

NAVIGATING AI'S WILD WEST

Trends, Ethics, and Backlash in Legal Practice

Mason Clark

Assistant Professor of Law, Englehardt Research Fellow

St. Mary's University School of Law

Mclarkr16@stmarytx.edu

2026 National Conference on Special Needs Planning
and Special Needs Trusts

Thursday, October 22, 2026

I. INTRODUCTION AND SCOPE

These materials¹ accompany the breakout session “Navigating AI’s Wild West: Trends, Ethics, and Backlash in Legal Practice.” They address the professional and ethical dimension of artificial intelligence in trusts and estates practice and the legal profession, broadly. It explores the current tool landscape, the duties implicated under the Rules of Professional Conduct, the particular problem of bias against beneficiaries with disabilities, and the reasons confidence in AI-assisted work product is eroding among judges, clients, and counsel.

The technical reliability of these tools, the taxonomy of tasks for which they are suited, the recurring substantive errors they produce in special needs matters, and the verification protocol that responds to those errors are treated in the companion materials prepared for the morning session, “Are the Machines in Charge? Using AI to Review Trusts and Other Things.” Those subjects are not repeated here.

¹ As many jurisdictions now require attorneys to provide disclosure and/or verification statements for documents prepared with assistance from or containing information generated by generative artificial intelligence, it seems fitting to include a specific and conspicuous disclosure and verification statement here: These materials were prepared with the assistance of generative artificial intelligence. The author used various large language models to organize some parts and subsections; to rephrase some content for brevity; to format citations; to summarize lengthy reports or articles, particularly because they pertain to an area of the law that is NOT the author’s area of expertise; and to compile research material. The author certifies that all citations have been manually verified to ensure their existence and confirm their propositions as accurate representations.

However, in Part V.C, the author intentionally included unverified citations to certain changes in the law impacting conference participants. The author invites participants to find incorrect AI outputs!

One connection between the two sets of materials should be stated at the outset, because it organizes much of what follows. A verification protocol of the kind described in the morning materials checks whether the output is accurate. It does not check whether the output is fair. The failures discussed in Part IV below would ostensibly pass a citation our output verification checklist.

II. THE CURRENT LANDSCAPE

Any inventory of specific products has a shelf life measured in weeks. What follows is therefore organized around categories and evaluation criteria rather than recommendations.

A. Categories of Tools

1. General-purpose frontier models accessed through consumer or enterprise interfaces.
2. Legal research platforms with integrated AI, grounded in a curated database of primary authority.
3. Contract and document review tools, generally oriented toward transactional review.
4. Trusts-and-estates-specific platforms performing estate plan summarization, visualization, and intake.
5. Firm-deployed internal tools built on enterprise interfaces, typically with contractual protection against training on inputs.
6. Agentic tools that take multi-step actions rather than answering a question, which raise supervision questions of a different order.

B. Adoption

Practitioners assessing where the profession actually stands should consult the ABA Legal Technology Survey Report, the Clio Legal Trends Report, and the Thomson Reuters Future of Professionals report, each updated annually. The most informative figure in these surveys is the gap between firm-sanctioned use and actual use.

C. Vendor Diligence

Before any tool touches client data, the following questions should be answered from the terms of service rather than from marketing materials: Does the vendor train on user inputs? Where is data stored, and for how long? Are terms negotiable? What becomes of prompts and outputs in litigation or on subpoena? Does the tool cite to source, and can the citation be opened? Does the contractual language match the sales representation?

D. Shadow AI

The most common deployment in law firms and trust departments is not a firm-procured platform. It is an associate, paralegal, or trust administrator using a personal consumer account on a personal device. Any policy that does not account for this is not a policy. It is a document.

III. PROFESSIONAL RESPONSIBILITY

The starting point for any practitioner is ABA Formal Opinion 512 (July 29, 2024), which addresses generative AI across the range of implicated duties. State authority has developed alongside it, and practitioners must consult their own

jurisdiction's guidance; opinions and task force reports have issued in a substantial number of states, and several state supreme courts have adopted interim rules or standing requirements.

A. Competence

Model Rule 1.1, comment 8, requires a lawyer to keep abreast of the benefits and risks associated with relevant technology. The duty runs in both directions. A lawyer who uses a tool she does not understand is not competent; and an emerging argument holds that a lawyer who declines available tools that would materially improve accuracy or reduce cost may likewise fall short. The second proposition is contested and should be presented as such.

B. Confidentiality

Model Rule 1.6 presents the sharpest issue in this practice area, because special needs files contain an unusual concentration of sensitive material: diagnoses, psychiatric and substance use history, capacity evaluations, sealed guardianship records, school records, and information concerning minors.

Three points warrant particular attention. First, consumer and enterprise terms differ materially on whether inputs are used for training, and the consumer default is frequently the less protective. Second, informed consent under Rule 1.6(a) requires that the client understand what is being disclosed and to whom; a general engagement letter recital that the firm “uses technology” is unlikely to suffice. Third, and most difficult, where the trustee is the client, the most sensitive information in the file frequently belongs to a beneficiary who is not the client, who

cannot consent, and whose interests in that information are not represented in the consent conversation.

