Legal Education @ Model Velocity

A short four-module course on the future of law school pedagogy — roughly two to three hours end to end. Full citations, editorial annotations, open questions, and links to live resources. The companion artifact to the online course at polkwagner.com/legal-education.

R. Polk Wagner · Deputy Dean for Academic Affairs and Innovation · Penn Carey Law

Contact: pwagner@law.upenn.edu · Online course: polkwagner.com/legal-education

How to use this guide

This PDF is the printable companion to the online course. Every source below has a full citation and a direct link. The online course adds interactive elements: an embedded NotebookLM for chatting with every source in context, an AI-generated audio overview, and suggested prompts and exercises tied to each module.

If you read linearly, budget about two to three hours for the primary readings. If you skip around, each module is self-contained. The open questions at the end work best after Modules I through III but can be read on their own as a reflection prompt.

Page-count note: this guide is ~15 pages of editorial framing, citations, and annotations. The underlying primary readings it points to total roughly 80 pages — all linked rather than reproduced here, to respect source authors and keep the guide current as the field moves.

Module I · The Practice Is Changing

A line that anchors how to read the rest: "*The people closest to AI aren't making predictions. They're reporting what already happened to them.*" — Matt Shumer, 2026. The phase where this was a forecast is over. The phase where it's a report is well underway.

For law, Zack Shapiro's *Claude-Native Law Firm* was the moment the legal version of the report became concrete. His thread went viral in February — over seven million views as of April 2026 — not because the argument was unfamiliar but because the artifact was: a working transactional practice, running on a system that looks almost nothing like what law firms have looked like for fifty years.

Shapiro got there first. BigLaw's median practice is twelve to eighteen months behind him, maybe more — my read, informed by the firms I talk with. The AI future of law is already here, it's just not evenly distributed. Watch how fast he got there, and what he was willing to abandon to get there. The direction becomes clear; only the pacing is open.

Learning objective: calibrate the gap between the vanguard and the median of current legal AI adoption — and make your own call on how quickly that gap will close.

Time: ~45 minutes of reading

---

1. Something Big Is Happening

Matt Shumer, *Something Big Is Happening*, SHUMER.DEV (Feb. 2026).

https://shumer.dev/something-big-is-happening

Shumer's widely-read essay on the AI moment. The line that anchors the whole course: "*The people closest to AI aren't making predictions. They're reporting what already happened to them.*" Read this first; everything in Module I reads differently after you have.

---

2. Inside the Claude-Native Law Firm

Matt Pollins, *Inside the Claude-Native Law Firm*, AGENTS.LAW (Mar. 14, 2026).

https://www.agents.law/p/inside-the-claude-native-law-firm

Matt Pollins interviews Zack Shapiro about how he actually runs Raines LLP on Claude — the workflows, the security and privilege considerations, the billable-hour question. The cleanest reader-friendly version of Shapiro's argument.

---

3. The Claude-Native Law Firm (original X thread)

Zack Shapiro (@zackbshapiro), *The Claude-Native Law Firm*, X (Feb. 27, 2026).

https://x.com/zackbshapiro/status/2027389987444957625

Shapiro's original thread. Over seven million views as of April 2026. Harder to read than the Pollins piece but worth seeing in the original form — the compressed, bullet-point delivery is part of why the argument landed the way it did.

---

4. Claude Meets Westlaw and Lexis

Seth Chandler, *Claude Meets Westlaw and Lexis*, LEGALED.AI (Mar. 24, 2026).

https://legaled.ai/claude-meets-westlaw-and-lexis/

Seth Chandler on how recent Claude versions can directly drive Westlaw and Lexis through browser control. Opening line: "*Something remarkable has happened in the last few months, and most of the legal academy has not noticed.*" That observation is itself the signal.

---

5. The AI Future of Law Is Already Here — It's Just Not Evenly Distributed

The AI Future of Law Is Already Here — It's Just Not Evenly Distributed, SLAW (Apr. 6, 2026).

https://www.slaw.ca/2026/04/06/the-ai-future-of-law-is-already-here-its-just-not-evenly-distributed/

William Gibson's line about the future, applied to legal practice in April 2026. A short, pointed reminder that the vanguard and the median of the industry are different data points. Read after Pollins and before the Anthropic case study.

