

First-Year Harvard Law Student's Anticipation of their Early Practice AI Use.

Queen Eche

INTRODUCTION

Since its introduction in 2022, Generative Artificial Intelligence (“AI”) has reshaped nearly every major industry in the United States.¹ Until recently, the legal profession has been among the slower sectors to adopt it.² There are two complementary reasons for this gap. First, older generations of lawyers have been wary of AI due to confidentiality and professional responsibility concerns,³ and second, technology companies have struggled to gain the credibility needed to innovate for such a resistant audience. However, that resistance is now giving way. Harvard Law Professor David Wilkins warns that AI “has the potential to transform the practice of law.”⁴ In fact, AI-driven structural pressures are prime to reshape the market. There is speculation that in-house clients will demand AI-enabled efficiency gains from large law firms, which in turn will compress hours billed by junior associates and shrink the margin for traditional on-the-job training.⁵ If true, nowhere is this more

¹ See Luona Lin, *About 1 in 5 U.S. Workers Now Use AI in Their Job, Up Since Last Year*, Pew Rsch. Ctr. (Oct. 6, 2025), <https://www.pewresearch.org/short-reads/2025/10/06/about-1-in-5-us-workers-now-use-ai-in-their-job-up-since-last-year>.

² See *Is the Law Playing Catch-Up with AI?*, Harv. L. Today (Jan. 16, 2025), <https://hls.harvard.edu/today/is-the-law-playing-catch-up-with-ai/>.

³ *The AI Adoption Curve in Law*, Best L. Firms (Dec. 26, 2025), <https://www.bestlawfirms.com/articles/the-ai-adoption-curve-in-law/7196>; *AI Adoption Is Growing, but Some Are Hesitant*, New ABA Tech Survey Finds, A.B.A. J. (Mar. 7, 2025), <https://www.abajournal.com/web/article/aba-tech-report-finds-that-ai-adoption-is-growing-but-some-are-hesitant>.

⁴ Jeff Neal, *Harvard Law Expert Explains How AI May Transform the Legal Profession in 2024*, Harv. L. Today (Feb. 14, 2024), <https://hls.harvard.edu/today/harvard-law-expert-explains-how-ai-may-transform-the-legal-profession-in-2024/>.

⁵ *Power Rankings, AI, and the New Big Law Tournament*, AI and the Future of Law (Am. Arb. Ass'n Apr. 28, 2026), <https://www.adr.org/podcasts/ai-and-the-future-of-law/power-rankings-ai-and-the-new-big-law-tournament/>; see Bjarne P. Tellmann, *GC 4.0 — The AI Era Legal Department*, in *Law in the Era of AI* 79 (2026).

consequential than for current law students, who are the first to face this reality having trained alongside AI before ever setting foot in a firm.

Within the Class of 2028 at Harvard Law School, AI use was pervasive during the first year of legal education. Students asked AI for doctrinal explanations, case brief creation, and real-time in-class clarifications. This class of students are not learning law in the absence of AI—they are learning alongside it. Yet despite growing speculation about how AI will reshape early legal careers, almost no empirical work has examined what prospective legal professionals themselves actually expect, want or fear from AI as they prepare to enter practice. This paper addresses that gap directly.

In April 2026, a survey was administered to Section 1 of Harvard Law School’s JD Class of 2028—students on the immediate threshold of their first summer associate positions—asking a question that existing scholarship has not answered: do law students entering elite practice intend to adopt AI tools in their early careers, and if so, for what purposes and under what conditions?

This report proceeds as follows. Part I presents a short literature review that surveys the existing scholarship on AI adoption within the legal industry. Part II describes the survey instrument, sample, and analytical approach. Part III establishes, across three parts, that AI adoption intent is broad but conditional: Part A examines whether and for what purposes respondents intend to adopt AI. Part B analyzes the hierarchy of professional wariness across five concerned dimensions. Finally, Part C identifies the conditions respondents demand as prerequisites for fuller adoption. Part IV concludes this report with a synthesis of these findings.

