differ greatly from that needed for clients, and the punishment for such errors could have more severe outcomes, though client-related punishments should not be taken lightly either.

is bound by and understands the professional obligations imposed on the legal profession under the ABA Model Rules and state rules.¹⁷³ The structure will be an organizational hierarchy. As the degree of severity in the AI use increases, so does the level of hierarchy required for proper review of the matter.¹⁷⁴

ii. *Core Responsibilities*

To ensure universal, consistent implementation of the FGIS framework firmwide, the AI Oversight Department should be entrusted with five core responsibilities: (1) tool vetting and approval, (2) usage protocols and logging, (3) compliance audits, (4) training and guidance, and (5) client disclosure and consent. Each responsibility is designed to ensure the successful operationalization of FGIS in the firm's AI usage.

a. *Tool Vetting and Approval*

Evaluation of the AI tools and models to be used by the lawyers of the firm is a critical element of the FGIS framework.¹⁷⁵ As previously stated, this vetting process ensures that the firm uses only the most trusted models for inputting what could be confidential information. The AI Oversight Department will need to determine whether a ban on certain public models is necessary to eliminate all risk of shadow AI, or if lawyers can be trusted to use only vetted sources following their extensive training process.¹⁷⁶ Additionally, the AI Oversight Department will want to inquire into the vendor's internal processes surrounding security of data.¹⁷⁷ Strong data encryption, both in transit and at rest, appears to be the focal point for vendors at this moment.¹⁷⁸

Following the approval of the actual model itself, finding that it is free of frequent hallucinations and has safety precautions up to par, the AI

173. Staffing entirely with licensed lawyers also eliminates the need for additional oversight to ensure nonlawyers are compliant with lawyer obligations. If this additional step of precaution was still needed, then it would arguably render the implementation of the AI Oversight Department pointless, as it would cut into the efficiency the AI Oversight Department is meant to foster.

174. See discussion *supra* Section V.A.2.

175. See discussion *supra* Section V.A.1.

176. If not a strict ban or blockage, then the Department might consider cautionary language pop-ups if practicable as another level of defense for the unwary lawyers who inevitably forget their training and believe no harm can come of using a public model.

177. Oliver Roberts, *Legal AI Unfiltered: Legal Tech Execs Speak on Privacy and Security*, NAT'L L. REV. (Mar. 31, 2025), <https://natlawreview.com/article/legal-ai-unfiltered-legal-tech-exec-speak-privacy-and-security> [<https://perma.cc/65T6-TBH6>] (emphasizing the importance of verifying vendor credentials and thoroughly reviewing policies related to data storage, confidentiality, limitation of liability, data retention, and breach notification requirements before deploying any AI tool).

178. See *id.*

Oversight Department will need to negotiate hard on the Vendor Term Sheet before adding the model to the firm-approved list. Non-negotiables are “no-learning,” extensive indemnification for breaches or leaks, and no retention of inputs.¹⁷⁹ Even with the “no-learning” clause, the firm will still want to negotiate the disposal of user inputs and the outputs to the greatest extent possible. Doing so will drastically drive down the chance of client data being held in the hands of a third party and compromised. While there is some validity to the vendor’s likely counter that they need to hold on to certain data or certain amounts of data to meet their legal requirements, there should be, at the minimum, extensive redaction of all potentially identifying information in any of the retained data.

b. Usage Protocols and Logging

The enforcement of the comprehensive catalog and on-the-fly FGIS checklist fleshed out in Section V will rest on the AI Oversight Department to guarantee FGIS compliance. The color-coded catalog clearly lays out what is required depending on the severity of the matter, and the AI Oversight Department should implement it as such and provide extensive training sessions to reinforce its importance in the course of practicing at the firm.¹⁸⁰ Additionally, the FGIS per-matter checklist provides an on-the-fly reminder for lawyers of what their duties entail for each respective undertaking.¹⁸¹ The AI Oversight Department should consider making it a requirement that the lawyer submits their completed checklist for compliance review regardless of the color level the matter is.¹⁸² An additional check the AI Oversight Department might consider is the encouragement of metadata tagging¹⁸³ of documents that include AI-generated content. The use and organized keeping of these logs serve as another compliance trail and can be used in audits and malpractice defense as a shield for the firm.

c. Compliance Audits

An FGIS-compliant firm is not a one and done implementation process. Rather, it is a continuous process of regularly conducting internal audit

179. *See supra* Part VI.

180. *See supra* Section V.A.2.

181. *See supra* Section V.A.3.

182. *See discussion supra* Section V.A.2.

183. *See* Amalie Widerberg, *Metadata Tagging – Best Practices*, FOTOWARE: BLOG (Nov. 12, 2024), <https://www.fotoware.com/blog/metadata-tagging-best-practice> [<https://perma.cc/V8HG-WXX4>]. “Metadata tagging is the process of adding descriptive information, or ‘tags,’ to digital files” for easier identification; a common practice lawyers already use. *See id.* However, creating certain signaling terms for AI use-specific matters should be considered to generate a more organized file system for the firm for future reference. *See id.*

checks on AI usage to ensure continued compliance with ethical obligations and internal policies as they evolve hand in hand with AI. The AI Oversight Department must adhere to such a standard if the firm wishes to have continued success in curbing ethical violations and maintaining the benefits it reaps from the integration of AI in its practices. Three methods the Department might consider implementing to achieve success in its compliance audits are: (1) randomized spot checks, (2) red flag detection tools, and (3) remediation protocols.

1. **Randomized Spot Checks:** These can be thought of as almost a road stop checkpoint—the Department will review the AI catalog logs across the various practice areas to verify that outputs were properly reviewed and edited.¹⁸⁴ The randomness of the checks promotes the idea that lawyers must always be in compliance for internal audit review purposes, but more importantly for the purposes of ethical obligations imposed on them as licensed lawyers.¹⁸⁵
2. **Red Flag Detection Tools:** These tools operate primarily to eliminate the risk of shadow AI usage by the firm’s lawyers. The AI Oversight Department, in coordination with the firm’s IT department, should develop internal tools or scripts to identify potential issues such as unverified citations or excessive verbatim use of AI-generated text. Such tools could also be used to pinpoint the possibility of copying and pasting AI output without first rendering proper professional judgment by keying in on excessive use of em dashes.¹⁸⁶ Additionally, these internally developed tools can help spot when client data is entered into public AI systems and allow for the AI Oversight Department to step in and stop that practice from spreading amongst the firm and putting confidentiality in harm’s way.

184. Firms should consider making practice area subject matter review committees within the AI Oversight Department to minimize chances of AI mishaps from slipping through the cracks. While this might present more difficulty for midsized firms, large firms should have no issue supporting a proposal of such specialized review.

185. *See infra* Section VI.C (explaining firms want to be careful not to implement this process through fear but should consider the possibility of implementing internal punishments for noncompliance).

186. *See* Nick Potkalitsky, *Why AI Can’t Stop Using Em Dashes*, LINKEDIN (Sep. 11, 2025), <https://www.linkedin.com/pulse/why-ai-cant-stop-using-em-dashes-nick-potkalitsky-phd-7slle> [<https://perma.cc/NLX6-G6LB>]. AI is notorious for the use of certain syntax, and em dashes are one of the top signals of an AI generated sentence. *See id.*

3. Remediation Protocols: AI Oversight Departments might consider establishing a tiered response system for violations among lawyers depending on the severity, which could range from retraining to disciplinary action.

d. Training and Guidance

Ongoing education is as essential to maintaining ethical compliance and adapting to new AI developments as internal compliance audits are. Considering this, AI Oversight Departments should develop and adopt comprehensive onboarding, continuing legal education courses and credits, and distribute ethics bulletins on a monthly basis to inform lawyers about recent developments in AI.¹⁸⁷ The preceding examples are surefire ways to keep lawyers in the know of cutting-edge developments in AI and ensure continuous ethical compliance.

e. Client Disclosure and Consent

As stated in nearly every ethics opinion that has been circulated since the rise of AI, informed client consent is non-negotiable.¹⁸⁸ The AI Oversight Department has the responsibility to draft the standardized language disclosing the use of AI tools, their potential role in the matter, and any implications this might have in billing or confidentiality.