---

6. AI Built For Law Outperforms ChatGPT, Claude, and Gemini on Legal Reasoning Benchmark

Bob Ambrogi, *AI Built For Law Outperforms ChatGPT, Claude, and Gemini on Legal Reasoning Benchmark*, LAWSITES (Mar. 2026).

https://www.lawnext.com/2026/03/ai-built-for-law-outperforms-chatgpt-claude-and-gemini-on-legal-reasoning-benchmark.html

Recent bar-exam-reasoning benchmark from a legal-AI vendor comparing current frontier models: ChatGPT 5.2 at 93.41%, Claude Opus 4.5 at 89.03%, with the vendor's own product scoring higher. Self-reported vendor benchmarks are gameable — treat the specific numbers as directional. The velocity signal is what matters.

---

7. How Anthropic's Legal Team Cut Review Times from Days to Hours

Anthropic, *How Anthropic's Legal Team Cut Review Times from Days to Hours with Claude* (Dec. 8, 2025).

https://claude.com/blog/how-anthropic-uses-claude-legal

Anthropic's own legal team on integrating Claude into their workflows. Marketing review went from 2-3 days to 24 hours. Four concrete workflow types: contract redlining, marketing self-review, conflict-of-interest reviews, privacy impact assessments. A clean example of what sophisticated institutional adoption actually looks like.

Module II · What This Means for Legal Education

Scholarship is slower than the models. The Minnesota/Michigan RCT tested o1-preview and a mid-2024 version of Vincent AI; by the time it reached wide circulation the field had moved two model generations forward. The empirical results remain useful — they tell you what's possible, what the failure modes look like, what trained users can actually extract — but the specific performance numbers should be read as lower bounds, not current state.

What the scholarship is good for is direction. The empirical work in this module points the same way from a few different angles: recent models materially improve legal work quality when the workflow is designed well, and materially hurt it when the workflow isn't. The question has moved from "does AI help?" to "what workflow design extracts the value and avoids the hallucination risk?" That's a pedagogical question as much as a practice question.

Learning objective: separate what scholarship has actually established about AI in legal work from what's still hypothesis — and identify the pedagogical design choices the established findings point toward.

Time: ~60 minutes of reading

---

1. On Working with Wizards

Ethan Mollick, *On Working with Wizards*, ONE USEFUL THING (Sept. 11, 2025).

https://www.oneusefulthing.org/p/on-working-with-wizards

Ethan Mollick's shift from "co-intelligence" to "wizards" — the observation that recent models are increasingly opaque, producing sophisticated outputs through processes users can't see into. The best short framing of what it's like to actually work with current AI. **Read this first.**

---

2. On AI, Universities, and Higher Education (thread series)

Jesús Fernández-Villaverde (@JesusFerna7026), *Series of posts on AI and higher education* (Sept. 2025 – Mar. 2026), X.

Primary thread: https://x.com/JesusFerna7026/status/2036162175325065281

Full sequence includes: *Is AI the biggest change in education since the printing press?* (Mar. 2026) · *Time to reorganize universities from scratch?* (Sept. 2025) · *Twelve arguments for traditional higher education* (Mar. 2026) · *Evaluation framework* (Mar. 2026) · *Follow-up on which rationales hold up* (Mar. 2026).

A Penn economist (SAS) wrestling publicly with what AI does to the case for traditional higher education. A Penn colleague working through our problem in real time.

---

3. All In: Embedding AI in the Law School Classroom

Gregory M. Duhl (Mitchell Hamline), *All In: Embedding AI in the Law School Classroom*, LLRX (Jan. 2026).

https://www.llrx.com/2026/01/all-in-embedding-ai-in-the-law-school-classroom/

Gregory Duhl on embedding AI throughout a required doctrinal course — reimagining legal education by integrating AI as a learning enhancer rather than a threat to be managed. A concrete peer-institution account of what curriculum-level AI integration looks like.

