AI, ACCESS TO JUSTICE, AND PROFESSIONAL JUDGMENT

The Justice Gap Has a New Unit of Measure

AI Lowered the Cost of Expression. It Did Not Lower the Investment Required for Sound Legal Judgment.

By Brian Edward Walters | www.TexasTowingLawyer.com

ABSTRACT

Artificial intelligence has made legal-looking language cheap enough that it feels almost free. A person who could not previously draft a demand letter, contract clause, intake summary, policy, research memo, or court filing can now generate something that looks legal in seconds. That is real access. It is not the same thing as making law cheap. Legal expression is not legal work. A document is not valuable because it resembles a lawyer's work product. It is valuable when it fits the facts, the law, the forum, the client objective, and the next step in the matter. AI lets users stand on accumulated legal language and move straight to production without necessarily acquiring the habits that make legal production safe: source criticism, factual restraint, verification, risk assessment, counseling, and judgment. It also creates an intoxicating feeling of ungated access. People experience the same rush when AI lets them write software, build a website, or publish something that used to require an expensive intermediary. Law now produces that sensation too. The old justice gap was measured in lawyer hours, specialized attention, research tools, institutional knowledge, and the ability to stay with a legal problem until it is actually solved. The new justice gap may add a new unit of measure: compute. The better-funded user may buy not only more lawyers, but more agents, more tokens, more retrieval, more context, more verification loops, and more disciplined passes at the same problem. Meanwhile, subsidized subscriptions and free tiers train lawyers and legal users to experience AI as abundance before the true professional meter appears. This article argues that AI may narrow the drafting gap while widening the justice gap if the market confuses cheap legal expression with legal work product that can be trusted in context.

AI lowered the cost of expression. It did not lower the investment required for sound legal judgment.

WHAT THIS ARTICLE COVERS

I. The Drafting Gap Is Not the Justice Gap

AI lowered the cost of legal expression. It did not lower the investment required for sound legal judgment.

That distinction is the center of the problem. The public story about artificial intelligence and law is often framed as access: people who could not afford legal help can now produce legal-looking documents. Lawyers can draft faster. Small firms can punch above their weight. Clients can understand their problems more clearly. People can enter negotiations, business relationships, administrative processes, or courts with documents that look closer to the form the legal system expects.

That is true as far as it goes. The trouble is that it does not go far enough.

The justice gap was never only a drafting gap. The Legal Services Corporation reported in its 2022 Justice Gap Study that 92% of civil legal problems reported by low-income Americans received no legal help or not enough legal help.[1] The World Justice Project estimates that billions of people globally have unmet justice needs, including people who cannot obtain justice for everyday legal problems and people excluded from the opportunities the law provides.[2] Those gaps are not caused only by a shortage of legal words. The system does not leave people stranded merely because they lack nouns and verbs. It leaves people stranded because advice, representation, knowledge, procedure, evidence, time, leverage, enforcement, and endurance cost money.

AI can reduce one barrier by helping people create legal language. But a legal-looking document is not yet a legal tool. It has to be fitted to a purpose, a client objective, a factual record, a governing body of law, and the institution or transaction in which it must operate. A clause is not risk allocation until it fits the deal. A memo is not advice until it answers the actual question. A policy is not compliance until it survives the rules that govern the business. A demand letter is not leverage merely because the paragraphs sound like they came from a lawyer. A court filing is not a case.

In a strange way, AI may finally prove something lawyers have been trying to explain for years. Lawyers were never valuable because they had secret access to forms. Forms mattered, and access to them mattered, but the form was never the whole service. The work was knowing which form belonged to the problem, what facts had to be true before it could be used, what language had to be changed, what risk it created, what law governed it, and what would happen after someone signed it, sent it, filed it, or relied on it. AI has made the form visible. It has also made the limits of the form visible.

The risk is that we will mistake the narrowing of the drafting gap for the narrowing of the justice gap. Those are different gaps. The next gap is not only economic. It is epistemic. It turns on who knows enough to ask the right questions, doubt the right answers, and recognize when the polished legal language has not actually solved the legal problem.

II. The Discipline Problem

There is a Jurassic Park problem in legal AI.

The danger is not that people stand on prior knowledge. Law has always worked that way. Lawyers cite older cases, reuse forms, borrow clauses, inherit doctrines, learn from treatises, and build arguments from the work of people who came before them. Legal knowledge is cumulative. Nobody starts from first principles.

