The Art of Delegation in the Era of AI: How to Brief an AI Like You'd Brief a Colleague

Table of Contents

You Already Know How to Do This

If you've opened an AI tool, typed a quick sentence, and felt underwhelmed by what came back, you're not alone. Many research professionals try it once, get generic results, and quietly conclude that AI isn't ready for serious work. The problem isn't the technology. It's the frame.

Most people approach AI like a search engine. They type a command ("summarize this study") and expect magic. When the output is shallow, they assume they lack some technical skill they haven't learned yet. But that's not what's missing. What's missing is a mental model that unlocks what you already know how to do.

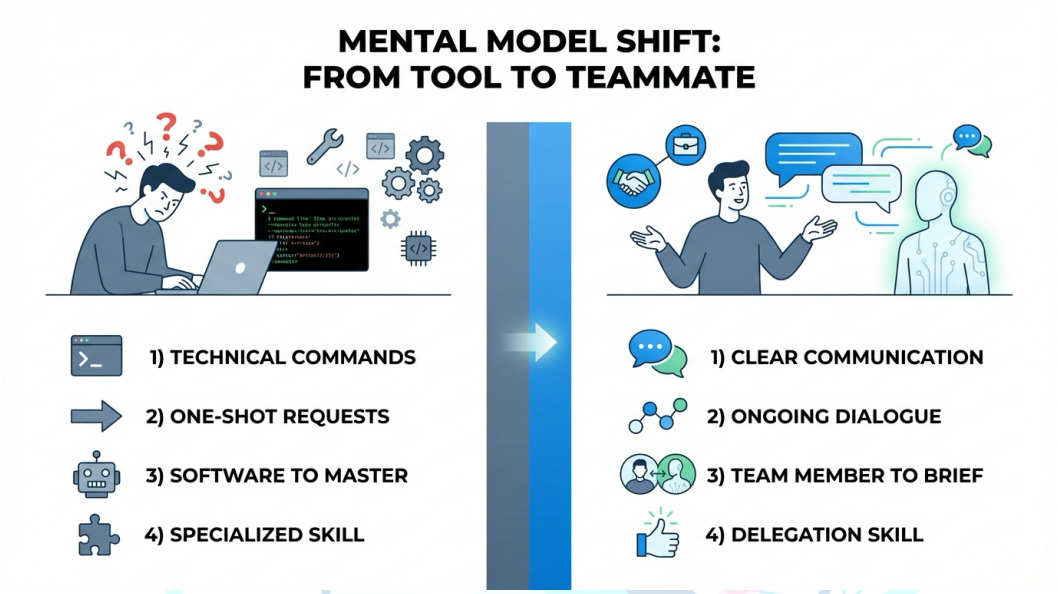

Mental Model Shift: From Tool to Teammate

The Shift: From Tool to Teammate

Here's the reframe: **Stop thinking of AI as a tool you use. Start thinking of it as a capable but completely context-free colleague you're briefing.**

When you hand a task to a new team member, you don't just say "analyze this dataset" and walk away. You explain who you are, what the project is, what stage you're at, and what you need. You describe what good looks like. You share why this matters and what decision it supports. That's not a special skill. That's just how professionals delegate work.

The same instincts apply to AI. The difference between mediocre and genuinely useful output isn't technical expertise. It's communication clarity. **If you've ever written a clear**

email to a colleague asking for help, you already have 80% of what it takes to get great results from AI.

Research on workplace delegation confirms this: effective delegation is fundamentally a communication skill, not a technical one. It requires clarity about objectives, trust in the process, and regular feedback. These are skills professionals use every day that transfer directly to AI interaction once the framing shifts.