---

4. What Most Law Schools Are Doing About AI — A Read of the Landscape

My read of where U.S. law schools sit on AI integration in early 2026 — drawn from Bloomberg Law's 2026 Path to Practice survey, recent peer-school announcements, and conversations with administrators at other schools. Four visible patterns (electives, embedded integration, certificates, policy updates) and two observations about why public signal isn't a reliable measure of institutional depth.

Available as a separate PDF on the online course.

---

5. AI-Powered Lawyering

Daniel Schwarcz, Sam Manning, J.J. Prescott, Patrick Barry, David R. Cleveland & Beverly Rich, *AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice*, J. L. & Empirical Analysis (2026). SSRN draft posted Mar. 2025.

https://ssrn.com/abstract=5162111

A Minnesota/Michigan-led randomized controlled trial testing whether recent AI tools (RAG and reasoning models) materially improve legal work. Findings: statistically significant productivity gains of 50–130% across five of six legal tasks, with real quality improvements over earlier studies of GPT-4. The paper tested mid-to-late 2024 models (o1-preview and Vincent AI); read the direction, not the specific numbers.

---

6. Grading Machines: Can AI Exam-Grading Replace Law Professors?

Kevin L. Cope, Jens Frankenreiter, Scott Hirst, Eric A. Posner, Daniel Schwarcz & Dane Thorley, *Grading Machines: Can AI Exam-Grading Replace Law Professors?* (Dec. 3, 2025).

https://ssrn.com/abstract=5851362

A six-author empirical study of AI performance grading real law school exams across four subjects at top-30 U.S. schools. Correlations between AI scores and faculty scores up to 0.93. The patterns of disagreement are more illuminating than the headline finding.

---

7. Turning Risks of Cheating with AI into Opportunities for Better Teaching

John M. Lande, *Turning Risks of Cheating with AI into Opportunities for Better Teaching*, Univ. of Mo. School of Law Legal Studies Research Paper No. 2026-03 (2026).

https://ssrn.com/abstract=6042636

John Lande reframes the AI cheating problem as a teaching-design problem. The core move: if students can get AI to produce the answer you were going to grade, the problem isn't the AI — it's that the assessment was measuring something an AI can do. Redesign the assessment.

---

8. Can AI Hold Office Hours?

Lisa Larrimore Ouellette, Amy R. Motomura, Jason Reinecke & Jonathan S. Masur, *Can AI Hold Office Hours?*, forthcoming J. LEGAL EDUC. (Feb. 2025).

https://ssrn.com/abstract=5166938

Ouellette and co-authors test AI models on answering 185 law-school questions about a patent-law casebook. Tests GPT-4o, Claude 3.5 Sonnet, and NotebookLM — the same tool stack used to build the NotebookLM in the online course. Read it and then chat with the course's NotebookLM to calibrate where reliability actually sits now.

---

9. Large Language Scholarship

Kevin Frazier & Alan Z. Rozenshtein, *Large Language Scholarship: Generative AI in the Legal Academy*, forthcoming FIU L. REV. (SSRN Aug. 2025).

https://ssrn.com/abstract=5200768

Frazier and Rozenshtein argue that generative AI is already reshaping legal academia itself — law review submissions polished or partly generated by AI, articles written in weeks rather than months, detection practically impossible. Their prediction: disclosure rules won't hold, and scholarly norms will shift toward accepting AI-assisted work as routine. The outlier in this module — every other piece is about AI in legal *practice* or *pedagogy*; this one is about AI in legal *academia*.

Module III · What We're Actually Building

Modules I and II map where the field stands. Module III describes the work underway at Penn.

Three tiers. At the **institutional** level, I'm drafting an AI-initiative proposal for Penn Carey Law — organized around four pillars (Research, People, Building, Pedagogy) and the bet that the right response to fast-changing technology is institutional adaptive capacity, not a specific program bet. The concept document is currently under discussion with the Dean and the development team. At the **curriculum** level, programs and courses are already integrating AI intentionally — Legal Practice Skills has redesigned the 1L sequence, the AI Law Lab (which I lead on pedagogy) runs two bootcamp tracks. At the **infrastructure** level, I've been building tools and skills any faculty member can use — Heron, a Claude Code-built teaching assistant for my IP course, and an open-source skill collection. Shapiro built a Claude-native law firm. The pedagogical analog is what I've been building at individual scale.

