
The Shadow AI Advantage

How law firms can turn individual AI use into lasting productivity



INDIVIDUAL INITIATIVE
SHARED CAPABILITY

Tyler Whitley

Original paper · October 8, 2025

Alternate revision prepared with Codex

September 2026 edition · Evidence checked September 9, 2026

THE OPPORTUNITY

Individual initiative. Lasting productivity.

A law firm can spend months choosing an AI platform. A lawyer with a draft due tomorrow has a shorter decision to make. Can a tool help organize the facts, improve an awkward paragraph, or get a useful first version onto the page?

That difference in urgency helps explain why AI adoption can advance at a lawyer's desk while the firm's strategy is still taking shape. In its 2026 research on the legal profession, Thomson Reuters found that 34 percent of surveyed law firm professionals were using AI tools their firms had not authorized. The survey covered 736 law firm professionals across 46 countries. It captures a substantial gap between what firms have approved and what people are doing.^[1]

There is a business opportunity in that gap, alongside a real responsibility to close it. Unapproved use can expose client information and leave firms unable to supervise the work. It can also reveal where people are looking for help. Leaders should want to know which tasks prompted that search, what the tools actually improved, and why an approved route was missing.

The argument for starting with individual AI is straightforward. Give lawyers access to suitable tools for specific work, teach them how to use and check those tools, and measure what changes. Expand the applications that survive that test. A firm should not need a technology transformation to find out whether its lawyers can prepare a better first draft in less time.

This approach gives individual initiative a legitimate place in the firm's development. The useful discovery can begin with one person. The firm supplies the conditions under which others can rely on it.

THE CENTRAL IDEA

The useful discovery can begin with one person. The firm supplies the conditions under which others can rely on it.

01 / START WITH THE TASK

Start where the work is

The attraction of a general AI assistant is easy to understand. A lawyer can ask it to reorganize an outline, compare two passages, or explain why a draft is difficult to follow. The question can be narrow enough to judge, and the interaction can happen while the lawyer is already working. A promising use does not have to wait for every practice group to agree on a new process.

Evidence outside law helps explain why that matters. In a six-month randomized field experiment involving 7,137 knowledge workers at 66 firms, researchers studied access to AI embedded in familiar work applications. Their revised paper reports that, in the second half of the experiment, workers who used the tool spent about two fewer hours a week on email. The researchers did not detect changes in the quantity or composition of tasks beyond those individual time savings. This was an enterprise-provided tool, and the study was not specific to lawyers.^[2]

Its relevance is the level at which change occurred. People could alter their own handling of email more readily than the larger organization could alter its work. That is a useful starting point for a law firm. An associate may be able to improve how they prepare a chronology before the firm is ready to redesign litigation support.

Individual use and enterprise provision can therefore coexist. A firm-managed account may offer the same direct interaction that made a personal tool appealing. Specialized software may add access to legal sources or controls that the task requires. The important distinction is whether the tool fits a real piece of work and whether the person using it knows how to judge the result.

Start by asking lawyers to show a recurring task they would like to improve. Watch the current process. Find out where the time goes and which mistakes are costly. Those observations give a firm something much more useful to test than a vendor's general promise of productivity.

START HERE

Ask lawyers to show a recurring task. Watch the current process. Find out where time goes and which mistakes are costly.

02 / THE EVIDENCE

Where legal AI stands in September 2026

69% of 8am survey respondents used general-purpose AI for work. [5]

41% of Thomson Reuters law firm respondents reported organizational generative AI use. [11]

Individual adoption is already substantial. In 8am's 2026 Legal Industry Report, 69 percent of respondents said they used general-purpose AI for work, up from 31 percent in the previous report. Its US survey included more than 1,300 legal professionals, with responses collected in September and October 2025.^[5]

Firm adoption is advancing too. Thomson Reuters's separate 2026 AI in Professional Services report found that 41 percent of law firm respondents said their organizations used generative AI. The figure for corporate legal departments was 47 percent. These are reports of organizational use, rather than measures of how often every lawyer uses a tool.^[11]

Those surveys ask different questions of different groups. Their percentages cannot be combined into one adoption rate. Taken together, they support a change in emphasis. A firm considering AI today is entering a profession where individual use and institutional adoption are both well underway. The decision increasingly concerns which work to support, how to supervise it, and what the firm expects to gain.

There is now reported experience to put alongside the earlier forecasts. In the 8am survey, 61 percent of legal professionals said AI saved them time each week; 38 percent reported saving one to five hours. These are respondents' estimates, not independently timed results or a promise of savings for the next purchaser.^[5]

READ THE NUMBERS CAREFULLY

Different surveys, populations, and questions. These figures cannot be combined into one adoption rate.