Why the Barrier Feels Technical (But Isn't)

The barrier to AI adoption isn't skill or access. It's culture. When organizations frame AI use as a technical capability, adoption fragments. People worry that using AI makes them look less competent or that they should already know how to "prompt engineer." But the real obstacle is this cultural framing, not the absence of technical skills

Healthcare organizations that have successfully integrated AI don't treat it as a specialized skill requiring training in technical jargon. They treat it as standard professional communication, like writing a clear brief or setting expectations with a vendor. When leadership models transparent AI use and communicates that it's valued, adoption accelerates.

Old Mental Model	New Mental Model
AI is a tool I use	AI is a teammate I brief
I need to learn technical commands	I use the communication skills I have
The first output is the verdict	The first output is a draft to refine
If it's not perfect, AI isn't ready	If it's not perfect, I haven't finished the conversation

Consider this scenario: You need a one-page summary of five cost-effectiveness studies for your VP, who's deciding whether to pursue a new intervention. If you handed this to a smart junior analyst who just joined your team, you wouldn't say "summarize these." You'd explain the context (the VP's decision, your threshold, the comparator in your portfolio) and describe the format (one page, bullets, emphasis on economic data). You'd review their draft, give feedback, and iterate.

That's exactly how effective AI interaction works. The conversation style, the detail, the refinement - it's the same skill set. The only difference is that your AI "colleague" starts with zero context, so you provide it upfront.

The 80% You Already Have

Research professionals are particularly well-positioned for AI adoption because they already practice structured communication daily. When you write an evidence brief for a payer, you're tailoring the message. When you explain study findings to medical affairs, you're choosing which data points matter. When you onboard a new team member, you're setting expectations and providing background.

All of these are delegation and communication skills, and they're exactly what AI responds to. The 20% you might be missing isn't technical. It's just practice applying these skills in a new context.

You don't need to learn a new language. You just need to use the one you already speak.

Set the Scene: Why Context Changes Everything

AI interaction is about delegation, not technical wizardry. Now comes the most impactful shift: learning to set the scene before you ask for anything.

Picture this: A new analyst joins your team. You hand them a dataset and say, "Analyze this." No background. No explanation of what the project is, who needs it, or what decision it supports. They produce something generic because they can't prioritize what matters. You're frustrated. They're confused. The output is unusable.

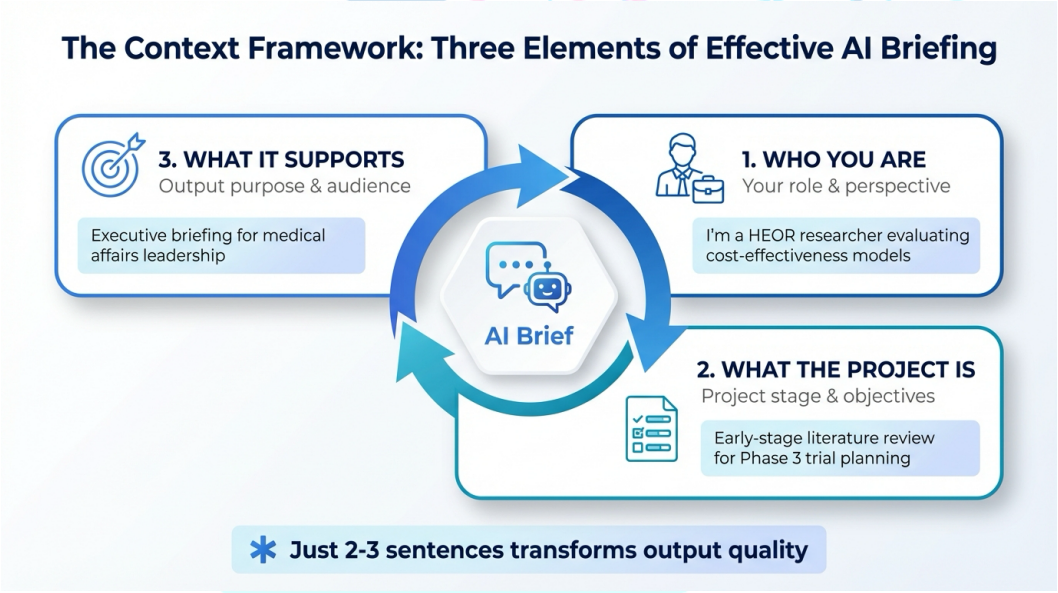

The Context Framework: Three Elements of Effective AI Briefing

You would never do this to a colleague. Don't do it to AI either.