B. BENEFITS & RISK MITIGATION

The establishment of a dedicated AI Oversight Department within law firms offers a strategic response to the ethical, operational, and reputational challenges posed by the integration of AI into legal workflows. The traditional law firm structure arguably does not possess the unique traits necessary to see strict compliance come to fruition and endure within the firm long-term, making the creation and implementation of an AI Oversight Department crucial to AI success in the legal industry. This section explores the multifaceted benefits of such a department and outlines how it serves as a critical risk mitigation mechanism.

1. Ethical Risk Mitigation: One of the most immediate benefits is ethical risk mitigation. As stated countless times throughout this Article, AI tools can hallucinate legal citations, misinterpret facts, or expose

187. See Gleb Tsipursky, *Justice Meets Algorithms: The Rise of Gen AI in Law Firms*, IND. ST. BAR ASS'N (May 13, 2025), <https://www.inbar.org/blogpost/1981821/510657/Justice-Meets-Algorithms-The-Rise-of-Gen-AI-in-Law-Firms> [<https://perma.cc/YL8U-CLLE>].

188. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18; Fla. Bar Ethics Op., Op. 24-1, *supra* note 21; COMM. IN PRO. RESP. & CONDUCT, *supra* note 34; N.Y.C. Bar Ass'n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19.

confidential client data if not properly vetted and supervised.¹⁸⁹ Again, the ABA's Formal Opinion 512 emphasizes that lawyers must understand the capabilities and limitations of AI tools and supervise their use as they would a junior associate or nonlawyer assistant.¹⁹⁰ By implementing standardized protocols and requiring human review of AI outputs, the Oversight Department helps prevent violations of Model Rules 1.1, 1.6, and 5.3.

2. **Enhanced Professional Accountability:** The Department reinforces a culture of accountability by establishing a clear stance of organizational commitment to the matter. Establishing a department with the sole responsibility of generating and enforcing AI oversight signals to the lawyers that it is not something to be trifled with, and that they must put forth the same effort in maintaining the high degree of ethical conduct as the firm. The Department mandating certain practices when utilizing AI throughout the course of client representation creates a deep-rooted respect for the process of upholding the ethical obligations of the legal industry in a time when it may be easier for lawyers to shirk these duties rather than abide by them. These practices have the bonus of providing the firm with an auditable trail it can use in defense of itself if there are claims raised in the future regarding the representation and use of AI in such representation.
3. **Operational Consistency:** Even with the recommended FGIS framework advanced by this Article, the actual implementation and practice of such process would likely vary across the firm's offices and practice areas. This is a risk that a firm wishing to be fully FGIS-compliant cannot leave to fate. Instead, the firm must use the AI Oversight Department posited in this Article to ensure all offices and practice groups follow uniform standards for tool approval, usage protocols, and training.¹⁹¹ This reduces internal friction and improves firmwide efficiency.
4. **Client Trust and Transparency:** Client trust is also bolstered by transparency. Clients increasingly expect to be informed when AI is used in their matters. The AI Oversight Department will develop the standardized engagement letter language and informed consent

189. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18; Fla. Bar Ethics Op., Op. 24-1, *supra* note 21; COMM. IN PRO. RESP. & CONDUCT, *supra* note 34; N.Y.C. Bar Ass'n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19.

190. A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18.

191. See *supra* Part VI.

protocols to ensure compliance with Rule 1.4.¹⁹² With the strict demand for informed client consent,¹⁹³ clear client communication about AI use is the new norm for best practice in firms that prioritize ethical innovation.¹⁹⁴

5. **Competitive Advantage and Market Differentiation:** From a competitive standpoint, firms that proactively manage AI risks differentiate themselves in the marketplace. They attract tech-savvy clients and top legal talent who value responsible innovation. Firms that embrace ethical AI governance are better positioned to lead in a rapidly evolving legal landscape.¹⁹⁵ The age of AI in legal practice is coming whether legal professionals are ready to accept it or not; those firms who embrace it and usher it in with the AI Oversight Department set themselves up to be leaders in the new-look legal industry.
6. **Agility in Responding to Regulatory Change:** As AI continues to evolve, so does its regulation. The AI Oversight Department provides agility in responding to new guidance. It can monitor changes in ABA opinions, state bar rulings, and vendor policies—updating internal protocols and training materials accordingly. Leaving these duties to the practicing lawyers alone will result in inefficiency of their practice, cut into firm profitability, and ultimately neutralize any efficiency gained from the use of AI in the first place. This ensures that the firm remains compliant and ahead of regulatory developments.

C. CHALLENGES & IMPLEMENTATION CONSIDERATIONS

While the creation of a firmwide AI Oversight Department offers clear ethical and operational benefits, its implementation presents a series of practical and cultural challenges that law firms must navigate carefully. These challenges are not insurmountable, but they require deliberate planning, leadership buy-in, and a phased approach to change management.

One of the most immediate hurdles is cost and resourcing. Establishing a new department staffed exclusively by licensed attorneys entails significant

192. See *supra* Section VI.2.

193. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18.

194. See *AI & Legal Ethics: A Guide to ABA Rules in 2026*, NEXLAW (July 11, 2025), <https://www.nexlaw.ai/blog/ethical-ai-litigator-guide-2025/> [<https://perma.cc/YR2W-3GKH>]. As AI progresses and its capabilities expand into new use methods for legal practice, it becomes omnipresent for all matters. See *id.* The saturation of presence in work product demands that there be continuous and unambiguous communications with the client regarding any new developments in their matter through the utilization of AI's new tools. See *id.*

195. See Tsipursky, *supra* note 188.

investment. Firms must budget for salaries, training, infrastructure, and potentially new software tools to support logging, auditing, and compliance. Unlike conflicts departments, which are now standard and often supported by decades of institutional knowledge, AI oversight is a novel function with few precedents. As a result, firms may struggle to define the scope of the Department and justify its initial expense, especially in smaller or mid-sized practices.

Another challenge is cultural resistance. Attorneys accustomed to autonomy in their workflows may view AI oversight as intrusive or unnecessary.¹⁹⁶ Some may perceive it as a constraint on creativity or efficiency, particularly if they are early adopters of AI tools.¹⁹⁷ Overcoming this resistance requires clear communication about the Department's purpose. Not as a punitive body, but as a safeguard for ethical compliance and professional integrity. Leadership must emphasize that the Department exists to protect lawyers, clients, and the firm itself from the risks of unsupervised AI use.

There is also the issue of rapid technological evolution. AI tools are updated frequently, and new platforms emerge regularly. This creates a moving target for oversight. The Department must remain agile, continuously evaluating new tools, updating protocols, and revising training materials. Static policies will quickly become obsolete. To address this, firms' AI Oversight Departments must be agile and possess the ability to amend the FGIS policy framework ad hoc so that it can be revised quickly and incrementally as technology and regulatory guidance evolve.

Training and education pose additional implementation challenges. Lawyers are not technologists, and many may lack the foundational knowledge needed to understand how AI tools function. The AI Oversight Department must develop accessible training programs that demystify AI without overwhelming attorneys with technical jargon. These programs should focus on practical risks, such as hallucinated citations, data leakage, and overreliance, and offer clear guidance on how to use AI responsibly.¹⁹⁸

196. See generally *Big Law Firms Struggle to Monetize AI, Facing Operational and Cultural Barriers*, LEGAL NEWS FEED (Aug. 18, 2025), <https://legalnewsfeed.com/2025/08/18/big-law-firms-struggle-to-monetize-ai-facing-operational-and-cultural-barriers/> [https://perma.cc/DZQ3-NAJY].