02 / THE EVIDENCE, CONTINUED

What experiments can tell us

Controlled experiments help establish what can cause an improvement, although their tools and settings must be dated carefully. A study published in 2024 by Jonathan Choi, Amy Monahan, and Daniel Schwarcz assigned law students realistic legal tasks with or without GPT-4. AI assistance produced large, consistent gains in speed, while improvements in the quality of legal analysis were smaller and uneven.^[3]

A later experiment by Schwarcz and colleagues, published in 2026, tested o1-preview and Vincent AI 2024. Both improved student work quality and retained efficiency benefits, with significant productivity gains on five of six tasks. The publication date does not make those systems current. The findings support testing both general and specialized assistance; they do not rank the products available in September 2026.^[4]

These experiments studied students completing bounded assignments. An experienced lawyer managing a live matter works with different incentives, obligations, and sources of difficulty. The findings justify a measured trial in practice. They cannot supply a guaranteed savings figure for every lawyer who receives a subscription.

The case for adoption no longer needs to rest on a forecast of hundreds of recovered hours or a sweeping claim that enterprise projects fail. Lawyers are reporting useful applications, firms are adopting tools, and experiments establish benefits on specific tasks. The next step is to determine which of those benefits a particular firm can reproduce.

FROM EVIDENCE TO PRACTICE

The findings justify a measured trial in practice. They cannot supply a guaranteed savings figure for every lawyer who receives a subscription.

03 / TRAINING

Training changes what access is worth

The gap between use and support is visible in current survey evidence. In the State Bar of Texas's 2026 survey, 50 percent of respondents reported that their firm had no official AI policy. Among AI users, 46 percent said their firm provided no AI training.^[6]

Giving someone a capable tool leaves a surprising amount undecided. What should they ask it to do? Which information belongs in the prompt? What would an adequate answer contain? How will they recognize an omission that sounds perfectly plausible?

A 2026 working paper by Benjamin Chen and Hong Bao makes the importance of instruction concrete. In an experiment completed by 164 University of Hong Kong law students, participants took a contract-law examination with no AI access, optional access to DeepSeek, or access accompanied by brief training. Reported AI use was about 41 percent in the trained group and 26 percent in the untrained group. Trained participants also earned higher scores than untrained participants and stated applicable rules more accurately. This was one examination at one institution, not a field trial of practicing lawyers.^[7]

That is a reason to budget for instruction alongside access. A subscription leaves each lawyer to discover an effective method. Training gives them a place to begin and a way to recognize when the method fails.

For a firm, that instruction should be close to the work. Take a permitted document and a task the team already understands. Show the starting material, the request, and the output together. Discuss where the model compressed a distinction, skipped an exception, or introduced a claim the source did not support. Then demonstrate how the lawyer checks and repairs it.

Teaching an associate to request a chronology is only part of the lesson. They also need to distinguish the date of an event from the date on which someone described it, retain conflicting accounts, and trace an entry to its source. The prompt should reflect those requirements. The review should test whether they were met.

As successful uses accumulate, the firm can preserve the instructions, examples, and review steps that made them work. Another lawyer can then begin with a tested approach and adapt it to the matter. Over time, that shared knowledge makes adoption less dependent on the enthusiasm of a few individuals.

04 / VERIFICATION

Build verification into the workflow

A fluent answer can conceal a difficult review problem. A summary may accurately describe most of a document while omitting the qualification that changes its meaning. A citation may point to a real case that does not support the proposition beside it. Reading for style will not catch either problem reliably.

Current evidence shows that review remains a practical concern. In the Texas bar's 2026 survey, 40 percent of respondents described themselves as not at all or only slightly confident in evaluating AI-generated documents. That measures confidence, not their actual ability to detect errors, but it gives firms a concrete reason to assess the review skills their workflows depend on.

[6]

Earlier technical research explains why the review needs substance. A preregistered 2024 evaluation found that specialized legal research tools reduced hallucinations relative to the general-purpose model tested, but did not eliminate them. Those findings concern the versions evaluated. A present-day purchase needs testing on the current product and the firm's intended tasks.^[8]

The lasting operational question is how easily the lawyer can check the answer. A useful workflow keeps source material close, identifies the passages behind important claims, and makes it possible to inspect omissions. An answer with impressive formatting but no practical route back to the evidence may take longer to verify than it saved in drafting.