The danger is that AI allows the user to take the next step without necessarily acquiring the discipline that makes the step safe.

AI also creates a kind of *Matrix* illusion. The user feels as though the knowledge has been downloaded directly into the brain. The document appears, the cadence is right, the terminology is familiar, and the user can feel as if the skill arrived with the output. But there is a difference between producing the appearance of legal competence and possessing the judgment that legal competence requires. The user may be able to imitate the moves. That does not mean the user knows where the next move should be, why it matters, or what happens if it is wrong. [3]

A lawyer does not learn only words. A lawyer learns when words are dangerous. A lawyer learns that a clause can shift risk without announcing it, that a fact can be assumed before it is proven, that a rule can change across jurisdictions, that a remedy can exist in theory but fail in practice, and that a confident paragraph can be wrong because it answers the wrong question.

That investment is slow. It is expensive. It is often invisible to the client. It is checking the source. It is reading the case after the headnote. It is asking whether the factual premise is true. It is knowing when the neat answer is too neat. It is understanding that the polished draft may be the easiest part of the job.

AI can simulate the surface of that discipline. It can produce the cadence of legal analysis, the structure of a memo, the vocabulary of risk, and the appearance of confidence. What it does not automatically provide is the user's respect for the thing being produced. That respect usually comes from doing the work long enough to know where it breaks.

That is why token cost is not mere procurement trivia. If the market teaches users that they can generate endlessly without feeling the cost, it may also teach them to treat iteration as a substitute for investment. More prompts can help, but only if the person directing them knows what to ask, what to doubt, what to verify, and when to stop.

III. The Intoxication of Ungated Work

Part of what makes this moment hard to see clearly is that AI feels liberating.

A person who could not code can now build a small piece of software. A business owner who could not build a website can publish one. A professional who used to pay someone else to translate an idea into a technical artifact can now sit down with a model and produce the artifact directly. That feeling is not fake. It is powerful because the gate was real.

Not every gatekeeper was adding value. Not every expensive quote reflected deep expertise. Many people have been excluded from useful tools by price, jargon, dependency, and the basic fact that they did not know the technical language of the field. AI cuts through some of that. It gives people an immediate sense of agency.

That sensation is intoxicating.

The user can mistake the collapse of the first barrier for the collapse of the whole discipline. A page loads, and the user feels like a software developer. A website deploys, and the user feels like a technology company. A legal memo appears, and the user feels like the law has become accessible in a way it never was before. In one sense, that feeling is right. The user has crossed a threshold. The old dependency has been weakened. The first artifact exists. But entry is not mastery.

Law and software share this structure. Both have visible artifacts and invisible systems. The website is visible. Security, dependencies, accessibility, hosting, uptime, maintainability, data flows, and failure modes are less visible. The legal document is visible. Issue spotting, adverse authority, procedural posture, evidentiary burden, remedy selection, factual assumptions, enforcement risk, and omitted questions are less visible.

The danger is that AI gives the user the feeling of access before it gives the user the discipline of evaluation. The person can ask for the thing, receive the thing, revise the thing, and publish the thing. That is a profound change in access. But it does not mean the user understands the corpus behind the thing.

In law, the unknown unknowns are often the whole game. The issue the user does not know to ask about may be the issue that controls the matter. The exception the user does not know exists may be the exception that defeats the claim. The missing fact may be the fact that changes the remedy. The procedural posture may be the reason the otherwise correct rule does not help.

No serious lawyer should pretend humans are perfect at this. Lawyers miss issues too. But legal training is supposed to build a habit of suspicion: what am I not seeing, what changes the analysis, who has the burden, what law controls, what happens next, and what would make this document fail?

A private consumer AI account may assist that process. It should not be mistaken for the full process. It can generate a plausible first map. It does not reliably know when the map omits the road that matters.

That is one of the sources of the coming shock. The technology feels like it has removed the gate. In some ways, it has. But the discipline behind the gate remains.

IV. The Old Unit Was Lawyer Time

The legal system has never been an economically level playing field, and it has never taken long for that fact to matter. Equal formal rights do not produce equal practical leverage when one side can fund more attention, more specialization, more patience, and more rounds of analysis than the other. A person or institution with more money can hire better lawyers, bring in specialists, run more research, review more documents, negotiate longer, test more options, and get a second opinion before committing to a course of action. That is not cynicism. It is the economics of litigation. Money buys attention. Money buys options. Money buys the ability to keep asking questions after the cheaper side needs to move on.