197. See generally *The Legal Industry Report 2025*, A.B.A. (May 6, 2025), https://www.americanbar.org/groups/law_practice/resources/law-technology-today/2025/the-legal-industry-report-2025/ [https://perma.cc/GF7Q-3WS4]. Lawyers are often adopting the use of AI outside the firm, but the willingness to be on board of firmwide adoption is lagging, primarily due to concerns of policy, ethical, and cultural concerns that are linked to AI use in practice. See *id.* So, the AI Oversight Department is a practical solution to bridge the gap of concern here.

198. Annual seminars where the firm's FGIS framework is the main focus of discussion is one program that should be central in these efforts. Considering that the FGIS framework needs to be

Firms may consider a route such as monthly or quarterly circulations of developments in AI and how they impact the firm's FGIS framework so that the lawyers are up to date and utilize the AI tools in an appropriate manner with the current times. Partnering with legal tech vendors or academic institutions may also help bridge the knowledge gap and distill the information into a more digestible size for the often non-tech-savvy lawyers.

Finally, firms must contend with jurisdictional variability. While the ABA has issued national guidance through Formal Opinion 512, state bars may adopt different or more stringent positions on AI use.¹⁹⁹ The Oversight Department must monitor developments across jurisdictions where the firm operates and tailor its protocols accordingly. This may include creating jurisdiction-specific disclosures, consent forms, or usage restrictions to ensure compliance with local rules.

To mitigate these challenges, firms should consider a phased implementation strategy. This might begin with a pilot program of the FGIS framework in one practice group, allowing the Department to refine its protocols and demonstrate value before scaling firmwide. Early wins—such as preventing a citation error or improving client communication—can build momentum and support. Firms should also leverage existing structures, such as risk management or professional responsibility committees, to support the Department's launch and integration.

While the path to establishing an AI Oversight Department is complex, it is navigable. With thoughtful planning, transparent communication, and commitment to ethical innovation, law firms can build a structure that not only mitigates risk, but also positions them as leaders in the responsible use of artificial intelligence.

VIII. ALIGNING INCENTIVES: ECONOMICS OF ETHICAL AI

A robust governance framework like FGIS can only be sustained if the economic incentives of the law firm align with its ethical mandates. From a law and economics perspective, ethical rules without aligned incentives are inherently unstable and prone to failure. The economic logic is straightforward: when compliance with ethical rules reduces firm profitability while violations go undetected or lightly punished, profit-

malleable, the seminar should cover not only the original FGIS framework, but also new amendments made in the past year which make certain parts of the original FGIS framework antiquated.

199. *Compare* A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, with Fla. Bar Ethics Op., Op. 24-1, *supra* note 21. *See AI and Attorney Ethics Rules: 50-State Survey*, JUSTIA (2025), <https://www.justia.com/trials-litigation/ai-and-attorney-ethics-rules-50-state-survey/> [<https://perma.cc/2M3V-HBXD>].

maximizing firms will systematically under comply. Sustainable ethical frameworks must therefore align ethical obligations with economic self-interest, making “doing the right thing” also “doing the profitable thing.”

The duties of reasonable fees are central to ethical GPAI implementation.²⁰⁰ The profound efficiency gains promised by GPAI create the “AI efficiency paradox”: the billable hour model, which rewards time spent, is fundamentally misaligned with a technology designed to eliminate time.²⁰¹ This paradox reflects a deeper economic insight: the billable hour is a price mechanism that, in traditional legal practice, roughly aligned lawyer incentives with client welfare by rewarding effort and expertise.²⁰² GPAI breaks this alignment because it creates a wedge between input (lawyer time) and output (legal work product).²⁰³ A firm relying solely on the billable hour while using GPAI faces what economists call a “revenue destruction problem”: the technology that makes the firm more productive simultaneously destroys its revenue base.²⁰⁴

A firm relying solely on the billable hour while using GPAI faces perverse incentives to either underutilize the technology or obscure its use and bill for time not actually spent by human lawyers.²⁰⁵ These are not merely ethical failures, but economically predictable responses to misaligned incentives.²⁰⁶ When firms’ profit-maximizing strategy involves either forgoing efficiency gains or engaging in fraudulent billing, the market structure itself is fundamentally broken. The optimal strategy for individual firms creates negative externalities (reduced client welfare) and increases

200. See MODEL RULES OF PRO. CONDUCT r. 1.5(a) (A.B.A. 2025) (prohibiting lawyers from “mak[ing] an agreement for, charg[ing], or collect[ing] an unreasonable fee or an unreasonable amount for expenses”).

201. See, e.g., Bruce MacEwen & Janet Stanton, *Pricing AI-Driven Legal Services: The Billable Hour is Dead, Long Live the Billable Hour*, THOMSON REUTERS (Sep. 10, 2024), <https://www.thomsonreuters.com/en-us/posts/legal/pricing-ai-driven-legal-services-billable-hour> [<https://perma.cc/KQ8S-S6KP>] (“[H]ourly billing misaligns law firm and client interests. The client presumably is seeking a fast and reasonably priced response, but the billable hour works against both those results . . .”).

202. See Jonathan H. Choi, *In Defense of the Billable Hour: A Monitoring Theory of Law Firm Fees*, 70 S.C. L. REV. 297 (2018).

203. See Nancy B. Rapoport & Joseph R. Tiano, Jr., *Fighting the Hypothetical: Why Law Firms Should Rethink the Billable Hour in the Generative AI Era*, 20 WASH. J.L. TECH. & ARTS 43, 55–57 (2025).

204. See *id.* at 46.

205. See Stuart L. Pardau, *Bill, Baby, Bill: How the Billable Hour Emerged as the Primary Method of Attorney Fee Generation and Why Early Reports of Its Demise May Be Greatly Exaggerated*, 50 IDAHO L. REV. 1, 6 (2013).

206. See generally Mark C. Palmer, *Be Reasonable, People! AI’s Impact on Legal Fees*, 1 PROMPT 1, 6–7 (2025) (highlighting the underlying concern of reasonable billing by attorneys and how it is exacerbated with the rise of AI inclusion and use in legal practice, leading to heightened risk of misaligned incentives of inputs versus outputs).

aggregate inefficiency (underinvestment in productivity-enhancing technology).²⁰⁷

This directly undermines both the independence safeguards and confidentiality protections established in the FGIS framework. To resolve this paradox, firms must adopt new standards for disclosure, billing, and documentation that reinforce their fiduciary duties.

This tension reflects a crisis of confidence accelerated by GPAI adoption. The billable hour has “routinely been criticized . . . as creating incentives for lawyers to be inefficient, pitting the lawyer’s financial interest against that of the client.”²⁰⁸ The 2002 ABA Commission on Billable Hours Report declared that “the unintended consequences of the billable hours model have permeated the profession.”²⁰⁹ Today, clients are demanding change: an “overwhelming 80% of corporate legal executives expect their outside counsel bills to be reduced as a result of Gen AI technology,”²¹⁰ while recent surveys show “nearly 50% of firms surveyed said clients are pushing back harder on their rates.”²¹¹ High-profile disputes like X Corp.’s litigation against Wachtell, Lipton, Rosen & Katz—alleging nearly \$90 million in fees with blank time entries and single-word descriptions like “strategy” for charges exceeding ten thousand dollars—have made billing abuses visible.²¹²

The resolution lies in three interconnected principles that operationalize ethical duties and reinforce the FGIS framework: mandatory transparent disclosure, prohibition of hour inflation, and memorialization of human value.