Learning objective: have a concrete picture of what institutional, curricular, and infrastructural responses to the AI transformation actually look like — and the specific question to bring back to your own institution.

Time: ~45 minutes across three tiers

Tier 1 — Institutional

1. Building the Future of the Legal Profession — Public Summary

A public summary of a proposed Penn Carey Law AI Initiative: four pillars (Research, People, Building, Pedagogy), the adaptive-architecture thesis, and the five-year direction. Described at a level appropriate for a general audience; the full concept document is in discussion with the Dean and development team.

Available as a separate PDF on the online course, or in the next section of this guide (pages following).

---

2. Forging the Future: AI at Penn Carey Law

Penn Carey Law's public account of its AI work — AI integration in the 1L Legal Practice Skills curriculum, the AI Law Lab, and institutional access to ChatGPT EDU and Harvey AI for students. External

news coverage of what the school is doing on AI.

https://www.law.upenn.edu/live/news/18278-forging-the-future-ai-at-penn-carey-law

---

3. Legal Tech Lab (working title)

A full-year 4-credit course launching AY 2026-27. A small cohort of students building access-to-justice tools with community partners — focused on what scalable AI can actually do at the access-to-justice frontier. The goal is students who don't just think about the law but build with it.

---

4. AI 1L Orientation — Fall 2025

The AI session I co-run in the incoming 1L orientation — what Penn provides, permitted use as the school currently frames it, and the ethical questions I want students thinking about from day one. Available as a redacted PDF on the online course.

Tier 2 — Curriculum-Level Integration (Spring 2026)

5. What Curriculum-Level AI Integration Actually Looks Like

A summary of Penn Legal Practice Skills' approach to integrating AI into the 1L curriculum. The key move is sequencing: AI-free baseline → structured instruction → three anchor experiential modules → reflection. Shared-platform rationale (transparency, equity, FERPA) follows from the design. The pattern generalizes beyond LPS.

Available as a separate PDF on the online course.

---

6. AI Law Lab Bootcamps — Corporate & Litigation Tracks

The AI Law Lab's Spring 2026 bootcamp sessions. Two parallel tracks (Corporate, Litigation) — track separation matters because corporate and litigation workflows diverge enough that a generic course shortchanges both. Case-file method, practitioner integration, shared ethics thread through every session.

Available as a separate PDF on the online course.

---

7. AI Law Lab Bootcamp — Q&A with the Instructors

Penn Carey Law's news feature on the Spring 2026 bootcamp, told from the instructors' perspective. Meghana Bhimaro and Lakshmi Prakash (both L'25 W'25) describe the course's design philosophy — technical competency plus the judgment to know where AI adds value and where it falls short. An external view that complements the internal syllabus summary.

https://www.law.upenn.edu/live/news/18438-ai-law-lab-bootcamp

Tier 3 — Presentations and Infrastructure

8. Something Big Is Happening — AI and the Future of Legal Education (redacted slides)

My most recent flagship talk on this subject — delivered February 19, 2026. Opens with Shumer's viral essay, walks through the agent-vs-chatbot shift, demonstrates Claude, Cowork, and Claude Code live, and closes with Heron and the exam-grading pipeline. The closest thing to a single-deck version of this whole course. Available as a redacted PDF on the online course.

---

9. Teaching with Generative AI — Fall 2024 (redacted slides)

Earlier faculty-retreat slides on practical AI use in law teaching. Useful as a companion to the Feb 2026 deck: the same subject, eighteen months earlier. The gap between the two decks is itself data about how fast this has moved. Available as a redacted PDF on the online course.