I. LITERATURE REVIEW

The majority of existing scholarship on AI in the legal field has largely focused on established legal professionals rather than students preparing to enter practice. In fact, Andreeva and Savova conducted the first known empirical study⁶ specifically targeting law students, surveying 54 students across multiple institutions on their AI usage, attitudes, and concern. Their findings reveal a student cohort meaningfully different from practicing lawyers. In comparing their results to a parallel study of senior attorneys,⁷ they found that while over 90% of firm partners expressed positive attitudes toward AI, only 50% of students shared that enthusiasm. Andreeva and Savova attribute this gap to differing levels of exposure to professional AI tools. Furthermore, Andreeva and Savova found that when students identified barriers to broader AI adoption in legal practice, accuracy and privacy emerged as primary concerns.⁸

However, Andreeva and Savova's study has two limitations that bear on its applicability here. First, the sample was recruited through a LinkedIn survey and drew responses from students in the United Kingdom, Nigeria, Pakistan, and India. As a result, the distribution's relevance is limited as to the United States legal market, where firm culture⁹ and AI governance frameworks differ materially.¹⁰ Second, the study treated professional wariness as a unified cluster rather than identifying the components that cause wariness, leaving open the question of which anxieties most powerfully condition adoption intent. This paper addresses both gaps: it surveys a cohort of students at an elite

⁶ Daniela Andreeva & Guergana Savova, *Artificial Intelligence in the Legal Field: Law Students' Perspective*, arXiv:2410.09937, 4–5 (Oct. 13, 2024), <https://arxiv.org/abs/2410.09937>.

⁷ Robert Mahari et al., *The Law and NLP: Bridging Disciplinary Disconnects*, in Findings of the Ass'n for Computational Linguistics: EMNLP 2023, at 3445 (Houda Bouamor et al. 2023).

⁸ Andreeva & Savova, *supra* note 6, at 5.

⁹ *Id.*

¹⁰ *Id.*

American law school who are at the cusp of entry into practice, and it measures wariness across five distinct professional concern dimensions to identify which concerns drive the conditional structure of adoption intent.

II. METHODOLOGY

This study utilized a survey, designed to collect data on AI adoption intentions among law students entering legal practice. The survey was administered within Section 1 of Harvard Law School's JD Class of 2028 between April 7 and April 10, 2026—a moment when students had completed their first semester of legal education and were either finished with or actively engaged in recruiting for their 2026 and 2027 summer internships. The students, many of whom are set to begin their 2026 summer internship in late May, were on the immediate threshold of professional practice.

Section 1 was selected as the sampling unit because Harvard Law School deliberately constructs each section to reflect demographic diversity among gender, race, geography, and undergraduate background. This reduces, though does not eliminate, the risk that the sample skews toward students with unusually homogenous characteristics.

The survey was distributed exclusively through the Section's online Slack channel. Slack was chosen as the medium of distribution because all 83 students of the section are members of the channel and the channel has served as a neutral communication space rather than one associated with any particular interest or viewpoint. Of the 83 students in Section 1, 29 responded, yielding a response rate of approximately 35%.

To further reduce self-selection bias, no description of the survey's subject matter was provided when students were invited to respond within the Section's slack channel. Respondents only

learned about the nature of the survey after clicking through to complete it. This was done in hopes of reducing the risk that students with unusually strong opinions about AI in either direction were disproportionately likely to participate.

The data was analyzed using descriptive statistics, and a comparative analysis of survey responses to test relationships between respondents' wariness and adoption intent.

The findings in this report are suggestive and descriptive rather than statistically generalizable, given the small sample size, the sample being from a single, elite institution, and the single point-in-time administration.

III. FINDINGS

This Part presents data collected from the survey and provides analysis and charts to illustrate the findings. Part A examines adoption intentions in two ways: first, whether respondents intend to adopt AI tools in their early legal careers, and second, for what purposes. Part B analyzes respondent's wariness of AI adoption across five professional concern dimensions. Part C examines the conditions respondents identified as prerequisites for fuller adoption.

A. AI Adoption Intentions.

1. Do They Intend to Adopt?

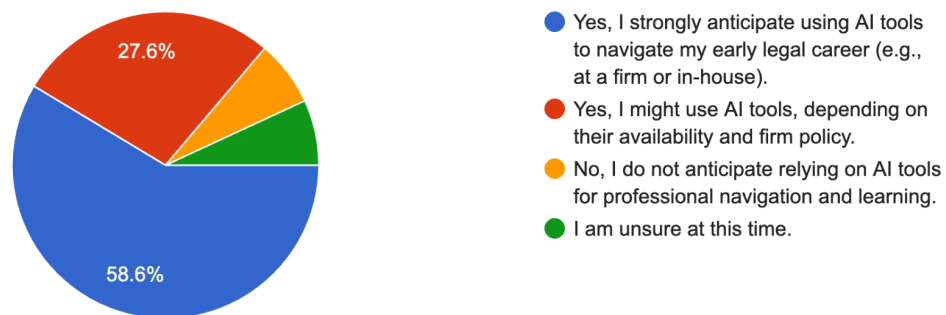
The data reveals broad but uneven adoption intent across the sample. As Table A.1.1 indicates, when asked to place themselves in one of four categories, 58.6% of respondents (17 of 29) selected "Yes, I strongly anticipate using AI tools to navigate my early legal career." 27.6% (8 of 29) of respondents selected "Yes, I might use AI tools, depending on their availability and firm policy." A small resistant minority—6.9% (2 of 29)—indicated they were unsure. Furthermore, an equally small