There is a reason the line from *A Civil Action* still lands: litigation can become a contest in which each side spends more than it should to force the other side to do the same, and "whoever comes to their senses first loses." [4] The line is memorable because it says openly what the economics of litigation usually dresses up. Procedure is not only a search for truth. It is also a test of endurance.

Before AI, much of that advantage was measured in lawyer time. A well-funded client could pay a team to review every document, structure the transaction, compare authorities, identify risk, prepare witnesses, revise the agreement, test settlement positions, map regulatory exposure, and update the advice as the facts changed. A poorer client might have a valid problem and still lack the resources to develop it. The difference was not merely access to legal language. The difference was the amount of professional attention each side could afford to pour into the problem.

That attention mattered because law is a comparison discipline. A lawyer compares facts to rules, rules to remedies, remedies to forums, forums to deadlines, documents to testimony, testimony to impeachment, and strategy to the client's actual objective. Good legal work is not merely producing a document. It is deciding which document matters and what must be true before anyone should rely on it.

AI changes the visible meter. It does not remove the work.

V. The New Unit May Be Compute

AI makes the first answer feel cheap because the first answer is often only a first impression dressed as a deliverable. It may be useful. It may be a good start. But it is not finished legal work. The professional cost appears when the user has to test that answer against authority, facts, procedure, client objectives, institutional realities, and consequences.

That is where the meter starts to matter. The real work begins when the lawyer asks the system to do what legal judgment requires: retrieve sources, compare authorities, check quotations, distinguish jurisdictions, test remedies, read the file, identify contrary facts, evaluate risk, revise the recommendation, and explain what remains uncertain.

That is not one prompt. That is a workflow.

Workflows consume tokens. They consume context. They use retrieval, output, cached input, long documents, model selection, tool calls, and repeated passes. The user may not see all of that inside a consumer subscription, but the provider does. The provider sees the difference between casual chat and a legal work machine.

The public pricing materials already show the structure. OpenAI publishes API prices by input tokens, cached input tokens, and output tokens, and states that API use is billed separately from ChatGPT subscriptions.[5] Anthropic publishes API rates for Claude models and explains that Claude Code cost depends on token consumption, model selection, codebase size, and usage patterns.[6] Google describes Gemini API pricing across free, paid, and enterprise tiers.[7] GitHub announced that Copilot is moving to usage-based billing because agentic workflows bring higher compute and inference demands than ordinary chat.[8]

The point is not that lawyers need to memorize model prices. The point is that the market is moving from subscription psychology to usage visibility. The companies are not measuring professional AI work by vibes. They are measuring tokens, credits, context, tool use, and compute.

That should change how lawyers think about the justice gap. The expensive part of AI-assisted law may not be the first draft. It may be the investment required to make the first draft worth relying on.

The legal system does not need AI that merely helps people produce isolated legal documents faster. It needs AI systems and workflows that help people produce coherent legal work that advances the matter toward an actual endpoint. In litigation, that means documents that fit the record, the procedural posture, the burden, the remedy, the forum, and the next move. In transactional work, it means documents that fit the deal architecture, risk allocation, regulatory setting, closing mechanics, and business objective. A document can look competent standing alone and still fail because it does not belong to the larger legal structure in which it has to function.

VI. The Subscription Teaches Abundance

The subscription experience teaches users to crank.

Crank out the draft. Crank out the rewrite. Crank out the cleaner version. Crank out the table. Crank out the client summary. Crank out the policy. Crank out the clause. Crank out the checklist. If the answer is not right, ask again. If the tone is wrong, ask again. If the format is wrong, ask again. The cost does not feel like money. It feels like patience.

AI is a megaphone. In complex work, it amplifies what the user knows and what the user does not know. A lawyer who understands the field can use AI to expand range, speed, memory, comparison, and production. The tool amplifies judgment because there is judgment to amplify. But a user who lacks the necessary context does not become competent simply because the prose improves. The tool may amplify the absence. It may produce work that looks confident to the user and immediately reveals to a specialist that the user does not understand the problem, the institution, or the body of law in which the work has to live.

That is an astonishing way to learn. It builds fluency quickly. A lawyer who experiments heavily with AI will learn faster than a lawyer who treats it like a delicate toy.

But it also teaches a false price signal. The user learns abundance before learning the cost of professional use. The system feels fixed-price until the rate limit, credit pool, usage bar, overage rule, or enterprise invoice appears.

