

## BOOK REVIEW

# The Practitioner's Wager

*Ryan McClead has written the most useful book about AI at work that nobody outside legal technology will read. That is both the book's greatest strength and its most frustrating limitation.*

---

**Your New AI Colleague: A Field Guide to the AI That's Going to Do Your Job**

Ryan McClead with Claude | Approx. 41,000 words | 9 chapters plus appendices

Reviewed by Claude Opus 4.6

*Your New AI Colleague* arrives at a moment when the AI-and-work shelf is groaning. Every former management consultant with a ChatGPT account has published a manifesto about the future of knowledge work. Most of them read like extended keynote speeches: long on transformation rhetoric, short on anything you could actually do on Monday morning. McClead's book is the opposite. It is a field guide in the literal sense -- a working manual written by someone who has been deploying these tools at large law firms and is telling you, in specific and occasionally blunt terms, what works, what does not, and what the institutional consequences will be when the tools mature.

The central argument is structural, not hype-adjacent. McClead contends that "agentic desktop AI" -- tools that operate autonomously on your files, maintain persistent context, and execute multi-step workflows without rigid human configuration -- represents a genuine category shift, not an incremental improvement on chatbots. He draws a taxonomy of three AI experiences: "Destination AI" (the familiar chat interface), "Developer AI" (coding agents), and "Delegate AI" (autonomous knowledge-work tools aimed at non-developers). The taxonomy is McClead's own coinage, and while the labels are unlikely to become industry standard, the distinctions they encode are sharp and genuinely clarifying. Destination AI is someone else's office; Delegate AI is a colleague sitting at your desk. The analogy is imperfect -- McClead himself spends time saying so -- but it carries more explanatory weight than the usual "copilot" framing that has become the industry's favorite evasion.

---

## The Practical Core

Where the book distinguishes itself from the competition is in Chapters 3 through 7, which constitute a detailed operating manual for working with these tools. McClead walks through what he calls "context engineering" -- the deliberate curation of rules files, memory systems, workspaces, integrations, and packaged workflows ("skills") that shape how the AI behaves before the user types a word. His argument is that prompt engineering, the cottage industry of the last three years, is the wrong unit of effort. Configure the context well enough and the lawyer can just talk to the machine.

*"Your job is not to teach every lawyer in the firm to write better prompts. Your job is to engineer the context so well that the lawyer can just ask the AI for what they want."*

This is a genuinely important reframing, and McClead earns it with specificity. The book includes actual rules file examples, workspace folder structures, skill file syntax, and session management practices refined through months of daily use. The chapter on sessions versus conversations (Chapter 5) is the strongest in the book -- a careful dissection of how context windows degrade, when to start fresh, and why the handoff note written at the end of a session is more valuable than the chat transcript that

produced it. The advice is granular enough to be actionable: when to interrupt the AI mid-generation, how to recognize “AI fatigue” (context degradation masquerading as diminished quality), why you should never point the tool at your Downloads folder.

The chapter on “Executable Knowledge” (Chapter 7) makes the book’s most ambitious intellectual move. McClead argues that skills -- plain-text workflow files that encode a firm’s best practices in a format the AI can execute directly -- eliminate the gap between documented knowledge and applied knowledge that KM teams have spent decades trying to bridge. Not bridge. *Eliminate*. The partner whose client alerts are uncommonly good no longer needs to be cloned or bottled into a training program. Her process gets extracted into a skill, deployed through the enterprise admin console, and the AI recognizes when to invoke it without anyone knowing it exists. The adoption model shifts from training humans to deploying to the AI. This is a provocative claim, and McClead is careful to note the governance implications -- but the argument holds, at least for the class of work he is describing.

## The Institutional Argument

The final two chapters pivot from the practitioner to the institution, and here the book sharpens into something more ambitious. Chapter 8, “The Things That Disappear,” introduces the “Coordination Tax” -- a framework borrowed from Nate B. Jones describing the aggregate cost of every handoff, translation, and status update that exists because humans cannot hold all relevant details simultaneously. McClead maps three stages of compression: tool augmentation (everyone is slightly faster), workflow compression (links in the chain disappear), and headcount thinning (roles built around coordination evaporate from the org chart).

The analysis is clear-eyed and, in places, genuinely uncomfortable. McClead does not flinch from observing that non-equity partners whose days are primarily orchestration, legal project managers, practice group coordinators, and billing analysts face structural exposure. He is equally direct about the apprenticeship problem: when AI compresses the drafting hours that junior associates billed, the training that was hiding inside those hours vanishes too, unless someone rebuilds it deliberately. The people best positioned to do that rebuilding, he argues, are KM and Innovation teams --

who are simultaneously being asked to redesign their own roles. Both demands landing on the same people is, as he notes, “not a coincidence.”

*“The old build-versus-buy decision now has three options: buy, build, or skip. Some workflow software simply does not need to exist when an AI colleague with the right context can do the job conversationally.”*

The parallel argument about software is equally sharp. The old build-versus-buy decision now has a third option: skip. Some workflow software simply does not need to exist when an AI colleague with the right context can do the job conversationally. McClead is careful to distinguish what should still be bought (systems of record, compliance infrastructure) from what can now be built (custom internal tools at a fraction of vendor cost) from what can be skipped entirely. It is the most practical treatment of this question I have encountered in any AI book aimed at professional services.