A. THE LAW AND ECONOMICS OF INFORMATION DISCLOSURE

Before examining specific disclosure requirements, it is essential to understand the economic rationale for mandatory disclosure rules.

207. See Widerberg, *supra* note 183.

208. See Pardau, *supra* note 206, at 6.

209. See Susan Saab Fortney, *I Don’t Have Time to Be Ethical: Addressing the Effects of Billable Hour Pressure*, 39 IDAHO L. REV. 305, 313 n.41 (2003) (citing *ABA Commission on Billable Hours Report, 2001-2002*, A.B.A., at ix (2002)).

210. See Geoffrey D. Ivnik, *The Disconnect Between Law Firms and Clients on Use of Gen AI*, LEXISNEXIS (Feb. 9, 2024), <https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/the-disconnect-between-law-firms-and-legal-departments-on-impact-of-gen-ai> [https://perma.cc/XHM4-F4PB].

211. See Andrew Maloney, *Rates Will Go Up (Again), But Here’s Why Profitability Might Not Be Maximized*, AM. LAW. (Oct. 18, 2024, at 6:00 CT), <https://www.law.com/americanlawyer/2024/10/18/rates-will-go-up-again-but-heres-why-profitability-might-not-be-maximized> [https://perma.cc/UY32-DNQA].

212. See Ryan Mac & Lauren Hirsch, *Twitter Sues Law Firm Over \$90 Million Payment in Elon Musk Deal*, N.Y. TIMES (July 7, 2023), <https://www.nytimes.com/2023/07/07/technology/twitter-sues-wachtell-lipton.html> [https://perma.cc/HP25-CGG9].

Information economics teaches us that markets function efficiently only when parties have sufficient information to make rational decisions.²¹³ In the legal services market, severe information asymmetries exist as lawyers possess expert knowledge and information about their work processes that clients lack. This creates potential for opportunistic behavior where lawyers exploit their informational advantage.

Mandatory disclosure serves multiple economic functions. First, it reduces information asymmetry, enabling clients to make more informed decisions about value and quality.²¹⁴ When clients know that AI was used and under what safeguards, they can better assess whether the work product meets their expectations and whether the fee is reasonable relative to the inputs used.²¹⁵

Second, disclosure creates market discipline. When AI usage is transparent, competitive pressure incentivizes firms to use AI efficiently and to implement robust governance frameworks like FGIS.²¹⁶ Firms that use AI poorly or without adequate safeguards will lose clients to better-performing competitors.²¹⁷ This market mechanism reinforces ethical compliance by ensuring that a firm's economic interests and its ethical obligations move in the same direction—firms that cut corners on AI governance risk both client loss and reputational harm, while those that invest in robust oversight are rewarded with greater client trust and competitive standing.²¹⁸

Third, disclosure generates information spillovers that improve overall market efficiency. When firms transparently share their AI usage patterns and governance frameworks, the legal services market collectively learns what works and what does not.²¹⁹ This reduces duplicative experimentation costs and accelerates the development of best practices. From a social welfare perspective, disclosure mandates create positive externalities that justify regulatory intervention.

213. See PARISI, *supra* note 58, at 101 (explaining the efficient market hypothesis).

214. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 8–9.

215. See *supra* Part IV (disclosing the general FGIS framework employed by the firm for GPAI use in its practice is sufficient to reduce the asymmetric gap that exists and ensure that clients can assess the AI tools are being utilized appropriately—or at least in theory are being utilized appropriately in carrying out representation of their matter(s)).

216. See *supra* Part VI (Benefits & Risk Mitigation).

217. See *supra* Part VI (Benefits & Risk Mitigation).

218. See *supra* Part VI (Benefits & Risk Mitigation).

219. See generally George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488 (1970) (demonstrating that information asymmetry in markets allows lower-quality providers to misrepresent their quality, driving out higher-quality competitors unless mechanisms exist to credibly signal quality distinctions to consumers).

Fourth, disclosure serves as a bonding mechanism that credibly signals quality.²²⁰ Firms implementing comprehensive FGIS frameworks have incentives to disclose this fact because it differentiates them from competitors who may be using AI less responsibly.²²¹ Mandatory disclosure ensures that high-quality firms can signal their quality, while preventing low-quality firms from free-riding on opaque information environments.²²²

The economic case for mandatory disclosure is thus grounded in both efficiency (reducing information asymmetry improves market function) and equity (preventing exploitation of clients through informational advantages). This theoretical foundation supports the specific disclosure requirements detailed below.

B. MANDATORY, TRANSPARENT DISCLOSURE

Model Rule 1.4 requires lawyers to “reasonably consult with the client about the means by which the client’s objectives are to be accomplished.”²²³ From an economic perspective, this rule addresses information asymmetry in the agency relationship: clients need sufficient information to assess whether their agents (lawyers) are acting in their interests. When GPAI is used, information asymmetry increases dramatically because clients now face uncertainty about both lawyer judgment and algorithmic influence. Disclosure requirements restore informational balance necessary for efficient contracting.

When GPAI is used for substantive work under an FGIS framework, it is a material “means” of representation requiring contemporaneous disclosure. Firms should indicate on invoices, at the task or phase level, that GPAI was used and under what controls (e.g., “Legal Research (AI-Assisted, FGIS Compliant)”).

ABA Formal Opinion 512 makes disclosure mandatory when clients inquire about work methodology, when engagement agreements require it, or when GPAI use “will influence a significant decision in the representation.”²²⁴ This requirement aligns with the economic principle that material information affecting transaction value must be disclosed to prevent adverse selection and moral hazard. When clients lack information about AI

220. See generally Michael Spence, *Job Market Signaling*, 87 Q.J. ECON. 355 (1973) (positing that signaling is an excellent form of differentiating oneself from the pack).

221. See *id.*

222. See *id.* Establishing a comprehensive framework as the Bar makes it so that those firms who do not rise to the occasion will be faded out by clients seeking transparent and best practices as AI enters legal practice.

223. MODEL RULES OF PRO. CONDUCT r. 1.4(a)(2) (A.B.A. 2025).

224. A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 8.

usage, they cannot properly value services, creating market inefficiency and potential for exploitation.

Florida Bar Ethics Opinion 24-1 and the New York City Bar Association Opinion 2024-5 emphasize that disclosure should be addressed in writing, explaining the firm's GPAI approach and associated costs.²²⁵ This aligns with the "Informed Client Consent" component of FGIS Pillar B, which requires matter-specific discussion of the GPAI tool's nature, scope of application, internal controls, and cost implications.²²⁶ This disclosure could frequently be included in the engagement letter.

Disclosure serves multiple purposes: satisfying Rule 1.4's communication duty, demonstrating implementation of FGIS verification protocols and treatment of AI vendors as supervised nonlawyer assistants under Rule 5.3,²²⁷ and creating a foundation for value-based conversations. Economically, disclosure also creates accountability mechanisms that reduce agency costs.²²⁸ When firms must explicitly disclose AI usage, they face reputational incentives to ensure such usage is both effective and ethically compliant. This transforms disclosure from a mere information-sharing requirement into a market discipline mechanism.²²⁹

Rather than presenting GPAI as merely cost-cutting, disclosure allows lawyers to explain how FGIS-compliant GPAI enables higher-quality work to be delivered efficiently. This reframing is economically significant: it shifts the value proposition from input-based (time spent) to output-based (results achieved), facilitating the transition to value-based pricing models that better align firm and client incentives.