THE REVIEW STANDARD

Keep source material close. Identify the passages behind important claims.
Make omissions possible to inspect.

04 / VERIFICATION, CONTINUED

Measure finished work

ILLUSTRATIVE CHRONOLOGY PILOT · NOT A CLIENT RESULT

Existing process

90 minutes

Assisted process, including review

55 minutes

35 minutes

illustrative saving

Consider a proposed pilot for preparing a chronology from a closed, suitably cleared matter. This is an illustration, not a reported client result. The team establishes how long the existing method takes, then uses an approved tool to produce a draft with source references. A lawyer checks the entries, looks for missing events, and resolves ambiguities. The team records all of that time, including corrections and final review.

Suppose the existing process takes 90 minutes and the assisted process takes 55 minutes, including review. The illustrative saving is 35 minutes. If the model generates its draft in two minutes but leaves 95 minutes of checking and repair, the process has become slower. The fast generation of text is interesting. The time required to finish acceptable work determines the practical result.

Quality also needs its own record. Count material omissions and unsupported statements, and examine their consequences. Compare several representative tasks, including awkward cases. A method that works only on the cleanest example is a poor basis for a wider rollout.

COUNT THE WHOLE PROCESS

If generation takes 2 minutes but checking and repair take 95 minutes, the assisted process takes 97 minutes. Faster generation alone does not establish a benefit.

05 / AUTHORIZATION

Give useful work an approved home

The phrase shadow AI describes use outside an organization's authorization or oversight. It does not tell us whether a tool is general-purpose or specialized, free or paid. A commercially licensed product can still be unauthorized. A familiar chat interface can sit inside an approved arrangement.

The firm's job is to make that arrangement clear enough to follow. People need to know which tools they may use, which information those tools may receive, and who can resolve an uncertain case. An account that requires every lawyer to interpret vendor terms and client obligations independently leaves too much unsettled.

ABA Formal Opinion 512 applies existing professional duties to generative AI, including competence, confidentiality, communication, and supervision. It calls for understanding a tool's capabilities and limitations and reviewing its output appropriately. It does not make an AI label a substitute for professional judgment.^[9]

Florida Bar Opinion 24-1 likewise permits AI use subject to professional obligations. It recommends obtaining informed consent before using a third-party generative AI program where confidential information would be disclosed. Its analysis covers the broad duty of confidentiality and the need to understand how the provider handles information. Removing a client's name, by itself, should not be assumed to resolve those questions.^[10]

The applicable jurisdiction, court requirements, and client instructions still govern the matter. A workable firm policy should translate those obligations into concrete decisions before people begin uploading documents. Who approved the use? What material is permitted? Where is it retained? Who reviews the result?

These decisions are part of implementation. Once they are made, lawyers can work within a known arrangement and bring new uses forward for assessment. A firm that understands why people reached for outside tools is better placed to provide an approved option they will actually use.

06 / BUSINESS VALUE

Turn saved time into a business result

The original attraction of AI is often expressed in hours. A law firm's return depends on what happens to those hours afterward.

For an hourly matter, completing the work faster can reduce the fee for that task. ABA Formal Opinion 512 is clear that a lawyer billing by the hour must charge for time actually spent. Time saved cannot be billed as though it had been worked. Alternative fee arrangements also remain subject to reasonableness requirements.^[9]

A firm may recover the economic benefit through additional matters, reduced backlog, lower delivery costs, or a service clients value more. Each route has conditions. Additional capacity helps revenue only if there is demand and the firm can turn that capacity into paid work. A fixed fee creates room for better margins only within a sound agreement and an acceptable result. Sometimes the immediate benefit will be a less punishing working day. That has value too, but it should be described honestly.

Measurement is still catching up with adoption. In the 2026 AI in Professional Services survey, only 18 percent of respondents said their organizations collected AI return-on-investment metrics. The survey covered 1,514 AI-familiar professionals across legal and other professional fields. It should not be read as a census of law firms.^[11]

For a modest pilot, measurement can remain modest. Record the time needed to complete the work, the reviewer's assessment, and the cost of the tool and training. Establish whether turnaround improved for the client and whether the change persisted after the novelty wore off. Measure the benefit the firm intended to create.

That makes the next purchasing decision easier. Repeated demand for the same document workflow may justify an integration. Difficulty finding authoritative sources may justify a specialized research product. A successful use shared across a practice group may justify more accounts and structured support. The investment follows a demonstrated need, with a clearer explanation of what the additional expense should accomplish.

07 / THE NEXT STEP

Make the next decision smaller

A firm does not need to resolve the future of legal practice before improving one part of it. It needs a suitable task, an approved tool, and people who can tell whether the result is worth keeping.