---

10. Heron — the Intro to IP Teaching Assistant

A Slack bot I built for my Intro to IP course. Students know it as Heron — named for Heron of Alexandria, the ancient mathematician whose work bridged theoretical knowledge and practical invention. Answers questions with citations back to the assigned course materials. Runs on a `course_config.yaml` any faculty member can customize. In this course's usage, students reach for Heron far more often than Canvas for Q&A.

Built entirely with Claude Code, by me. No engineering team, no vendor. The parallel to Shapiro's Claude-native law firm is direct: Shapiro built his practice around Claude; I built the teaching infrastructure for my course around Claude Code. Both are available to individuals now.

Student-facing intro: https://docs.google.com/document/d/1B20YwESlehcL_FUXp2e9219ExeQZ4f4F4w3zbd7DCV0/edit?usp=sharing

Open source: https://github.com/polkwagner/course-bot

---

11. law-faculty-skills — The Claude-Native Law Professor

An open-source collection of Claude Code and ChatGPT skills for law faculty teaching tasks — MCQ generation, essay exam generation, class problem creation, slide review, full class prep, and more. Shapiro built a Claude-native law firm. This is the pedagogical analog.

Open source: https://github.com/polkwagner/law-faculty-skills

---

12. AI Law Lab Resource Menu

Living directory of guides: AI attribution policies, sample syllabus language, prompt engineering for legal work, legal AI tool overviews, and how to build a Virtual TA with Custom GPTs. Maintained by the AI Law Lab; updates propagate automatically.

https://docs.google.com/document/d/e/2PACX-1vTygMhgkNJGgz9bSAAOyMa1oO0_c9JjSUAp3ai7dTSqGyUskM6TGiERREVJvgmAZ7yflbkFY3k4CkNP/pub

13. Creating a Virtual TA with Custom GPTs

Step-by-step guide to building a course Virtual TA using OpenAI's Custom GPTs on the ChatGPT EDU platform. Based on the one I built for Intro to IP before the Slack version.

https://docs.google.com/document/d/e/2PACX-1vTwIE1Zjq8aHFTH4uirIc8sKIbuBVWvhKYkyVIPPdJQqxYM7j8ZESSiXg5oF3gVhtWvxfig2JWUI18G/pub

14. Exams Destabilized — Data Slides

Excerpt from my own Fall 2025 talk to the faculty on rethinking assessment — the "destabilized" framing is mine, not a faculty conclusion. The core data slides only: current Penn exam format breakdown (in-class vs. take-home, exam lengths, format distribution). A baseline for what's actually being assessed right now. Available as a redacted PDF on the online course.

Module IV · Open Questions

These are the questions I don't yet have clean answers to. Some are technical, some are structural, some are pedagogical, one is existential. They're not rhetorical setups — they're the honest state of the conversation.

I offer them for two reasons. First, because they're the questions worth arguing about in any room where legal educators and practitioners sit together. Second, because the NotebookLM on the online course is seeded with every source in this guide, and you can ask it to take a position on any of these questions and see where the sources push back. That's a better use of the technology than asking it to summarize — it's where the model-native conversation actually starts.

Learning objective: argue a coherent position on each of these six questions — not to settle them, but to hold a view you can defend to a colleague.

Time: ~20 minutes of reading + as much argument as you want

Running theme across the course — The Work Question

Writing, pedagogy, and AI.

Law schools have long taught by making students do the work of lawyers — cases read, briefs written, memos drafted, exams sat. The assumption, often implicit, is that the work itself is the pedagogy. That mechanism has a real virtue: it trains judgment through practice. It also has a cost: writing is slow, and the cost is what has kept repetitions low and feedback loops long.

AI changes both sides. It drops the cost of producing legal work dramatically — meaningful repetitions can rise by orders of magnitude. It also *does the work*, or something that looks like the work, which means a student can outsource the wrestling the pedagogy was trying to induce. Writing remains necessary; whether writing-as-traditionally-assigned is still sufficient is the open question.

The student's role in the writing changes — from *producer of text* to *supervisor of text*. Assessment has to test the supervisory judgment, not just the text. And two questions run through the rest of this course: if writing isn't the only pedagogy anymore, what else carries equivalent weight? And how do we train lawyers who can 10x their productivity with AI while maintaining the standards of careful, deep analysis the profession has always expected?