C. THE PROHIBITION OF HOUR INFLATION

ABA Formal Opinion 512, Florida Bar Opinion 24-1, and New York City Bar guidance establish a critical bright line: if GPAI saves time, lawyers cannot "back-fill" hours or bill for time the task would have taken manually.²³⁰ Lawyers billing hourly may only charge actual time spent

225. See Fla. Bar Op. 24-1, *supra* note 21, at 6; N.Y.C. Bar Ass'n Comm. on Pro. Ethics, Formal Op. 2024-5, *supra* note 19, at 6.

226. See *supra* Part V.

227. See *supra* Part V.

228. See Jensen & Meckling, *supra* note 64, at 305 (establishing the agency-cost framework showing how information, monitoring, and reputational mechanisms constrain opportunistic agent behavior).

229. See Carl Shapiro, *Premiums for High Quality Products as Returns to Reputations*, 98 Q.J. ECON 659 (1983) (showing that disclosure is a key aspect of reputation in a market and is thus a valuable signaling mechanism for firms when used accordingly).

230. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 12 ("If a lawyer uses a GAI tool to draft a pleading and expends 15 minutes to input the relevant information

crafting prompts, verifying outputs, and integrating work product.²³¹ From a law and economics perspective, hour inflation represents what economists call “negative-sum” behavior: the firm captures value (inflated fees) by destroying value elsewhere in the system (client overpayment, erosion of trust, market inefficiency).²³² When aggregated across the profession, such behavior reduces total social welfare by misallocating resources and reducing market efficiency.

This reaffirms longstanding principles from Formal Opinion 93-379: attorneys may not bill for more time than actually spent, and “a lawyer should . . . certainly never be motivated by anything other than the best interests of the client when determining how to staff or how much time to spend on any particular project.”²³³

This prohibition is critical given documented billable hour abuses. The 2002 Report recognized the billable hour creates penalties for efficiency and conflicts of interest between the lawyer and the client.²³⁴ In economic terms, the billable hour creates a classic principal-agent problem where the agent’s (lawyer’s) interests diverge from the principal’s (client’s).²³⁵ The agent maximizes compensation by maximizing hours, while the principal desires efficient completion of legal work.²³⁶ GPAI dramatically exacerbates this misalignment by making hour inflation both easier (clients cannot observe how much AI was used) and more tempting (the gap between manual and AI-assisted time widens).

With associates facing 1,800 to 2,000 hour annual requirements, a junior associate using GPAI to complete document review in five hours instead of twenty faces severe pressure: reporting only five hours costs three-quarters of the weekly billable target.²³⁷ The temptation to inflate is obvious; the ethical prohibition must be enforced through restructured compensation systems. This represents a failure of incentive design at the organizational level. When firms set billable hour requirements without adjusting for AI-

into the GAI program, the lawyer may charge for the 15 minutes as well as for the time the lawyer expends to review the resulting draft for accuracy and completeness.”).

231. *See id.*

232. *See* PARISI, *supra* note 58, at 311 (explaining variable-sum games).

233. A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 93-379, at 6 (1993).

234. *See* Fortney, *supra* note 209.

235. *See* Jensen & Meckling, *supra* note 64.

236. *See id.*

237. *See generally* Jordan Rothman, *0.1 Time Entries Definitely Add Up*, ABOVE L. (Aug. 11, 2023), <https://abovethelaw.com/2023/08/0-1-time-entries-definitely-add-up> [https://perma.cc/H9S8-DV7T] (“Depending on the type of firm, an associate’s billable hour requirement might average around 1,900 to 1,950 hours a year, although some firms have requirements in the 1,800-hour range, and other law firms have billable-hour requirements that exceed 2,000 hours.”).

driven efficiency gains, they create systematic pressure for hour inflation. The economically efficient solution requires restructuring compensation to reward productivity and quality rather than time spent.

D. ALIGNMENT THROUGH ALTERNATIVE FEE ARRANGEMENTS

This creates powerful incentives toward Alternative Fee Arrangements (AFAs).²³⁸ Fixed fees, value-based fees, or subscription models align firm financial interests with client efficiency expectations.²³⁹ By pricing outcomes rather than inputs, AFAs let firms retain benefits from FGIS-compliant GPAI investment while providing client cost certainty.²⁴⁰

From a law and economics perspective, AFAs solve the incentive compatibility problem that plagues hourly billing. Under AFAs, firms maximize profits by minimizing inputs (time and cost) while maximizing outputs (quality and value).²⁴¹ This directly aligns firm incentives with client interests: both parties benefit from efficient, high-quality legal work.²⁴² GPAI becomes an economic asset rather than a revenue threat because productivity gains increase firm margins under fixed pricing while simultaneously reducing client costs.

For example, a firm historically charging twenty billable hours for contract review can now complete it in five hours with FGIS-compliant GPAI, spending those five hours on higher-value work: issue-spotting, strategy, and client-specific refinement.²⁴³ A fixed fee reflecting twelve hours rewards the firm's FGIS investment through increased margin while saving the client forty percent compared to traditional billing. Both parties benefit from genuine productivity improvement.

This represents what economists call a "Pareto improvement": a transaction structure where both parties are better off compared to the status

238. See Shannon Kirk, *Alternative Fee Arrangements Put the Focus on Value, Not Time*, A.B.A. (Sep. 1, 2023), https://www.americanbar.org/groups/law_practice/resources/law-practice-magazine/2023-september-october/alternative-fee-arrangements-put-focus-on-value-not-time [<https://perma.cc/P7ZP-DU2H>].

239. Clients increasingly request AFAs because they understand the perverse incentive created by the billable hour. By shifting from time-based to outcome-based pricing, AFAs align firm incentives with client interests in cost predictability and efficiency, allowing firms to retain the economic benefits of technological investments while meeting client demands for transparent, competitive pricing. See Adams & Matheson, *supra* note 50, at 35–36.

240. See *id.*

241. See *id.* at 36–37.

242. See *id.*

243. See Katherine Medianik, *Artificially Intelligent Lawyers: Updating the Model Rules of Professional Conduct in Accordance with the New Technological Era*, 39 CARDOZO L. REV. 1497, 1506–07 (2018) (“[AI] permits [lawyers] to spend more time and money on high value advisory tasks and complex legal matters. These elements of implementing AI technology generate margins superior to competing firms, thereby creating a competitive advantage.”).

quo.²⁴⁴ The firm increases its profit margin per matter (from billing twelve hours of value while working five hours), and the client reduces costs (paying for twelve hours instead of twenty). The efficiency gain created by GPAI is shared between firm and client rather than being captured entirely by one party or hidden through hour inflation. This sharing of gains creates stable, sustainable market equilibrium where both parties have incentives to continue the arrangement.

Moreover, AFAs create proper incentives for investment in FGIS infrastructure. Under hourly billing, investments in AI governance represent pure cost with no revenue upside (and potentially revenue destruction if AI increases efficiency).²⁴⁵ Under AFAs, FGIS investments become profit drivers: better governance enables more confident AI deployment, which increases efficiency and further increases margins on fixed-fee engagements. This transforms ethical compliance from a cost center into a profit center, fundamentally aligning economic incentives with professional obligations.

E. THE MEMORIALIZATION OF HUMAN VALUE

In FGIS-compliant GPAI practice, lawyer value shifts from performing rote tasks to exercising critical judgment, strategic oversight, and rigorous verification.²⁴⁶ To justify fees, demonstrate FGIS human verification protocol compliance, and defend against malpractice claims or sanctions, this human value-add must be documented. From an economic standpoint, documentation serves to reduce information asymmetry about the quality and nature of legal services. When lawyers document their verification process, strategic judgments, and value-added contributions, clients can better assess the value they receive.²⁴⁷ This information disclosure improves market efficiency by enabling more accurate pricing and reducing adverse selection problems.