That pricing psychology matters because discipline has always been connected to friction. Legal work forces the lawyer to slow down. Read the source. Check the record. Ask the client another question. Compare the exception. Revise the theory. A tool that removes friction can be wonderful when it removes waste. It becomes dangerous when it removes the very resistance that used to make the user think.

GitHub's Copilot announcement says the quiet part in product language. Copilot is becoming an agentic platform, and under the prior model a quick chat question and a multi-hour autonomous coding session could cost the user the same amount while GitHub absorbed much of the escalating inference cost.[9] That is exactly the difference that matters for law. A casual question and a full legal workflow are not the same product.

AI companies can subsidize the habit. They cannot subsidize legal consequences.

VII. The Lexis and Westlaw Lesson

Lawyers have seen this kind of habit formation before.

Westlaw and Lexis became part of legal practice because they were useful, but usefulness was only part of the story. They were also built into the formation of lawyers. Law students learned to research inside those systems while access felt broad, fast, and low-friction. Compared with Shepardizing through books, the databases felt like a door opening. The student learned the habit before the student felt the market price. Later, in practice, the tool did not feel like a vendor choice. It felt like how legal research was done.

The analogy is not that legal research companies did something uniquely sinister. The analogy is that subsidized access at the training stage shapes professional dependency.

LexisNexis still describes broad post-graduation access, including six months of Lexis+ access for graduates and longer access for some public-interest graduates.[10] It also rolled out free access to Lexis+ AI for students at ABA-accredited law schools, with promotional material emphasizing that students need to learn the tools firms are already using.[11] Thomson Reuters likewise operates a law school portal built around Westlaw, Practical Law, CoCounsel, training, certifications, and summer-associate preparation.[12]

Those materials are not hidden. They are the business model wearing a name tag. Give future lawyers the tool while they are learning the profession. Make it feel like competent practice. Let the habit form before the commercial price arrives.

AI repeats that pattern with a stronger engine. Westlaw and Lexis trained lawyers where to look. AI trains lawyers how to begin. It does not merely return cases. It produces the draft, the theory, the clause, the letter, the summary, the checklist, the strategy, and the revised version when the first one feels wrong.

There is an important difference. Westlaw and Lexis still required the user to do much of the professional synthesis. To find the answer, the lawyer still had to go through the cases. The lawyer's brain still had to absorb the information, compare the authorities, identify the competing rules, and make a judgment call. Those steps were not waste. They were safeguards.

AI weakens those safeguards in a deceptive way. It can make confidence appear to be research. It can encourage the user to accept a solution because the solution is neatly written, even when the user has not verified whether that solution actually exists in the law, fits the facts, or survives the next procedural step. AI moves closer to the final legal artifact. It does not just open the library. It hands the user something that looks like the work product that would have come out of the library.

The research databases made lawyers dependent on a place to find the law. AI may make lawyers dependent on a machine to start the work. That is the trap.

VIII. The Justice Gap May Move

The access-to-justice story is only half true if it counts entrance but ignores completion.

AI may give more people legal words, but legal words are not the same as legal capability. A person who can generate a demand letter, contract clause, research memo, internal policy, or court filing has crossed one barrier. That matters. The old drafting barrier was real. But the harder barriers remain: knowing what problem actually matters, what facts change the answer, what evidence exists, what rules apply, what risks are acceptable, what deadlines control, what the counterparty will do, and what the legal system can actually provide.

The justice gap does not disappear merely because the document becomes easier to create. It may move.

The visible gap narrows at the front end. More people can produce legal-looking papers. The hidden gap may widen at the back end. Well-funded users can buy better AI systems, more source access, longer context, more agents, more retrieval, more verification, more expert review, and more human lawyers supervising the output. They can run the problem again and again until the weak points surface.

One user may use AI to draft a letter or clause. Another may use AI agents to review the whole file, summarize the correspondence, compare the contract set, check the authorities, map the factual gaps, model the negotiation, identify regulatory issues, and run every draft through a source and record check.

Both users are using AI. They are not using the same machine in any meaningful sense.