C. Supervision

Model Rules 5.1 and 5.3 supply the governing analogy: the tool is treated as a nonlawyer assistant whose work must be supervised and whose output the lawyer remains responsible for. The obligation extends to the conduct of the firm's staff in their use of these tools, which is where most unsupervised use occurs.

This analysis should be extended beyond the lawyers in the room. Pooled trust administrators, corporate trust officers, and care managers make consequential decisions with AI assistance and are subject to no analogue of Rules 5.1 and 5.3. They are frequently the individuals handling the most sensitive information in the file.

D. Communication and Disclosure

Whether use must be disclosed to the client under Model Rule 1.4 depends on the nature of the use. Opinion 512's treatment is more nuanced than either the disclose-always or the disclose-never position, and the variables it identifies — whether the client has asked, whether the use bears on fees, whether confidential information is input, whether the output is material to the representation — are what practitioners should be applying.

E. Fees

Model Rule 1.5 prohibits billing for time not expended. A task that formerly took three hours and now takes twenty minutes may be billed at twenty minutes.

Whether the efficiency gain belongs to the client is a question firms should resolve deliberately and disclose in the engagement letter rather than resolve by default. Costs of AI tools passed through to clients should be handled as any other disbursement.

F. Candor and Unauthorized Practice

Fabricated citations implicate Model Rules 3.1 and 3.3, and the reported decisions treat them as misconduct rather than as embarrassment.²Separately, Model Rule 5.5 and state unauthorized practice statutes are increasingly implicated by AI-generated trust instruments produced by nonlawyer fiduciaries and consumer platforms.

IV. BIAS AND VULNERABLE BENEFICIARIES

The general observation that AI systems encode bias is familiar and not particularly useful. The disability-specific bias is concrete and directly relevant to this practice.

A. Mechanism

These models were not trained principally on writing produced by disabled adults describing their own lives. They were trained on the far larger available corpus of clinical literature, caregiver material, risk-management guidance, and institutional policy—text produced about disabled people by the people and

² See *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023); *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024).

institutions managing them. A model reproduces the stance that dominates its source material.

B. Manifestations in Practice

Asked to analyze a distribution request, a model will reliably produce output exhibiting a recognizable set of features: deficit framing and concessive constructions; treatment of a competent adult beneficiary's stated preference as a data point requiring family corroboration; the introduction of a family member as a de facto decision-maker where no guardianship exists; risk analyses that typically only count the risks of acting and rarely the costs of denial; staged, conditional, and revocable approvals for ordinary adult expenditures; and monitoring conditions attached to routine purchases.

The corresponding drafting risk, when these tools are used to generate distribution standards or trustee guidance, is the importation of paternalistic assumptions about what a beneficiary should be permitted to spend money on.

C. Historical Precedent

This population has encountered opaque automated decision-making before, in the benefits context, with documented results.³

³ See *K.W. v. Armstrong*, No. 1:12-cv-00022-BLW (D. Idaho) (challenge to an undisclosed algorithmic budget tool used to reduce home and community-based services, in which the state asserted the formula was a trade secret); *Arkansas Department of Human Services v. Ledgerwood*, 530 S.W.3d 336 (Ark. 2017) (litigation arising from algorithmic reallocation of in-home care hours).

Today, the technology is more sophisticated and with increased potential for automated judgments. The distinction now is that algorithmic decision-making systems have moved into the practitioner's and the trustee's own workflow, where it is subject to no due process challenge at all because it never issues a decision. It only advises.

D. Detection

Two practices are worth adopting. First, ask who is absent from the analysis; where the output describes a person who was never consulted, that is the flag. Second, apply a substitution test: replace the disabled beneficiary with a non-disabled person of the same age and resources and re-read the analysis. Where the analysis changes and nothing about the request was disability-related, the difference is bias rather than prudence. Improved prompting mitigates this and does not cure it.

V. THE FIDUCIARY OVERLAY

Trustee use of AI raises duties that exist independently of the Rules of Professional Conduct and that apply to corporate and pooled trust fiduciaries who are subject to no professional responsibility regime at all.

1. Prudent administration. Uniform Trust Code § 804 requires administration as a prudent person would, exercising reasonable care, skill, and caution. A trustee who adopts AI-generated reasoning without independent evaluation has arguably not satisfied that standard.