FGIS-required auditable workpapers should record the GPAI's initial output, the lawyer's independent verification process, corrections made, and strategic judgments about integrating GPAI-assisted analysis into final work product. This documentation also creates what economists call "credible commitment": firms commit ex ante to meaningful human oversight, and

244. See PARISI, *supra* note 58, at 215.

245. See generally Toby Brown, Bruce MacEwen & Janet Stanton, *Exploring the Impact of AI on the Economics of BigLaw*, ALM (Apr. 2024), https://images.law.com/contrib/content/uploads/documents/1/Law_com-Pro-Fellows-Report_Exploring-the-Impact-of-AI-on-the-Economics-of-Big-Law_April-2024.pdf [<https://perma.cc/F7Z5-B6FA>] (positing that implementation of AI in the legal practice will cut into profitability of firms due to increase in efficiency under the billable hour model).

246. See Medianik, *supra* note 243.

247. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 8–9.

documentation provides verifiable evidence of such commitment ex post.²⁴⁸ This bonding mechanism reduces client uncertainty and justifies premium pricing for FGIS-compliant services.

This documentation proves compliance with multiple duties. First, it demonstrates Rule 1.1 competence by showing lawyers “have a reasonable understanding of the capabilities and limitations” of the GPAI.²⁴⁹ Second, documentation reduces monitoring costs for all stakeholders: clients can verify lawyer contributions, disciplinary authorities can assess ethical compliance, and courts can evaluate professional conduct.²⁵⁰ By making previously hidden information observable, documentation reduces agency costs and improves overall market efficiency. Third, as courts become hypersensitive to unverified GPAI content, workpapers become essential evidence that the lawyer actively applied judgment, fulfilling duties under Rule 3.3 and earning the fee charged.²⁵¹ Documentation also addresses quality control challenges. A Stanford HAI study found legal GPAI tools from LexisNexis and Westlaw “each hallucinate between 17% and 33% of the time,” with errors including misunderstood holdings, false attribution, and fabricated information.²⁵² Systematic documentation of outputs and corrections allows firms to identify error patterns and continuously refine verification protocols. This transforms documentation from defensive compliance into affirmative quality management.

F. THE ECONOMIC TRANSITION: FROM BILLABLE HOURS TO VALUE-BASED PRICING

These three principles are necessary for ethical GPAI use under FGIS but are insufficient to resolve the tension between GPAI efficiency and firm profitability.²⁵³ That requires fundamental economic model transformation. The billable hour must yield to pricing structures that align firm incentives with client value and reward FGIS investments.²⁵⁴

The economic case for this transition rests on several foundations:

248. See PARISI, *supra* note 58, at 72.

249. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18, at 2–3.

250. See generally *id.* (discussing the numerous ethical impacts of AI use under the current Model Rules of Professional Conduct).

251. See *supra* Part I.

252. Magesh et al., *supra* note 4, at 216.

253. William Josten, *Pricing AI-Driven Legal Services: Alternative Fee Arrangements Are Almost Inevitable*, THOMSON REUTERS (June 3, 2024), <https://www.thomsonreuters.com/en-us/posts/legal/pricing-ai-driven-legal-services-alternative-fee-arrangements> [https://perma.cc/FFK2-29JE] (arguing that the current billable hour fee structure will fall apart because the price of services depends on the hours and resources put into it, rather than the value of the outcome to the client).

254. See *id.*

Transaction Cost Economics: The billable hour imposes substantial transaction costs on both firms and clients.²⁵⁵ Clients must audit bills and monitor time entries.²⁵⁶ Firms must track time meticulously, justify billing decisions, and manage client disputes.²⁵⁷ These activities consume resources without creating value. Value-based pricing reduces these transaction costs by shifting focus from inputs to outputs, enabling both parties to deploy resources more productively.²⁵⁸

Incentive Theory: As discussed throughout this section, the billable hour creates systematic incentive misalignment. Value-based pricing realigns incentives by making firm profitability dependent on efficiently delivering high-value outcomes. This is particularly crucial in the GPAI context, where efficiency gains are dramatic and the temptation for hour inflation is correspondingly strong.²⁵⁹

Competitive Dynamics: Markets tend toward efficient pricing mechanisms through competitive pressure. As more firms adopt GPAI and realize efficiency gains, competitive pressure will force transparent pricing that shares these gains with clients.²⁶⁰ Firms that cling to hourly billing while using GPAI will face adverse selection: sophisticated clients will migrate to competitors offering value-based pricing, leaving hourly-billing firms with less sophisticated clients or matters where AI provides minimal advantage.²⁶¹

Risk Allocation: Value-based pricing enables more efficient risk allocation. Under hourly billing, clients bear the risk of matter complexity and firm inefficiency.²⁶² Under fixed-fee arrangements, firms bear these risks and are incentivized to manage them through better technology, processes, and expertise.²⁶³ This risk reallocation is efficient when firms have better information and capabilities to manage such risks—precisely the situation with GPAI deployment.

This transformation is underway where work permits predictability. Transactional practices with standardized documentation suit fixed fee or

255. See Susan Saab Fortney, *The Billable Hours Derby: Empirical Data on the Problems and Pressure Points*, 33 FORDHAM URB. L.J. 171 (2005).

256. See *id.* at 181.

257. See *id.*

258. See *Value-Based Pricing: Aligning the Cost and Value of Legal Services*, DELOITTE, <https://www.deloitte.com/us/en/services/tax/articles/value-pricing-legal-services.html> [<https://perma.cc/Q535-U54V>] (last visited Apr. 13, 2026).

259. See discussion *supra* Section VII.C.

260. See discussion *supra* Section VII.A.

261. See discussion *supra* Section VII.A.

262. See Fortney, *supra* note 255, at 181 (establishing that under the billable hour, clients must often audit and scrutinize billing to ensure that it is being done fairly and the firm is using its time on the matter in an efficient manner).

263. See *supra* Section VII.D.

tiered pricing.²⁶⁴ Economic theory predicts that value-based pricing will first emerge in legal markets with the following characteristics: (1) low variance in matter complexity, enabling predictable pricing; (2) high potential for AI-driven efficiency gains; (3) sophisticated clients who can evaluate value propositions; and (4) competitive markets where firms must differentiate through pricing innovation. These conditions increasingly characterize large portions of the legal services market.

A task requiring twenty hours historically might take five hours of higher-value work. A fixed fee reflecting twelve hours of traditional work captures increased margin while significantly discounting client costs. The firm's FGIS investment is economically rewarded. The AI-driven productivity improvement creates economic surplus (the gap between five hours of cost and twenty hours of traditional billing). The firm captures part of this surplus (twelve minus five equals seven hours of pure margin), and the client captures part (paying twelve instead of twenty). Both parties benefit from the productivity improvement, creating stable economic equilibrium.

Regulatory compliance presents similar opportunities. Firms implementing FGIS safeguards—treating platforms like SALI and UTBMS as supervised nonlawyer assistants, ensuring economic neutrality through flat subscriptions, maintaining fiduciary-grade confidentiality through private deployment—can offer subscription-based pricing providing predictable client costs while serving more clients with existing staff.²⁶⁵ For example, a basic “Environmental Compliance Package” might include quarterly updates, annual filings, and routine consultations for fixed monthly fees, with premium tiers for complex strategic counseling requiring deeper human expertise. This aligns perfectly with FGIS economics: upfront governance investment creates platforms for efficient standardized service delivery while preserving attorney capacity for high-value, judgment-intensive work justifying premium pricing.