The first step is to look closely at work already happening. Ask where AI is being used, including outside approved systems. Learn what people were trying to accomplish. Test the promising uses under conditions the firm can defend, and give the successful ones a route into ordinary practice.

This is the opportunity behind shadow AI. Individual experimentation can reveal needs that a broad technology plan overlooks. Training and supervision can turn a promising use into a repeatable method. Thoughtful investment can then extend that method to the people and matters that benefit from it.

Start with the next piece of work that can be done better. Give the lawyer doing it the tools, instruction, and authority to find out. Let the result inform the next decision.

A PRACTICAL SEQUENCE

Find a task.

Test a method.

Measure the result.

Extend what works.

READING NOTES / 1 OF 2

Sources and reading notes

Evidence checked September 9, 2026. Publication year and survey field dates differ. The 8am survey was fielded September-October 2025; the Thomson Reuters AI in Professional Services survey was fielded October-November 2025; the Texas bar survey closed February 10, 2026; and the Thomson Reuters legal Future of Professionals survey was fielded March-April 2026. These are the cited reports available at review, not live measurements of the profession on September 9. This is an independent synthesis and proposed approach, not an original empirical study. Survey findings, controlled experiments, and illustrative examples are identified separately. Vendor surveys provide useful context but are not neutral population censuses. Historical model evaluations do not establish the performance of current products. Ethical opinions are cited for their stated guidance, not as a complete account of every jurisdiction's requirements.

[1] Thomson Reuters Institute. Future of Professionals Report 2026: Actionable insights for law firm leaders. See the shadow AI discussion and About this report. Data collected March-April 2026. Read the report.

[Read source](#)

[2] Eleanor Wiske Dillon, Sonia Jaffe, Nicole Immorlica, and Christopher T. Stanton. Shifting Work Patterns with Generative AI. Working paper, arXiv version 4. The reported email result concerns users among those offered access, not all workers assigned access. Microsoft researchers coauthored the study. Read the paper.

[Read source](#)

[3] Jonathan H. Choi, Amy B. Monahan, and Daniel Schwarcz. Lawyering in the Age of Artificial Intelligence. Minnesota Law Review 109, 147 (2024). Randomized student experiment with blind assessment of work. Read the article.

[Read source](#)

[4] Daniel Schwarcz, Sam Manning, J. J. Prescott, Patrick Barry, David R. Cleveland, and Beverly Rich. AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice. 2025 working paper; published in 2026. Findings cited from the authors' institutional abstract. Institutional record and published article.

[Read source](#) · [Related source](#)

[5] 8am. 2026 Legal Industry Report. Published March 5, 2026. Vendor survey of more than 1,300 US legal professionals. Adoption and time savings are self-reported. Report and methodology and publisher's account of time savings.

[Read source](#) · [Related source](#)

READING NOTES / 2 OF 2

Sources, continued

[6] State Bar of Texas Department of Research and Analysis. 2026 AI in the Practice of Law Survey. See printed pages 10-11 on policy, training, and confidence. The survey reports 1,553 participants and a 5.2 percent response rate; voluntary response can affect representativeness. These are Texas survey findings. Read the report.

[Read source](#)

[7] Benjamin M. Chen and Hong Bao. Training for Technology: Adoption and Productive Use of Generative AI in Legal Analysis. Working paper, revision 3, June 5, 2026. See experimental design, results, and limitations. The training comparison should not be generalized into a guaranteed gain for practicing lawyers. Read the paper.

[Read source](#)

[8] Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, and Daniel E. Ho. Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools. 2024 evaluation. See query design, results, and limitations; the benchmark concerns the versions and tasks evaluated. Read the paper.

[Read source](#)

[9] American Bar Association Standing Committee on Ethics and Professional Responsibility. Formal Opinion 512: Generative Artificial Intelligence Tools. July 29, 2024. See competence, confidentiality, supervision, and fees, especially pages 12-14 on charging for work and tools. Read the opinion.

[Read source](#)

[10] The Florida Bar. Ethics Opinion 24-1. January 19, 2024. See confidentiality and oversight, pages 2-4. Recommendations about consent depend on the information and the proposed use. Read the opinion.

[Read source](#)

[11] Thomson Reuters Institute. 2026 AI in Professional Services Report. Pages 3 and 16-18 on measurement, page 5 on organizational adoption, and page 23 on methodology. Data collected October-November 2025 from respondents screened for familiarity with AI. Read the report.

[Read source](#)