2. Delegation. Uniform Trust Code § 807 permits delegation of functions a prudent trustee of comparable skills could properly delegate, with duties of care in selecting the agent, establishing the scope of the delegation, and monitoring performance. Whether the core discretionary distribution decision is delegable at all is a serious question; delegation to a system that cannot explain its reasoning, cannot be examined,⁴ and produces different output on different runs satisfies none of the three monitoring obligations.
3. Impartiality and the duty to inquire. A trustee must consider the circumstances of the beneficiary. Analysis that substitutes a statistical profile of a diagnostic category for inquiry into the actual beneficiary does not discharge that duty.
4. Records and accountings. Where the reasoning supporting a denial exists only as a chat log the trustee did not write, the trustee's ability to defend the exercise of discretion on an accounting is materially impaired.
5. Abuse of discretion. Even under sole and absolute discretion, a trustee must act in good faith and consistently with the purposes of the trust. See Restatement (Third) of Trusts § 87. A supplemental needs trust exists to enhance quality of life beyond what public benefits supply,

⁴ Lou Blouin, *AI's Mysterious "Black Box" Problem, Explained*, U. Mich.-Dearborn News (Mar. 6, 2023), <https://umdearborn.edu/news/ais-mysterious-black-box-problem-explained>.

and a pattern of denials generated by a system disposed toward restriction is difficult to reconcile with that purpose.

VI. WHY TRUST IN THESE TOOLS IS FRAYING

Skepticism about AI-assisted work product is not irrational, and practitioners who treat it as mere technophobia will handle client conversations badly.

A. The Sanctions Record

Beginning with *Mata* and *Park v. Kim*, the volume of reported decisions involving fabricated AI-generated authority has grown substantially. Because the count changes continuously, practitioners should consult a maintained database rather than any static list; the AI Hallucination Cases database compiled by Damien Charlotin⁵ is updated on an ongoing basis and is the most useful current source.

B. Judicial Response

Courts have responded with standing orders requiring disclosure or certification of AI use. These orders are inconsistent across and within jurisdictions, and compliance requires checking the individual judge's requirements. The reputational consequence is asymmetric: a single fabricated citation attaches to a lawyer's name permanently and is trivially discoverable by opposing counsel and clients.

⁵ Damien Charlotin, *AI Hallucination Cases*, <https://www.damiencharlotin.com/hallucinations/> (last updated Sept. 23, 2026).

C. Institutional and Client Pushback

Outside counsel guidelines increasingly restrict or prohibit AI use; professional liability carriers have begun asking about it at renewal; and corporate trustees have adopted vendor policies governing what their outside counsel may use.

D. Cultural Backlash

Practitioners should understand the broader challenges to the proliferation of AI. This includes pending copyright litigation, objections concerning energy consumption, concern about labor displacement, and a general fatigue with low-quality machine-generated content (colloquially known as “AI slop”).⁶ A client’s skepticism frequently originates outside the legal context entirely.

E. Disability Community Distrust

This deserves to be named separately rather than folded into general skepticism. It has its own history, described in Part IV.C, and it is well founded.

VII. A FRAMEWORK FOR WEIGHING USE

A workable internal policy sorts uses along two axes: the sensitivity of the data involved and the reversibility of an undetected error.

1. Prohibited. Client-identifying or beneficiary health information in consumer tools; any use generating a legal conclusion relied upon

⁶ For an excellent law review article on “AI slop,” see Jessica Silbey & Woodrow Hartzog, *AI Slop* (Sept. 2026) (unpublished manuscript), https://scholarship.law.bu.edu/faculty_scholarship/4308.

without independent verification; delegation of discretionary distribution decisions.

2. Permitted with mandatory verification. Document review, summarization, cross-document consistency checking, and drafting assistance, subject to the verification protocol in the morning materials and to a documented file note.
3. Unrestricted. Tasks involving no client data and no legal conclusion — internal research orientation, formatting, general drafting of non-client material.

The policy should address staff and non-lawyer fiduciaries explicitly, should state what is permitted rather than only what is forbidden, and should be written on the assumption that personal-device use will occur unless a sanctioned alternative is genuinely easier to use.

VIII. “DID YOU WRITE THIS WITH AI?”

Practitioners should prepare an answer to this question before it is asked. The answer should not be defensive, and it should not overclaim. In most cases the accurate answer is that the tool was used for a specified task, that the output was independently reviewed, and that the professional judgment is the lawyer’s own.

Familiar analogies—form books, document assembly systems, a junior associate’s first draft—are incomplete. Each differs in that the intermediate work product remains inspectable and the human assistant can be asked to explain a

choice. Acknowledging where the analogy fails is considerably more persuasive to a skeptical listener than insisting that it holds.

Two further points are worth consideration. First, evaluating potential AI use requires you to ask if a machine made the decisions, or if a machine just touched the document. Clients almost always mean the first and almost always ask the second, and separating them usually resolves the concern. Second, the appropriate answer differs by audience. A skeptical adult child of a beneficiary, opposing counsel in a surcharge proceeding, and a judge operating under a standing order are asking for different things and are entitled to different responses.

IX. CONCLUSION

The competence question and the fairness question are distinct, and the second is the harder of the two. A practitioner who verifies every citation, confirms every statutory reference, and documents the review has addressed accuracy. She has not addressed whether the framing she adopted was one she would have chosen. In a practice serving beneficiaries who are frequently unable to detect or contest how they are being characterized, that second question is the one that ultimately governs whether these tools are used well.