Even litigation is adapting quickly. AI-powered document review platforms like Everlaw have demonstrated dramatic efficiency: one Am Law

264. Platforms such as Westlaw and LexisNexis already offer robust libraries of standardized contract templated for various transactions. Integrating GPAI with these databases allows swift population of templates by accurately inserting party-specific details, dates, and deal-specific conditions, significantly reducing preparation time and minimizing errors. See Bob Ambrogi, *It's the Battle of the AI Legal Assistants, As LexisNexis Unveils Its New Protégé and Thomson Reuters Rolls Out CoCounsel 2.0*, LAWSITES (Aug. 12, 2024), <https://www.lawnext.com/2024/08/its-the-battle-of-the-ai-legal-assistants-as-lexisnexis-unveils-its-new-protége-and-thomson-reuters-rolls-out-cocounsel-2-0.html> [<https://perma.cc/6QAB-ZP9D>].

265. See generally Patrick Fuller, *The Law Office of the Future: Adapting to and Incorporating Technology While Keeping the Human Touch*, N.J. LAW., Feb. 2025, at 12, 15 (“Practices with defined rule bases, such as tax, employment, and compliance, are highly conducive to subscription offerings.”).

100 firm processed 126,000 documents with fifty to sixty-seven percent time reduction using only one-quarter typical staffing while maintaining ninety percent plus precision.²⁶⁶ Troutman Pepper Locke reduced time and costs by over forty percent and time to completion by over fifty percent in a \$20 million litigation case.²⁶⁷ These efficiencies enable phase-based or task-based litigation pricing. Firms can offer fixed fees for discrete tasks where FGIS-compliant GPAI enables predictability while reserving hourly or hybrid billing for inherently variable phases like depositions or trial.

The economic logic is compelling: firms adopting FGIS-compliant GPAI and transitioning to value-based pricing simultaneously increase profitability and decrease client costs, creating genuine economic surplus.²⁶⁸ The firm completing in five rigorous FGIS-compliant hours what previously took twenty, charging a fixed fee reflecting twelve hours of value, more than doubles its effective hourly rate while saving clients forty percent. This is genuine productivity improvement benefiting all stakeholders, made possible by robust GPAI governance investment and willingness to capture value rather than merely sell time.

G. MARKET STRUCTURE CONSIDERATIONS

The transition to value-based pricing and FGIS-compliant AI usage will not occur uniformly across the legal services market. Market structure considerations will significantly affect adoption rates and patterns:

BigLaw vs. Small Firms: Large law firms have greater resources to invest in FGIS infrastructure and AI technologies, potentially creating competitive advantages. However, they also face greater organizational inertia and more entrenched billable-hour cultures.²⁶⁹ Small firms may be nimbler in adopting value-based pricing but lack resources for comprehensive FGIS implementation. This suggests market segmentation where different firm sizes optimize different strategies.

Client Sophistication: Sophisticated corporate clients with in-house counsel can better evaluate AI usage and value propositions, creating

266. See Petra Pasternak, *Am Law 100 Firm Slashed Doc Review Time by Two-Thirds With GenAI*, EVERLAW (Apr. 17, 2025), <https://www.everlaw.com/blog/case-studies/am-law-100-firm-slashed-doc-review-time-by-two-thirds-with-genai> [https://perma.cc/SHG3-48EX].

267. See Kristy Esparza, *Harnessing Generative AI: 4 Law Firms Share Lessons Learned*, RELATIVITY BLOG (Oct. 23, 2024), <https://www.relativity.com/blog/harnessing-generative-ai-4-law-firms-share-lessons-learned/> [https://perma.cc/3VF8-PJ2K].

268. See Adams & Matheson, *supra* note 50, at 41.

269. See generally Robert J. Couture, *The Impact of Artificial Intelligence on Law Firms' Business Models*, HARV. L. SCH.: CTR. ON LEGAL PRO. (Feb. 24, 2025), <https://clp.law.harvard.edu/knowledge-hub/insights/the-impact-of-artificial-intelligence-on-law-law-firms-business-models/> [https://perma.cc/EK78-RYFG].

pressure for transparency and value-based pricing.²⁷⁰ Less sophisticated clients may lack the information or expertise to demand such changes, potentially creating a two-tiered market where pricing innovation occurs primarily in serving sophisticated clients.

Practice Area Dynamics: Some practice areas (transactional, regulatory compliance) lend themselves more naturally to value-based pricing due to matter predictability.²⁷¹ Others (complex litigation, appellate work) involve greater uncertainty, making value-based pricing more challenging. Market evolution will likely proceed at different rates across practice areas.

Network Effects and Standards: As more firms adopt FGIS-type frameworks, network effects may emerge. Vendors will be incentivized to offer FGIS-compatible services if enough firms demand them.²⁷² Clients will come to expect FGIS-level protections as they become industry standard. These network effects can accelerate adoption once a critical mass is reached but create barriers to early adoption before such mass exists.²⁷³

Regulatory Pressure: Bar associations and courts are increasingly requiring disclosure and oversight of AI usage.²⁷⁴ This regulatory pressure will tend to level the playing field, requiring even reluctant firms to adopt governance frameworks. However, the extent and timing of such pressure will vary by jurisdiction, creating geographic variation in adoption patterns.²⁷⁵

These market structure considerations suggest that the transition described in this Article will be gradual and uneven. However, the economic logic—that aligned incentives create stable equilibrium while misaligned incentives create unsustainable tensions—suggests that value-based pricing

270. See, e.g., Josten, *supra* note 253 (“In the Thomson Reuters Institute’s . . . Report, we examined how corporate general counsel have been working to control their costs and found that nearly half of GCs surveyed cited an increasing preference for AFAs. . . . many law firms have built up robust teams of experienced professionals to support more creative and profitable AFAs that meet the goals of both law firms and their clients.”); see Adrian D’Amico, *AI-Ready Billing: Rethinking Legal Pricing in the Age of Automation*, FENNEMORE (Aug. 5, 2025), <https://www.fennemorelaw.com/ai-ready-billing-rethinking-legal-pricing-in-the-age-of-automation/> [<https://perma.cc/QK2M-8ALQ>]; *Value-Based Pricing*, *supra* note 258.

271. See Medianik, *supra* note 243.

272. See, e.g., SGS, *AI Governance in the Legal Industry* 13 (2025), <https://www.sgs.com/-/media/sgscorp/documents/corporate/white-papers/sgs-ba-ai-governance-in-the-legal-industry-en.cdn.en.pdf> [<https://perma.cc/N9EE-S5GR>].

273. See INT’L CHAMBER COM., *AI GOVERNANCE AND STANDARDS: POLICY PAPER* (2025), <https://iccwbo.org/wp-content/uploads/sites/3/2025/07/2025-ICC-Policy-Paper-AI-governance-and-standards.pdf> [<https://perma.cc/48UM-846M>] (positing that industry-wide standards accelerate adoption and innovation so that use of services remain applicable to the changing real-world scenarios).

274. See A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, *supra* note 18.

275. See, e.g., *AI and Attorney Ethics Rules: 50-State Survey*, *supra* note 199.

coupled with robust AI governance will eventually become the dominant model in most segments of the legal services market.

H. THE PATH FORWARD

Aligning disclosure, billing, and documentation with FGIS-compliant GPAI use is a strategic imperative determining which firms thrive. Firms failing to adapt face impossible choices: underutilize GPAI to protect billable hours and become uncompetitive, or use GPAI extensively with traditional hourly billing, inviting ethical violations and client backlash when efficiency gains become apparent.

The path forward requires institutional courage: acknowledging the billable hour model may not optimally serve the next generation; transparency with clients about GPAI and governance frameworks; substantial investment in FGIS infrastructure; and fundamental rethinking of performance measurement, shifting from hours generated to value delivered.

The encouraging reality is that the ethical framework already supports this evolution. GPAI is not creating new obligations; it makes longstanding ones impossible to ignore. The firm using FGIS-compliant GPAI to complete work in five rigorous hours, billing twenty based on historical norms, with no disclosure violates longstanding prohibitions on unreasonable fees and dishonest communication. Conversely, the firm implementing comprehensive FGIS, documenting verification in auditable logs, transparently disclosing GPAI use and governance, and pricing based on value delivered simply fulfills obligations that have always defined competent, honest, client-centered practice.