group of 6.9% (2 of 29) indicated they did not anticipate relying on AI tools at all. Taken together, more than four in five (25 of 29) respondents expressed some degree of adoption intent, with the majority expressing strong intent.

TABLE A.1.1. Respondents' Anticipation of Using AI

Your Anticipation of Using AI in Your Early Legal Career

29 responses




While such distribution seems straightforward on face, the numerical likelihood scores display a more complicated story. As shown in Table A.1.2., on a scale of 1 to 10 measuring how likely respondents were to use AI to help learn the practical skills of working at a law firm, the sample produced a mean score of 5.66 and a median of 6, with a standard deviation of 3.05. However, the total band of scores ranged from 1 to 10 and were noticeably bimodal. Scores cluster at the poles—with 13.8 of respondents selecting 1 (4 of 29) and 17.2 selecting 10 (5 of 29)—rather than converging near the mean. This polarization suggests that the sample is not merely divided on degree of enthusiasm in using AI, but that there are two large subgroups of opinion. One group anticipates deep, active reliance on AI to learn and the other intends to use AI either more discreetly or tasks irrelevant to learning legal practical skills and processes. This split exists even among the 86.2% (25 of

29) who indicated intent to adopt AI, revealing that broad willingness to adopt does not translate into uniform depth of adoption.

TABLE A.1.2. Respondents on Likelihood of using AI as a Learning Tool.

On a scale of 1 to 10, how likely are you to use AI to help you learn the practical skills and processes of working at a law firm?

29 responses




2. For What Purposes?

The data reveals an interesting division in the type of tasks respondents intend to use AI on compared to the ones not. As shown in Table A.2.1., when respondents were asked to identify the tasks for which they would most want AI assistance during their first years of practice, respondents were clustered overwhelmingly around peripheral, lower-stakes functions. The most selected task was “providing quick summaries of complex legal fields I am unfamiliar with,” chosen by 86.2% of respondents (25 of 29). “Document review and summarization” ranked second at 72.4% (21 of 29). “Drafting non-substantive professional communications” ranked third at 48.3% (14 of 29), and “learning legal research platforms and advanced query strategies,” ranked fourth at 44.8% (13 of 29).

By contrast, tasks involving assistance with direct client interaction or adversarial professional judgment did not get as much interest. “Simulating client interactions or deposition practice,” which was the only task within the survey question that involved client-facing work,¹¹ placed fifth out of the eight options and tied with “understanding and applying firm-specific processes and workflows.”

TABLE A.2.1. Tasks where AI Adoption is Most Likely.



The distribution suggests that respondents have distinguished the kinds of work they are willing to delegate to AI from the kinds of work they are not. Tasks that appear informational, preparatory, or administrative attracted the highest levels of interest. While, tasks that seem to require direct client contact or real-time professional judgement attracted less attention. This pattern is

¹¹ See *Law Firm Client Management Is a Competitive Edge*, Thomson Reuters Legal Blog, <https://legal.thomsonreuters.com/blog/law-firm-client-management-is-a-competitive-edge/> (last visited Apr. 28, 2026); Moray McLaren, *Defining the Perfect Partner*, The Practice (Nov. 2023), <https://clp.law.harvard.edu/knowledge-hub/magazine/issues/the-perfect-partner/defining-the-perfect-partner/>; *Mastering Client-Facing Roles: Tips for Junior Attorneys to Build Confidence and Trust*, Vault Law Blog, <https://vault.com/blogs/vaults-law-blog-legal-careers-and-industry-news/mastering-client-facing-roles-tips-for-junior-attorneys-to-build-confidence-and-trust/> (last visited Apr. 28, 2026).

significant because the tasks respondents are least willing to delegate to AI are precisely the ones more consequential to professional development and advancement at a law firm.¹²

There are two explanations that could explain this finding. The first is that respondents feel more capable handling high-stakes tasks independently and see less need for AI assistance there. However, this seems implausible because first-year law students have never taken a deposition, never managed a client relationship, and by their own admission¹³ are unfamiliar with many areas of the law. So, it would be strange for respondents to feel so confident in their client-facing abilities that they don't want AI help. The second, and more compelling, is that respondents are calibrating their AI use to the cost of error. They are comfortable delegating tasks where a mistake seems recoverable, and reluctant to delegate tasks where an AI error could have direct professional consequences. The wariness data examined in Part B is consistent with this explanation.