When organizations implement structured onboarding that prioritizes context setting upfront, they experience 50% greater new-hire productivity. The reason is simple: context enables better judgment. When someone understands the landscape, they make smarter decisions about where to focus and what to emphasize.

The same principle applies to AI. A two-to-three sentence context preamble transforms output quality. Context gives AI the reference points to make judgment calls about what matters and what to prioritize.

What Context Actually Does

Think of context as the background briefing you'd give any team member before assigning work. Research on cognitive load demonstrates that when speakers provide context, listeners perform significantly better because ambiguity is reduced. When context is missing, AI defaults to generic responses.

Context shapes every choice AI makes: which details to highlight, what tone to use, how much depth to provide, and what to leave out.

Consider two versions of the same request:

Without Context	With Context
"Summarize this cost-effectiveness study."	"I'm a health economics researcher preparing a one-page brief for our VP of Medical Affairs. We're evaluating whether to add this intervention to our formulary. Our cost-effectiveness threshold is $150k/QALY. Summarize the key findings with that lens."

The first gives AI a task. The second gives AI a scenario: who you are, what you're working on, who will use this output, and what decision it supports. That scenario changes everything.

The Three Elements of Effective Context

When you brief a colleague, you instinctively provide three types of context. Do the same with AI:

1. **Who you are and what role you're in.** ("I'm a health economics researcher on the payer value team.")
2. **What the project is and what stage you're at.** ("We're preparing a formulary review for Q2. I've already gathered clinical data; now I need economic synthesis.")
3. **What this output will support.** ("This summary will inform our VP's decision about whether to pursue coverage.")

This is two to three sentences. But those sentences give AI the judgment layer it needs to prioritize correctly. When summaries include context about audience and purpose, comprehension and decision-making improve dramatically.

A Real Example: Before and After

You need to synthesize findings from five clinical papers on a diabetes intervention. Here's how most people start:

Weak Brief (No Context):

"Summarize the key findings from these five papers on diabetes interventions."

The AI produces a generic overview. It lists findings from each paper but doesn't help you make a decision because the AI didn't know what decision you're making.

Strong Brief (With Context):

"I'm a health economist preparing a value dossier for payers. We're positioning a new diabetes intervention against the current standard of care. I need to summarize the cost-effectiveness findings from these five papers. Focus on data points that support or challenge a cost-effectiveness threshold of $120k/QALY. The audience is payers who challenged our threshold last year."

Now the AI knows your role, the project, the stage, the decision context, and the audience. The output will prioritize cost-effectiveness data relative to the $120k threshold and flag outliers that might undermine your position. It won't waste space on details that don't support this decision.

The task is the same. The context changed the output from generic to actionable.

Why Context Works

Organizations that establish clear expectations and communication norms upfront during onboarding see dramatically better results. When you set the scene, you're doing what effective managers do: giving someone the information they need to succeed .

Cognitive load increases when contextual cues available to the speaker are not appropriate references for the listener. If you assume AI knows your project background, audience, or strategic goals, you're asking it to work with incomplete information.

Context setting mirrors how HEOR teams already work. When you prepare communications for different audiences (payers, healthcare providers, medical affairs), you tailor the message. You wouldn't send the same evidence summary to a health plan executive and a clinical guideline committee because the context differs. AI responds to the same tailoring when you make that context explicit.

Start With the Scene, Not the Task

The next time you open an AI conversation, resist the urge to jump straight to the request. Set the scene first. Answer these three questions in two to three sentences:

1. Who am I, and what's my role?

2. What project am I working on, and what stage am I at?

3. What will this output support or inform?

Then make your request.