By aligning disclosure, billing, and documentation with FGIS-compliant GPAI realities, firms ensure their economic model reinforces rather than undermines ethical obligations and competitive positioning. The alternative, clinging to the billable hour while GPAI erodes its foundations, is neither economically nor ethically sustainable. The question is whether firms adapt proactively and systematically, implementing robust governance like FGIS and restructuring economics to reward that investment, or reactively under pressure from client defections, bar discipline, malpractice claims, or sanctions. Firms choosing the former will thrive in the GPAI era, delivering superior client value, maintaining healthy margins, fulfilling professional responsibilities, and leading a transformed legal market. Those choosing the latter may find the billable hour has finally met a technology it cannot survive.

IX. OBJECTIONS AND REBUTTALS

The proposal of a restrictive standard like FGIS will inevitably encounter objections from those who view it as an overly cautious or burdensome impediment to innovation. However, these objections typically underestimate the unique nature of the risks posed by GPAI and the fundamental, non-negotiable character of a lawyer's fiduciary duties. By anticipating and addressing these counterarguments, the necessity and prudence of the FGIS standard becomes even clearer.

Objection 1: "We need access to frontier, general-purpose models to remain competitive; legal-specific models lag in capability, and FGIS is too restrictive."

Rebuttal: This objection misconstrues the nature of a lawyer's professional obligations. The duty of professional independence under Model Rule 5.4 is not a feature that can be traded for a higher-performing model; it is a categorical imperative. The FGIS standard does not forbid the use of frontier models. Rather, it requires that if a firm chooses to use such a powerful tool for substantive client work, it must do so within an environment of maximum control and safety; specifically, a private deployment with a no-training guarantee and a rigorous, auditable verification process. The core of the legal profession's value is trusted judgment, not raw computational power. If a "frontier" model is only available through a public-facing interface that compromises confidentiality and independence, then a fiduciary's duty is to refuse to use it for substantive client matters, opting instead for a less powerful but fully controllable alternative. Professionalism demands that ethics govern technology, not the other way around.

Objection 2: "Clients are demanding greater speed and lower costs, and the procedural hurdles of FGIS are too slow and expensive."

Rebuttal: This is a false dichotomy. FGIS does not forbid automation; it forbids *uncontrolled* automation. The true source of sustainable efficiency and client value lies in delivering high-quality, reliable, and verifiably accurate work product in a predictable and transparent manner. Cutting corners on confidentiality and independence is a short-sighted strategy that invites far greater costs down the road in the form of sanctions, malpractice claims, fee disputes, and reputational damage. The costs associated with a single waived privilege or a court-imposed sanction for submitting unverified work will dwarf any marginal savings from using a cheaper, public-facing AI tool. Furthermore, the FGIS framework, by encouraging a shift to Alternative Fee Arrangements, directly addresses clients' demands for cost

predictability.²⁷⁶ A transparently priced fixed fee for a well-defined project, completed efficiently using a secure and verified AI workflow, aligns firm and client interests. FGIS is not a barrier to efficiency; it is the blueprint for achieving it responsibly.

Objection 3: “The current ethics opinions from the ABA and major states do not explicitly require a default ban or mandate private deployments. FGIS goes too far.”

Rebuttal: This objection confuses the floor with the ceiling. Ethics opinions articulate the *minimum* standards of conduct necessary to avoid discipline; they do not define the full scope of a lawyer’s fiduciary duty to act in the client’s best interest. While it is true that no opinion has yet used the precise language of a “default ban,” the clear and unmistakable *trend* of all recent guidance runs in one direction: toward stricter controls on data, greater demands for verification, and explicit warnings against feeding confidential information to public, third-party tools. The Florida Bar’s suggestion to use in-house AI to mitigate risk and the California judiciary’s prohibition on putting nonpublic data into public systems are two examples of the direction in which the rules are evolving.²⁷⁷ Adopting the FGIS standard now is not an act of over-compliance but one of prudent risk management. It is far wiser and more cost-effective to build a robust, conservative governance framework today than to be forced into a reactive, expensive retrofit after the rules have become even more stringent—or worse, after a catastrophic ethical failure has already occurred.

X. CONCLUSION

The integration of GPAI into the practice of law presents a challenge unlike anything the profession has previously faced. GPAI is not merely a more efficient tool; it is a new form of collaboration with a non-human, commercially motivated entity that operates with a degree of opacity that is fundamentally at odds with the legal profession’s core values. To focus solely on the risk of AI “hallucinations” is to miss the forest for the trees. The more profound, systemic dangers lie in the subtle erosion of the lawyer’s independent judgment and the waiver of the client’s most sacred protections: confidentiality and privilege.

A profession whose primary asset is trusted, independent judgment cannot ethically outsource that judgment to an algorithmic “black box” controlled by a third-party vendor. A profession bound by a fiduciary duty to protect client secrets at every peril cannot justify feeding those secrets into

276. See discussion *supra* Section VII.D.

277. See *supra* Section II.B.

public systems that reserve the right to use them for commercial gain. The bargain that lawyers strike with clients and with the courts is predicated on these foundational principles of independence and confidentiality. Uncontrolled use of GPAI breaks that bargain.

The FGIS standard proposed in this Article offers a path forward. It is not a Luddite rejection of technology but a robust, fiduciary-driven framework for its responsible adoption. By establishing a default prohibition on the use of GPAI for substantive work, the standard forces a conscious, deliberate, and documented decision to proceed only when specific, auditable safeguards are in place. The twin pillars of the standard are (1) Independence Safeguards that neutralize vendor influence and (2) Fiduciary-Grade Confidentiality Controls that protect client data. Both pillars are a direct operationalization of the duties already enshrined in the Model Rules of Professional Conduct.

The choice before the legal profession is not between innovation and stagnation. It is between responsible automation that augments human judgment and a reckless abdication of professional duty that replaces it. The FGIS standard provides a clear, bright-line rule that honors this distinction: if a law firm cannot prove, with tangible evidence, that its use of GPAI preserves professional independence and guarantees fiduciary-grade confidentiality, it must not use it on a client's matter. This is not an anti-technology stance; it is the only stance a true fiduciary can take. It is how a profession absorbs the immense power of automation without surrendering its soul.

From a law and economics perspective, the integration of GPAI into legal practice presents a paradigm case of technological disruption creating market failures that require regulatory intervention. The core economic problems are threefold: (1) severe information asymmetries between clients, lawyers, and vendors; (2) principal-agent problems compounded by third-party influence; and (3) misaligned incentives embedded in both AI vendor business models and traditional law firm pricing structures.

These market failures will not self-correct through competitive pressure alone. Information asymmetries are too severe, agency costs too high, and switching costs too substantial to enable efficient market responses without regulatory frameworks. The FGIS standard proposed in this Article represents necessary regulatory intervention to address systematic market failure.

However, sustainable regulation must align economic incentives with ethical obligations. This Article's economic analysis demonstrates that when firms adopt comprehensive AI governance frameworks (FGIS) and transition to value-based pricing models (AFAs), ethical compliance becomes

economically rewarding, rather than costly. This alignment creates stable market equilibrium where profit-maximizing behavior coincides with ethical behavior—the optimal outcome from both economic efficiency and professional responsibility perspectives.

Thus, the choice before the legal profession is fundamentally economic as well as ethical: will the market evolve toward efficient, incentive-aligned structures that benefit all parties, or will it remain trapped in unsustainable models that create systematic pressure for ethical violations? The answer will determine not only the profession's ethical integrity, but also its economic viability in an AI-transformed market.