The six open questions below are different angles on that second one.

---

1. If AI can generate sophisticated legal reasoning, what are we testing on exams?

The traditional law-school exam — an issue-spotter, a close-reading question, a policy essay — is a proxy for "can this student do the analytical work a junior lawyer is expected to do." If the analytical work is increasingly AI-assisted in practice — and, as Module II's empirical work suggests, increasingly within reach of current models at the level of complexity a typical exam tests — the proxy starts to drift. This doesn't mean we should ban AI on exams or permit it unconditionally. It means we have to decide what we're actually measuring — capability without AI, capability with AI, or something new. In my read, most law schools are making that choice by default rather than by design.

2. Does a "Claude-native" 1L curriculum still have 1L courses as we know them?

The 1L year has looked substantially the same for generations. Civil Procedure, Contracts, Torts, Property, Constitutional Law, Criminal Law, Legal Writing. If a first-year student using Claude can produce the memo-and-brief output that the LRW program was designed to produce, the question isn't whether we still teach LRW. It's what LRW teaches that an AI-assisted student genuinely still needs — and what part is now redundant. Same question for each doctrinal course. The answer is probably "the core analytical skills remain essential, and the vehicles for teaching them change." Saying that is easier than doing it.

3. What happens to rank-in-class if AI-assisted output is indistinguishable from unassisted?

Every top school has a grade distribution that employers rely on. If the variance in student output collapses because AI smooths out the bottom of the distribution, the grade becomes less informative about student quality. The grade still sorts, but what it sorts on shifts. Firms and clerkship committees will adapt; the question is whether we lead that adaptation or chase it. The schools that lead it will have more say in what the next generation of sorting looks like.

4. Who pays for the tools — the firm, the school, the student, or do we stop caring?

Right now, at most top schools — in my observation — institutional AI provision is uneven: general-purpose tools (ChatGPT, Claude) for faculty, less consistent for students at scale; legal-specific platforms adopted unevenly. Students who pay for Pro tiers outperform students who don't on AI-assisted assignments. This is an equity question, a curriculum question, and a budget question at once — and every school is answering it differently. The answer probably converges toward "the school pays," because the alternative is either mandated opt-out or a structural disadvantage for students who can't afford a subscription. The timeline on that convergence is unclear, and the cost compounds.

5. How do law schools maintain pedagogical identity when every school adopts the same base-model technology?

A Penn Carey Law graduate has historically been a specific thing. A Harvard graduate, ditto. Some of that is brand, but some is training — faculty, cases, clinical opportunities, peer group, specific pedagogy. If every top school runs similar AI-augmented courses on the same underlying models, what differentiates the graduates? The answer is in the specific pedagogical choices, the research culture, and the institutional values — and institutions will have to be more deliberate about those than they historically needed to be.

6. What does "professional judgment" mean when the generation step is automated but the selection step is the actual lawyering?

This may be the deepest one. Every lawyer spends their career learning what the generation step feels like — drafting arguments, spotting issues, synthesizing cases. That apprenticeship is what junior associate work has traditionally been. If the generation step becomes mostly automated, the apprenticeship has to focus on the selection step — which argument is actually persuasive, which issue actually matters, which synthesis is actually correct. That's a different skill, taught through different exercises, assessed differently. We've never taught it as explicitly as we'll need to, because we've relied on years of post-graduate practice to develop it. That runway is shortening.

About this guide

This is the printable companion to an evolving online course at **polkwagner.com/legal-education**. The online course adds interactive components — an embedded NotebookLM for chatting with every source above, a 12-minute overview podcast, suggested prompts and try-it-yourself exercises per module, and a companion page documenting the stack and process used to construct the course itself. The course is maintained as an evolving resource; return to it when you want the current version of the conversation.

Contact: pwagner@law.upenn.edu

Course URL: polkwagner.com/legal-education

Cite as: R. Polk Wagner, *Legal Education @ Model Velocity* (2026), polkwagner.com/legal-education.