You'll notice the difference immediately. The output will feel less generic and more tailored. It will emphasize what matters to your specific scenario. It will save you time because you won't need to sift through irrelevant details or rewrite sections that miss the point.

Context isn't extra work. It's the work that makes everything else work better.

Describe What Good Looks Like: Setting Output Expectations

You've given AI the context it needs. Now comes the next step: **describing what you actually want to receive.**

Most people stop too early. They provide context, then say "summarize this" or "analyze this dataset." The AI has background, but it doesn't know what shape your output should take. Should it produce a one-pager or a ten-page analysis? Bullet points or narrative? Technical language or plain English?

Without these specifications, AI defaults to generic formats. And generic rarely matches what you need.

The Instinct You Already Use

You already do this when you delegate to colleagues. When you brief a team member, you naturally say:

- "I need a one-page summary, not a full report."

- "Keep it high-level for the VP. She doesn't need methodology details."

- "Focus on cost-effectiveness results, not clinical data."

- "Make it conversational. This is for an internal deck, not a journal."

These aren't technical instructions. They're output expectations. You're clarifying format, depth, tone, and audience. **The same instincts work directly with AI**.

Vague Request	Specific Output Expectation
"Summarize this study."	"Create a 1-page summary in bullet format. Audience: VP of Medical Affairs. Tone: executive (not technical). Focus: cost-effectiveness vs. comparator."
"Analyze this dataset."	"Produce a 2-page narrative analysis. Highlight trends in patient adherence by age cohort. Include 1-2 data tables."

	Audience: health plan executives."
"Draft a summary of these five papers."	"Draft a 500-word summary comparing methods and findings. Format: short paragraphs, not bullets. Audience: HEOR peers familiar with CEA terminology."

The difference is specificity. Vague requests produce outputs that might be correct but aren't usable. Specific expectations produce drafts ready for refinement.

Four Dimensions of Output Expectations

When HEOR teams prepare communications for different stakeholders, they tailor the format and emphasis to match audience needs. A formulary dossier for payers looks nothing like a clinical manuscript. That tailoring happens because the team knows the output expectations before they start.

Apply this framework to AI. Specify four dimensions:

1. Format and Length

Do you need a one-page summary, slide deck outline, bullet points, narrative report, or comparison table? "Keep it to 300 words" or "Use a two-column comparison table" tells AI exactly what structure to use.

2. Depth and Detail

High-level or comprehensive? Executive summary or technical deep dive? "Focus on key takeaways only" or "Include methodology details and limitations" sets the depth.

3. Tone and Style

Is this for a clinical audience comfortable with technical language, or a business audience that needs plain English? "Use conversational tone" or "Maintain formal academic style" shapes the voice.

4. Audience and Purpose

Who will read this, and what do they care about? "This is for payers evaluating cost-effectiveness" or "This is for our internal team deciding which studies to cite" gives AI the lens to prioritize information.

A Real Example

You're preparing an executive briefing on a new diabetes intervention for your VP. You have five cost-effectiveness studies to synthesize.

Weak Brief:

"Summarize the cost-effectiveness findings from these five studies."

This tells AI what to do, but not what to deliver. The output might be five paragraphs, each 200 words, full of technical terminology. Accurate, but not what you need.

Strong Brief:

"Create a 1-page executive summary of cost-effectiveness findings from these five studies. Format: 4-5 bullet points highlighting key data. Tone: executive-friendly (avoid jargon). Audience: VP of Medical Affairs who approved our current comparator two years ago and wants to know if this intervention changes our value positioning. Emphasize findings relative to our $120k/QALY threshold."

Now AI knows exactly what you want. The output will be concise, non-technical, decision-focused, and formatted for quick review.

Same task. Different output expectations. Dramatically different result.