The early court data already points in this direction. In a March 2026 working paper, Anand V. Shah and Joshua Y. Levy studied more than 4.5 million non-prisoner federal civil cases and 46 million PACER docket entries from fiscal years 2005 through 2026.[13] They found that self-represented federal civil cases rose from a long-run average of about 11% to 16.8% in fiscal year 2025, with the increase concentrated in case types characterized by more formulaic document production and absent from more complex, attorney-intensive categories.[14] They also found that the total volume of docket entries per court generated by self-represented cases in the first 180 days rose 158% from pre-AI means to 2025, while AI-generated text appeared in more than 18% of sampled complaints in early 2026.[15]

Litigation is only one legal domain, but it is where the data is becoming visible first. The lesson is broader than court filings. Cheaper legal expression can increase entry, volume, and activity without creating a matching supply of judgment. The words get easier to produce. The institution still has to absorb them. Someone still has to decide what they mean.

The old advantage was: I can pay more lawyers for more hours.

The new advantage may become: I can pay more lawyers, using more agents, consuming more compute, for more passes at the same legal problem.

That is not legal equality. It is the old inequality with a new unit of measure.

IX. The Compute Layer Is Not Neutral

Digital access tools can improve justice, but digital systems can also deepen exclusion if they are built for people who already know how to use the system. The OECD's 2025 Toolkit for Access to Justice and People-Centred Justice Systems recognizes the opportunity in digital transformation, including online dispute resolution, digital legal information platforms, virtual legal aid, and AI-driven triage tools.[16] But it also warns that rigid digital-first strategies can sideline people with limited legal knowledge, low literacy, disabilities, rural barriers, low income, limited internet access, or digital exclusion.[17]

That warning applies with extra force to AI legal work.

The problem is not only whether a person can access a chatbot. The problem is whether the person can afford enough reliable AI assistance to use it well, knows when the answer is incomplete, has access to authoritative sources, can verify the output, and can survive the next procedural step. Cheap text may help the person begin. It may not help the person finish.

Compute is not neutral if some users can afford one answer and others can afford a thousand checks. It is not neutral if some users can afford only expression while others can afford the investment required for judgment. It is not neutral if one user gets a general chatbot and another gets private models connected to curated legal data, internal records, contract repositories, billing systems, litigation analytics, expert systems, and lawyers trained to supervise them. It is not neutral if the cheaper user experiences AI as a drafting tool while the richer user experiences AI as a legal operations layer.

That is why the justice gap has a new unit of measure. It is still money. It is still lawyer time. But now it may also be tokens, agents, context windows, retrieval passes, source access, and verification loops.

X. The Public Receipts Are Already Here

The point is not speculative. The industry is already producing receipts.

The first public testimony comes from the vendors themselves. TechCrunch reported in January 2025 that Sam Altman said OpenAI was losing money on its $200-per-month ChatGPT Pro plan because people were using it more than the company expected.[18] That is the subscription problem in one sentence. A price that looks premium to the user can still sit below the true cost of heavy use.

The second testimony comes from enterprise operators. Fortune's CEO Daily reported in May 2026 that companies are treating tokens as a workplace currency, with executives using token consumption as a sign of adoption while also confronting developers who can burn through monthly allocations in hours.[19] That is not consumer panic. It is management discovering that AI fluency has a billable unit.

The broader commentary is moving the same direction. Industry writers have begun describing the end of the "all-you-can-eat" period in AI subscriptions and the shift toward visible usage economics.[20] Those pieces are useful not because they prove the exact future price of AI, but because they show the market noticing the same transition at the same time.

OpenAI's API pricing page lists specific prices per million tokens for its models and says API use is billed separately from ChatGPT subscriptions.[21] Anthropic's usage-credit documentation says Team and seat-based Enterprise users can continue after included limits, but later usage is billed at standard API rates.[22] GitHub is replacing premium request units with AI Credits and aligning usage with token consumption.[23] These are not academic abstractions. They are product rules.

There are also public examples of scale. A report on Peter Steinberger's OpenClaw project described 100 Codex instances running for 30 days, with a reported OpenAI API bill of about $1.3 million and 603 billion tokens across 7.6 million requests.[24] The same report stated that disabling a faster execution mode would reduce the monthly cost to about $300,000, still a striking figure for one extreme agentic workflow.[25]

No ordinary law office should treat OpenClaw as its expected bill. That would be a category mistake. But it is a useful stress test. It shows what happens when agentic work removes the human typing speed limit.

Academic research points in the same direction. A 2026 arXiv study of agentic coding tasks found that agentic tasks consumed 1000 times more tokens than code reasoning or code chat, that token usage could vary by up to 30 times on the same task, and that more token use did not necessarily translate into greater accuracy.[26]

The legal translation is obvious. A legal agent reviewing a file, checking sources, comparing options, and drafting work product is not a longer chatbot answer. It is a compute-intensive process with uncertain cost and uncertain quality unless it is managed.