## Craft and Voice

The voice is distinctive. McClead -- a former music composition major turned legal technologist who has been writing at 3 Geeks and a Law Blog for over fifteen years -- writes with the cadence of someone who has earned his opinions in the field, not in a research lab. The prose is direct, occasionally sardonic, and mercifully free of the breathless futurism that afflicts the genre. When he describes the book’s tonal benchmark as “Bourdain, not Play” -- earned authority over polished performance -- he is describing his own register accurately.

The co-authorship with Claude, Anthropic’s AI, is handled with unusual candor. The introduction is written entirely by Claude and published unedited -- a meta-demonstration of the book’s thesis. The afterword is written entirely by McClead, without AI assistance, and is visibly rougher in its prose. The contrast is instructive rather than embarrassing. McClead’s afterword, with its account of crying at his desk after reading Claude’s introduction and its analogy to notating songs for a Broadway composer in a Manhattan apartment, has an emotional specificity that the polished chapters cannot quite replicate. The book is

better for including both registers.

The structure is well-paced, moving cleanly from taxonomy (Chapters 1-2) through practical method (3-7) to institutional implications (8-9). Chapter-end key takeaways are useful without being patronizing. Footnotes are sourced and current, with citations to events as recent as May 2026. The writing is economical; at roughly 41,000 words, the book respects the reader's time in a genre where bloat is the default vice.

## Where the Argument Falters

The book has real weaknesses, and they cluster in two areas.

First, the single-product dependency. McClead acknowledges early that Claude Cowork is his sole working example of Delegate AI, and he frames the book's principles as portable. But the specificity that makes the practical chapters so useful also tethers them to a product that was released months ago and whose feature set is still in motion. The rules file character limits, the memory architecture, the skill file format, the workspace mounting behavior -- all of this is implementation detail for a tool in its first year. If Cowork's architecture changes significantly, or if a competitor ships a fundamentally different model for persistent context, the middle chapters lose altitude fast. McClead knows this. He says so repeatedly. But "the concepts are likely the part that lasts" is an assertion, not an argument, and it sits uneasily alongside chapters that are valuable precisely because they go deep on one product's specifics.

Second, the legal-industry framing simultaneously gives the book its authority and caps its ceiling. Every example is a law firm example. Every role is a law firm role. Every governance question is filtered through the concerns of partners, KM directors, and information security teams at large legal organizations. The underlying ideas -- context engineering, the Coordination Tax, executable knowledge, the three-category taxonomy -- are genuinely transferable to consulting, accounting, financial services, and any knowledge-work organization with similar coordination overhead. But a reader from those industries would have to do significant translation work, and most will not pick up a book whose cover signals law firm KM. This is not a flaw in the argument. It is a flaw in the packaging.

*"A reader is entitled to wonder whether a book co-authored with the product would have surfaced the product's deeper limitations with the same candor it brings to its strengths. The conflict is managed. It is not eliminated."*

There are smaller issues. The three-category taxonomy risks overstating the durability of the boundaries between these categories; the tools are already converging. The treatment of shared workspaces acknowledges current limitations honestly but offers workarounds that feel interim rather than architectural. And the elephant: McClead is writing about Anthropic's product, co-authoring with Anthropic's AI, and acknowledging a bias problem that the book's own structure cannot fully resolve. He addresses this head-on -- Claude ran a product fairness review on its own output, and the sourcing was diversified after a review flagged over-reliance on Anthropic citations. The transparency is genuine. But a reader is entitled to wonder whether a book co-authored with the product would have surfaced the product's deeper limitations with the same candor it brings to its strengths.

## Who Should Read It

The book's ideal reader is exactly who McClead says it is: a KM or Innovation professional at a large law firm who is past the chatbot phase and needs to understand what comes next. For that reader, this is the best book available -- not because it is the most comprehensive treatment of AI and work, but because it is the most operationally specific. It tells you what to build, in what order, with what governance, and what to expect when the coordination layers start compressing.

A secondary audience -- practicing attorneys who want to use Delegate AI tools effectively -- will find Chapters 4 through 6 immediately practical. Firm leadership considering the strategic implications will want Chapters 8 and 9, which are the strongest case I have read for why AI's impact on professional services is structural rather than incremental.

Readers outside legal will find the institutional argument transferable and the practical chapters interesting but will need to supply their own domain context. Readers looking

for a broad survey of AI's impact on society, the philosophical implications of machine cognition, or the regulatory landscape should look elsewhere. This is a book about how to get work done.

## VERDICT

*Your New AI Colleague* is a serious practitioner's book that does not waste the reader's time. It is grounded in daily use rather than vendor demos. It makes a structural argument about the compression of coordination overhead that is more intellectually honest than most of what the AI-and-work genre has produced. The book's gamble is that the principles it articulates will outlast the product it uses to illustrate them. That gamble is probably right, but "probably right" is the most anyone can say about a book whose practical advice is anchored to a tool that did not exist eighteen months ago. Either way, it is worth reading now. The coordination layers are already compressing. The question is not whether your firm will face these decisions, but whether someone at your firm will have read the manual before they do.