Why Specificity Matters

Research demonstrates that decision-makers more easily understand summaries tailored to specific audiences and purposes than generic comprehensive reviews. When format and audience are specified upfront, comprehension and decision-making improve.

Vague requests produce outputs that cover everything but prioritize nothing. Specific expectations produce drafts that make choices about what matters most. If you handed an analyst five papers and said "summarize these," they'd produce something. But if you said "I need a one-pager for our VP focused on cost data, not clinical outcomes," they'd produce something tailored. The second version saves revision time.

The same logic applies to AI. Specificity upfront reduces iteration cycles later.

Practical Application

Before hitting enter, ask yourself:

1. **What format do I need?** (Length, structure, visual layout)

2. **How much depth?** (Overview or detailed analysis)

3. **What tone?** (Technical, conversational, formal)

4. **Who's the audience, and what do they care about?**

Add two to three sentences answering these questions. Here's a template:

"Create a [format/length] on [topic]. Tone: [style]. Audience: [who and what they care about]. Focus on [what matters most]."

For example:

"Create a 1-page bullet summary of this HEOR study. Tone: conversational. Audience: internal team deciding which evidence to include in our value dossier. Focus on cost-effectiveness data and study limitations."

Four sentences. But those sentences transform a generic task into a specific deliverable.

Vague requests produce generic outputs. Specific expectations produce usable drafts. The difference isn't technical expertise. It's communication clarity.

You already know how to brief a colleague on what you need. Now you're applying that same instinct to AI.

Share Your Thinking: The Power of Explaining Why

You've set the scene and described what good looks like. Now comes the element that transforms AI from a task-executor into a thinking partner: **sharing your reasoning.**

Most people skip this step. They provide context, specify the format, and make their request. But they don't explain *why* this task matters or what decision it supports. Without that reasoning, AI has no judgment layer to prioritize what's most important.

When you explain the "why" behind your request, everything changes.

The Difference Is Enormous

Consider two versions of the same request:

Task-Focused Request	Reasoning-Focused Request
"Summarize the cost-effectiveness findings from this paper."	"I'm updating our value dossier for payers who challenged our $120k/QALY positioning last year. Summarize the cost-effectiveness findings, focusing on data points that support or challenge that threshold. We need to defend our position or revise it."

The first tells AI what to do. The second tells AI what you're trying to accomplish and why it matters. That reasoning gives AI the judgment layer to prioritize correctly.

If the study shows $118k/QALY, the reasoning-focused request ensures AI highlights that prominently because it directly supports your threshold defense. If an outlier shows $200k/QALY, AI flags it as potentially undermining your position.

This mirrors how delegation works with colleagues. When delegators explain the purpose behind a task, delegatees make significantly better judgment calls about priorities. When you tell a team member "I need this summary for a payer meeting where we're defending our threshold," they understand what data matters most. **AI works the same way**.

Why Purpose Changes Prioritization

When colleagues understand the decision context, they naturally make better choices about what to highlight and what to downplay. They tailor their work to support the decision you're making.

AI responds to explicit reasoning. When you explain that you're preparing a slide for a VP deciding whether to pursue an intervention, AI focuses on high-level implications, not methodology details. When you explain that you're defending a cost-effectiveness position, AI emphasizes data that supports your threshold and flags outliers.

Including your "why" transforms AI from a tool that executes tasks into a partner that helps you think strategically.

A Real Example

You're synthesizing five diabetes intervention studies for a formulary review. Your VP wants to know if it's worth pursuing given budget constraints.

Weak Brief (No Reasoning):

"Summarize the clinical outcomes from these five studies. Focus on glycemic control and adherence data."

AI produces a competent summary listing findings from each study. But it doesn't help you decide because it doesn't know what decision you're making.

Strong Brief (With Reasoning):

"I'm preparing a recommendation for our VP on whether to pursue this diabetes intervention for our formulary. Our budget is constrained, so we can only add interventions demonstrating both strong clinical outcomes *and* improved adherence (which reduces downstream costs). Summarize the clinical outcomes, emphasizing the relationship between glycemic control and adherence. Flag any studies where adherence improvements weren't sustained beyond six months."