XI. Enterprise Pricing May Be the Honest Price

Enterprise pricing sounds like the expensive version of AI. In some ways it is. But it may also be the honest version.

A consumer subscription hides the meter from the user. An enterprise system may expose it. It can show usage, set budgets, pool credits, apply rate cards, enforce permissions, provide admin controls, manage privacy, and let a firm understand what different workflows cost. That may feel less magical, but magic is not a management strategy.

OpenAI's help center describes credits for advanced features, shared credit pools for Enterprise and Edu workspaces, usage alerts, overage limits, and Business workspaces that can add credits for Codex activity.[27] Anthropic's usage-credit documentation lets Team and seat-based Enterprise organizations set spend limits and continue beyond included limits through usage credits billed at standard API rates.[28] GitHub's move to AI Credits similarly gives organizations budgets, included pools, and the option to allow or cap additional usage.[29]

For legal work, this is not procurement trivia. A firm that cannot see usage cannot price matters correctly. A firm that cannot separate casual drafting from deep research cannot know whether AI is saving money or merely moving cost into a hidden account. A firm that cannot audit usage cannot know whether a new workflow is an efficiency gain or a new form of uncontrolled spend.

The cheap plan teaches the habit. The enterprise plan reveals the habit.

XII. What Lawyers Should Watch

Lawyers do not need to become AI economists. They do need to stop pretending the subscription price is the cost of the work.

The useful questions are practical. What work are we actually asking the model to do? Is this drafting, research, record review, citation checking, contract comparison, client communication, or strategy? How much source material must be read? How many times does the workflow need to run before the answer is reliable? Who is verifying the output: a lawyer, another model, a legal database, or no one at all? Who can see the usage before the bill arrives? What happens if the model is wrong?

Those questions matter because legal work does not fail in the abstract. It fails in a contract, a lien, a deadline, a closing, a sanction, a lawsuit, a bankruptcy, a family dispute, a regulatory audit, a broken deal, or a client's balance sheet.

The first draft may be cheap. Competence is not.

XIII. Conclusion: The New Math of Access

AI did something real: it made legal words cheap.

That should not be minimized. Cheap legal words can help people understand problems, organize facts, draft first versions, and get unstuck. They can help lawyers move faster and help clients participate more intelligently in their own matters.

But cheap legal words are not reliable legal work product, and the feeling of access is not the same thing as sound judgment.

The system was never balanced. Money already bought better lawyers, more hours, deeper research, more options, more pressure, and more professional attention. AI does not erase that. It changes the unit of advantage. The new unit may be tokens, credits, context, agents, retrieval, source access, and compute layered on top of human legal judgment.

That is the new math of access. AI may narrow the drafting gap while widening the justice gap.

The legal profession should use AI. It should use it aggressively, intelligently, and honestly. But it should not mistake a subsidized drafting experience for a new economics of justice.

The price of legal-looking language has collapsed.

The price of knowing what it should say, whether it is true, whether it is useful, and what happens next has not.

AI lowered the cost of expression.

It did not lower the investment required for sound legal judgment.

ENDNOTES

1. Legal Servs. Corp., *The Unmet Civil Legal Needs of Low-Income Americans (2022)*, The Justice Gap Report, https://justicegap.lsc.gov/resource/2022-justice-gap-report/ (last visited May 26, 2026). back
2. World Justice Project, *Measuring the Justice Gap*, https://worldjusticeproject.org/our-work/research-and-data/measuring-justice-gap (last visited May 26, 2026). back
3. *The Matrix* (Warner Bros. 1999). back
4. *A Civil Action* (Touchstone Pictures 1998), quoted in IMDb, *A Civil Action (1998) - Quotes*, https://www.imdb.com/title/tt0120633/quotes/?item=qt3414668 (last visited May 26, 2026). back
5. OpenAI, *API Pricing*, https://openai.com/api/pricing/ (last visited May 26, 2026). back
6. Anthropic, *Claude API Pricing*, https://platform.claude.com/docs/en/about-claude/pricing (last visited May 26, 2026); Anthropic, *Manage Costs Effectively*, Claude Code Docs, https://code.claude.com/docs/en/costs (last visited May 26, 2026). back
7. Google, *Gemini Developer API Pricing*, Google AI for Developers, https://ai.google.dev/gemini-api/docs/pricing (last visited May 26, 2026). back
8. Mario Rodriguez, *GitHub Copilot Is Moving to Usage-Based Billing*, GitHub Blog (Apr. 27, 2026), https://github.blog/news-insights/company-news/github-copilot-is-moving-to-usage-based-billing/. back
9. Mario Rodriguez, *GitHub Copilot Is Moving to Usage-Based Billing*, GitHub Blog (Apr. 27, 2026), https://github.blog/news-insights/company-news/github-copilot-is-moving-to-usage-based-billing/. back
10. LexisNexis, *The LexisNexis Law Schools Application Form*, https://www.lexisnexis.com/grad-access/Default.aspx (last visited May 26, 2026). back