B. The Structure of Professional Wariness.

The survey asked respondents to rate their wariness across five professional concerns on a five-point scale from “Not Wary at All,” “Slightly Wary,” “Moderately Wary,” “Very Wary,” and “Extremely Wary.” The respondents were asked to respond to these professional concerns: (1) client confidentiality and data security, (2) accuracy and reliability of AI-generated legal content, (3) job displacement, (4) ethical compliance and unauthorized practice of law, and (5) institutional acceptance and firm policy.

¹² McLaren, *supra* note 11.

¹³ In the third question of the survey, some respondents marked “Providing quick summaries of complex legal fields I am unfamiliar with” applicable to them. Implicitly, the respondents marking such a box admits that there are legal fields that are “unfamiliar” to them.

Accuracy and reliability ranked first with a mean wariness score of 4.24 out of 5. 79.3% (23 of 29) of all respondents selected either “Very Wary” or “Extremely Wary” —with nearly half of all respondents (14 of 29) selecting the latter. No respondent selected “Not Wary at All.” This suggests accuracy anxiety is a universal cause of wariness within this sample.

Client confidentiality ranked second with a mean of 4.00. 69% (20 of 29) of all respondents selected “Very Wary” or “Extremely Wary.” Similar to the sample's response to the accuracy concern, no respondent selected “Not Wary at All.”


Ethical compliance ranked third with a mean of 3.45. The distribution is more spread, with 55.2% (16 of 29) selecting “Very Wary” or “Extremely Wary.” To this professional concern, 6.9% (2 of 29) respondents selected “Not Wary at All.”

Job displacement ranked fourth with a mean of 3.07. Unlike the top two professional concerns, the displacement concern produced a divided sample: 41.4 (12 of 29) respondents selected “Not Wary at All” or “Slightly Wary,” while 37.9 selected “Very Wary” or “Extremely Wary.” Displacement anxiety exists in this sample, but it is neither dominant concern nor universally reflected among respondents.

Institutional Acceptance ranked last with a mean of 2.38. Respondents were largely unconcerned about whether their employers will permit AI use.

TABLE B.1. Wariness Across Professional Concerns.

Please rate your level of wariness regarding the use of AI in a professional legal context.



This hierarchy establishes a clear ordering across all five dimensions. Notably, it directly challenges the prevailing narrative that junior lawyers are primarily anxious about job displacement.¹⁴ In this sample, displacement ranks fourth. Students are not primarily afraid of being replaced. They are afraid of being wrong.

The open-ended responses reinforce this finding. When asked to identify the single most important capability an AI assistant must have, respondents converged on verification, citation accuracy, and hallucination prevention. 41.4% (12 of 29) explicitly named verification or citation checking—with responses included “a way to quickly and easily verify any information it gives the user,” “citation/hallucination checker,” and “easy to verify and very clear about each logical step it makes.”

¹⁴ See *Top Six AI Legal Issues and Concerns for Legal Practitioners*, Law Tech. Today (ABA Law Prac. Div., 2025), https://www.americanbar.org/groups/law_practice/resources/law-technology-today/2025/ai-legal-issues-and-concerns-for-legal-practitioners/; Megan Doyle, *Is AI Finally Going to Take Our Jobs? Meeting Client AI/Technological Demands While Supporting Junior Lawyers' Development*, Int'l Bar Ass'n (Nov. 29, 2024), <https://www.ibanet.org/is-AI-finally-going-to-take-our-jobs>.

Comparing how individual respondents rated accuracy wariness against numerical likelihood scores does not produce a strong directional finding. Respondents with high accuracy wariness (selecting “Very Wary” or “Extremely Wary”) produced a mean likelihood score of 5.57, compared to 6.00 among those with lower accuracy wariness. Due to the small sample size, this difference is not dispositive. How accuracy anxiety shapes the conditions of adoption—rather than the likelihood of it—is examined in Part C.