Now AI understands your decision criteria: clinical outcomes alone aren't enough; you need adherence data tied to cost implications. AI structures the summary around that relationship and calls out the specific red flag that matters to your VP's decision.

Same task. Different reasoning. Output becomes strategic instead of descriptive.

The Three Questions That Unlock Judgment

Before making your request, answer three questions:

1. What decision will this output support?

("This will inform our VP's decision about adding this intervention to our formulary.")

2. What's at stake or what constraints matter?

("Our budget is limited, so we can only pursue interventions with strong adherence data.")

3. What would make this output more or less useful?

("If adherence gains aren't sustained past six months, that's a dealbreaker.")

Add these answers as a short reasoning preamble. Two to three sentences give AI the lens to prioritize what matters.

Here's a template:

"I'm [your role/goal]. We're trying to [decision or action]. This output will help us [specific outcome]. Focus on [what matters most] because [why it matters]. Flag anything that [specific concern]."

Example:

"I'm a health economist preparing a payer presentation. We're demonstrating that our intervention is cost-effective relative to current standard of care. This summary will support our value positioning. Focus on incremental cost-effectiveness ratios and budget impact because payers care most about total cost. Flag any studies where the comparator differs from our market's standard of care."

Four sentences transform AI from a summarizer into a strategic thought partner.

Why This Feels Unnatural (At First)

When you use Google, you don't explain why you're searching. Many people carry that habit into AI interaction. They type a task and expect magic.

But AI isn't a search engine. It's closer to a capable colleague who needs context and reasoning to do their best work. The professionals getting exceptional results treat it that way.

If explaining your reasoning feels like extra work, remember: you already do it when you delegate to colleagues. When you say "I need this for the VP meeting on Thursday" or "focus on cost data because that's what payers care about," you're providing reasoning and decision criteria.

You already have this skill. Now you're applying it intentionally.

From Task Executor to Thinking Partner

The shift from "summarize this" to "I need this summary to support a decision about X, so focus on Y because Z matters" is the difference between using AI as a tool and partnering with it as a collaborator.

When you share your reasoning, AI stops producing generic outputs and starts producing outputs tailored to your specific strategic need. It prioritizes correctly, flags what matters, and helps you think, not just execute.

That's the power of explaining why.

Iterate Like You Would With Any Draft

You've provided context, set clear expectations, and explained your reasoning. Now comes the step most people skip: **treating AI's first output as a draft, not a verdict.**

When you delegate work to a colleague, you don't expect perfection on draft one. You review, identify what's working and what's not, and provide feedback. They revise. You iterate. **AI works the same way.** But most people who conclude "AI isn't ready" stopped after the first output. They didn't finish the conversation.

The Problem With Stopping Too Early

Picture this: You brief a junior analyst to summarize five cost-effectiveness studies. They deliver a draft that's three pages when you needed one. The tone is too technical. The cost data is buried in paragraph three.

Do you conclude the analyst isn't ready? No. You give feedback: "Cut this to one page. Use bullet points. Lead with cost data, not clinical findings. Make the language more executive-friendly."

They revise. Draft two is much closer. You give another round: "Better. Move the comparator data higher. Drop the methodology section entirely." Draft three is exactly what you need.

This is normal professional collaboration. It's also how effective AI interaction works. Yet when people try AI, they often treat the first output as the final word. If it's not exactly right, they assume the technology failed. In reality, they stopped before the conversation was finished.

How Iteration Changes Output Quality

Consider this scenario: You ask AI to draft an executive summary of an HEOR study for your VP.

Draft 1: AI produces a three-page summary. Comprehensive but too detailed. Formal, academic tone. Cost-effectiveness data appears midway through.