11. Carolyn Bach, *LexisNexis Rolls Out Free Access to Lexis+ AI for Law Students*, LexisNexis (Jan. 12, 2024), https://www.lexisnexis.com/community/insights/legal/b/product-features/posts/lexisnexis-rolls-out-free-access-to-lexis-ai-for-law-students. back

12. Thomson Reuters, *Law School Portal*, https://lawschool.thomsonreuters.com/ (last visited May 26, 2026). back

13. Anand V. Shah & Joshua Y. Levy, *Access to Justice in the Age of AI: Evidence from U.S. Federal Courts* 1 (Mar. 20, 2026) (unpublished manuscript), https://ssrn.com/abstract=6766859. back

14. Shah & Levy, supra note 12, at 1. back

15. Id. back

16. OECD, *Toolkit for Access to Justice and People-Centred Justice Systems* 72 (2025), https://doi.org/10.1787/aecf7f78-en, https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/toolkit-for-access-to-justice-and-people-centred-justice-systems_b1dd9b25/aecf7f78-en.pdf. back

17. Id. at 72-73. back

18. Kyle Wiggers, *OpenAI Is Losing Money on Its Pricey ChatGPT Pro Plan, CEO Sam Altman Says*, TechCrunch (Jan. 5, 2025), https://techcrunch.com/2025/01/05/OpenAI-is-losing-money-on-its-pricey-chatgpt-pro-plan-ceo-sam-altman-says/. back

19. Diane Brady, *CEOs Are Handing Out AI Tokens Like Paychecks and Figuring Out How to Justify the Spend*, Fortune (May 20, 2026), https://fortune.com/2026/05/20/ceos-ai-tokens-paychecks-justify-spend/. back

20. *The Time Bomb Went Off: AI's All-You-Can-Eat Era Just Ended in Real Time*, The State of Brand, https://www.thestateofbrand.com/news/ai-subscription-price-subsidiation-ending (last visited May 26, 2026); Shrsv, *The Subsidy Era Is Over: A Reality Check on AI-Powered Dev Tool Pricing*, DEV Community, https://dev.to/shrsv/the-subsidy-era-is-over-a-reality-check-on-ai-powered-dev-tool-pricing-51dn (last visited May 26, 2026). back

21. OpenAI, supra note 4. back

22. Anthropic Help Ctr., *Manage Usage Credits for Team and Seat-Based Enterprise Plans*, https://support.claude.com/en/articles/12005970-manage-usage-credits-for-team-and-seat-based-enterprise-plans (last visited May 26, 2026). back

23. Rodriguez, supra note 7. back

24. Alina Maria Stan, *OpenClaw Creator's $1.3 Million Monthly OpenAI Bill Reveals the Real Cost of Autonomous AI Coding at Scale*, The Next Web (May 18, 2026), https://thenextweb.com/news/openclaw-peter-steinberger-1-3-million-openai-token-bill. back

25. Id. back

26. Longju Bai et al., *How Do AI Agents Spend Your Money? Analyzing and Predicting Token Consumption in Agentic Coding Tasks* 1 (Apr. 29, 2026) (unpublished manuscript), https://arxiv.org/abs/2604.22750. back

27. OpenAI Help Ctr., *Flexible Pricing for the Enterprise, Edu, and Business Plans*, https://help.openai.com/en/articles/11487671 (last visited May 26, 2026); OpenAI Help Ctr., *Managing Credits and Spend Controls in ChatGPT Business*, https://help.openai.com/en/articles/20001155-managing-credits-and-spend-controls-in-chatgpt-business (last visited May 26, 2026). back

28. Anthropic Help Ctr., supra note 21. back

29. Rodriguez, supra note 7. back