C. The Demand for Verification as a Condition for Fuller Adoption.

The survey asked respondents to select up to three assurance features that would make them more comfortable using AI tools in early legal practice. The results, ranked by frequency selection, are as follows. Guaranteed data encryption ranked first, selected by 65.5% of respondents (19 of 29). Built-in verification mechanisms for AI-generated content ranked second selected by 62.1% (18 of 29). Endorsement and official training from the law firm ranked third at 55.2% (16 of 29). AI specifically designed for the legal industry ranked fourth at 48.3 (14 of 29). A clear ethical framework from the ABA or State Bar ranked fifth at 41.4% (12 of 29). Features focused on administrative tasks ranked sixth at 31% (9 of 29). A designated human supervisor ranked seventh at 24.1% (7 of 29). Free access to high-quality tools ranked last at 20.7% (6 of 29).

TABLE C.1. Features that Would Make AI Adoption More Comfortable.

What features or assurances would make you feel more comfortable using AI tools in your early legal practice? (Select up to three most important factors)

29 responses



Two features of this distribution deserve attention. First, the top two selections—encryption and verification—together suggest that respondents are not just asking for better AI performance. Rather, they are asking for protective infrastructure around AI outputs: one layer that keeps their data secure, and another that checks whether what the AI tells them is actually true. Second, respondents are not asking for free access, human supervisors to their AI assisted work products nor administrative-only restrictions to their AI tools. They are asking for tools that are secure, verifiable, institutionally endorsed, and purpose-built for law. This suggests that the respondents do not want to compromise on their AI use, they want their AI tools to rise to their expectations.

IV. DISCUSSION

The findings across Parts A, B, and C tells a story with two dimensions for the students in this sample.

A. The Barrier to Full Adoption is Epistemological, Not Attitudinal.

In Part A(2), students self-limited their intended use to the profession’s periphery—document review, field summaries, and routine communications—and showed notably less interest in client-facing or judgement-dependent tasks. In Part B, accuracy and reliability produced the sample’s highest wariness scores. In fact, accuracy and reliability were the only two concerns on which no respondent selected “Not Wary at All,” suggesting near-universal anxiety among the sample. In Part C, built-in verification mechanisms ranked as the second most selected assurance feature at 62.1%, and 51.7% of respondents named some form of verification, citation accuracy, or hallucination prevention as the single most important capability an AI assistant must have. Across all three finding sections, the same pattern emerges: students want to use AI, they intend to use AI, but they will not use it for anything consequential until they have a basis for trusting its outputs.

This pattern is not due to outright resistance to AI—it is professional risk calibration. Errors attributable to uncritical reliance on AI output carry distinct professional risks, including sanctions that courts have begun to impose specifically in that context.¹⁵ Therefore, the barrier to fuller adoption is not caused by technophobia. Rather, it is caused by the absence of a verification infrastructure that would make reliance on AI outputs professional defensible.

B. The Market Has Not Yet Supplied What Students Are Asking For.

The assurance preference data reveals something important about the current state of the legal AI market. Respondents are not asking for AI to be faster or more capable. They are asking for a different kind of AI—one with a verification layer built into its architecture rather than bolted on as a

¹⁵ See Nicole Black, *The Real Error Is Human: AI Can’t Cure Carelessness*, Above the Law (May 29, 2025), <https://abovethelaw.com/2025/05/the-real-error-is-human-ai-cant-cure-carelessness/>.

feature. The responses naming “citation/hallucination checker,” “fact checking/citation checking,” “easy to verify and very clear about each logical step it makes,” and “clear citations for all output so everything can be double checked” are not describing marginal improvements to existing tools. They are describing a product that does not yet reliably exist in the legal market.

This has direct implications for how legal AI products should be built and how bar associations should respond. For product design, the data suggest that competing on capability before solving verification is addressing the wrong problem. The students who will be the profession’s primary AI users within the next three to five years have already decided that verification is the threshold condition for adoption.

CONCLUSION

This study asked whether first-year law students at Harvard Law School intend to adopt AI in their early careers and on what terms. The answer is yes, but conditionally. Students are broadly willing to adopt AI, but have a precondition for deeper reliance on it: they need a reliable means of verifying what the AI tells them before they will trust it with anything consequential. This insight directly implies that the legal AI market has not reliably supplied a verification infrastructure. So, until it does, the adoption of AI by the profession’s incoming generation will remain broad in form but shallow in substance.