Your Feedback: "This is too long and too technical. Cut it to one page. Use bullet points. Lead with the cost-effectiveness findings because that's what my VP cares about. Make the tone conversational, not academic."

Draft 2: One-page bullet summary. Cost data leads. More accessible tone. But the comparator information is still buried at the bottom.

Your Feedback: "Much better. Move the comparator data to bullet two. My VP approved that comparator two years ago, so she'll want to see the comparison immediately. Drop the limitations section."

Draft 3: Exactly what you need. One page. Bullets. Cost data and comparator comparison front-loaded. Conversational tone. Ready for review.

Same task. Three iterations. The difference between "AI isn't ready" and "this saved me two hours."

Iteration	What Changed	Result
Draft 1	No feedback yet	Too long, too technical, wrong emphasis
Draft 2	Feedback on length, tone, and structure	Better format, but comparator data buried
Draft 3	Feedback on sequencing and content cuts	Exactly what was needed

Research on iterative feedback loops confirms that regular feedback cycles are a cornerstone of high-performing teams. Apps that actively incorporate user feedback see retention rates climb by 40% and customer satisfaction increase by 30%. Iteration isn't a sign of failure. It's how quality work gets done.

What Good Feedback Looks Like

Vague feedback produces vague improvement. Specific feedback produces rapid refinement.

Weak Feedback: "This isn't quite right. Try again."

AI doesn't know what to fix. It might shorten the output, change the tone, or restructure entirely, but you're guessing whether it will address what actually bothered you.

Strong Feedback: "The structure is good, but the tone is too formal. Make it more conversational. Move the cost data to the top because that's what my audience cares about most. Cut the methodology section; we don't need it for this executive audience."

Now AI knows exactly what to change: tone (more conversational), sequencing (cost data first), and scope (drop methodology). **Specific feedback produces specific improvement**.

After AI delivers a draft, ask two questions:

1. What's working?

2. What needs adjustment?

Then provide feedback: *"[What's working] is good. But [what needs adjustment]. Specifically, [concrete change 1], [concrete change 2], and [concrete change 3]."*

Why Most People Stop Too Early

Many people approach AI like Google: type a query, get an answer, move on. When the answer isn't perfect, they conclude the tool failed. But AI isn't a search engine returning facts. It's closer to a capable colleague producing work that benefits from feedback.

Research on delegation confirms this: ongoing communication is the secret to successful delegation. High-performing teams don't expect perfection on draft one. They build feedback loops into the process. Organizations using formalized feedback loops report up to 30% productivity increases.

If your first AI output is 70% of what you need, that's not failure. That's a strong draft. Provide specific feedback. Get draft two. By draft three, you'll typically have something highly usable.

The professionals getting exceptional results from AI aren't lucky. **They're finishing the conversation**.

Effective delegation isn't one-and-done. You brief, review, redirect, and refine. You treat outputs as drafts requiring iteration, not verdicts requiring acceptance or rejection. AI responds to the same workflow. That's not a technical skill. It's a collaboration skill. And you already have it.

Know When to Take Over: The Boundaries of Good Delegation

You've learned to set context, specify expectations, share your reasoning, and iterate. But here's the final essential skill: **knowing when to stop delegating and take over yourself.**

The best delegators make clear distinctions about which tasks benefit from delegation and which require their own expertise. AI delegation follows exactly the same principle.

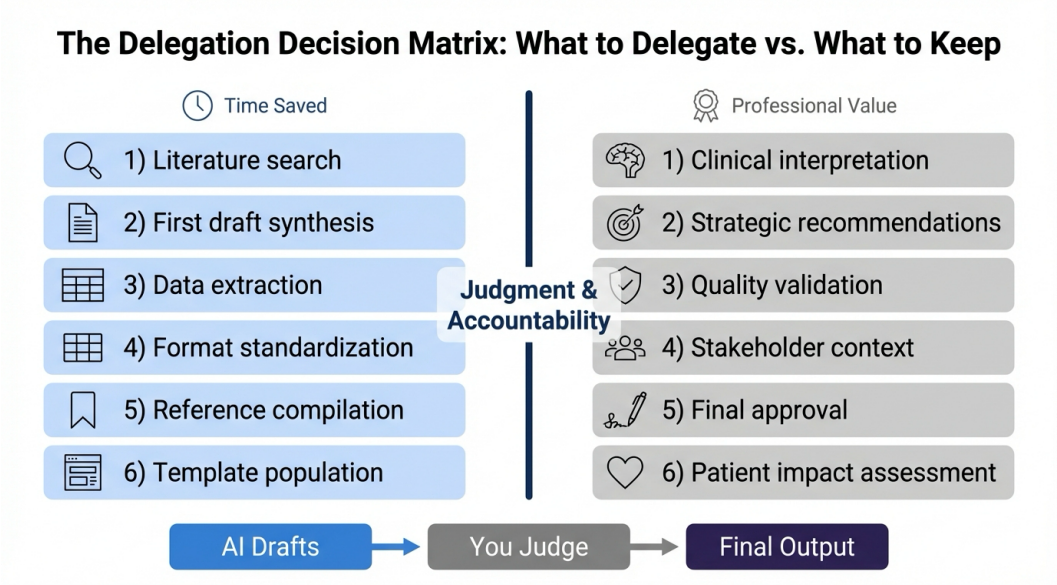

The Delegation Decision Matrix: What to Delegate vs. What to Keep

When you brief a junior analyst, you don't delegate the final interpretation of whether your data supports a coverage recommendation. You delegate the data synthesis and formatting, then apply your judgment and sign off. **AI works the same way**. The goal isn't to automate your thinking. It's to free up more time for the thinking only you can do.

The Delegation Decision Framework

Research on effective delegation identifies clear boundaries: delegate tasks where mistakes are reversible or where someone can achieve roughly 70% of your performance level. For AI, this translates into a simple framework:

Delegate to AI	Keep for Yourself
Structuring information into clear formats	Final judgment on data interpretation
Drafting content for your review	Decisions about which evidence matters most
Synthesizing findings from multiple sources	Sign-off on clinical or economic conclusions
Generating formatting options	Determinations about what's accurate in your field
Compiling references and citations	Final accountability for professional work

This isn't a technical distinction. It's a judgment distinction. Use AI for tasks that free your time. Reserve your expertise for decisions that carry professional weight.

What AI Does Well (and What It Doesn't)

When healthcare organizations implement AI for documentation and workflow automation, the goal is to free clinicians for high-judgment work, not replace that judgment. Apollo Hospitals aims to give clinicians back two to three hours per day by automating documentation tasks, but physicians still make every clinical decision.

The same boundary applies to research work. AI excels at structuring content, drafting summaries, synthesizing information, and formatting. AI should *not* handle interpretation (what does this data mean for your formulary decision?), judgment calls (which studies are most relevant to your context?), or professional sign-off (is this conclusion defensible?).

The line isn't arbitrary. It's the same line you draw when delegating to a colleague: you hand off execution, but you retain judgment.

A Practical Example

You're preparing a value dossier. Here's how to divide the work:

Delegate to AI:

- "Generate five formatting options for presenting cost-effectiveness data: table, bullet summary, infographic structure, narrative, and visual comparison."

- "Draft the methodology section summarizing our economic model structure, data sources, and key assumptions."

- "Compile all references in APA format and verify completeness."

Keep for yourself:

- Review the five formatting options and choose which best serves your payer audience.

- Read the methodology draft and determine if it accurately represents your model and whether assumptions need clarification.

- Verify references are accurate and citations support the claims made.

AI handles structure and formatting. You apply judgment and ensure accuracy. This mirrors how you'd work with a capable junior team member.

Why This Boundary Matters

Your value isn't in formatting slides or compiling references. Your value is in knowing what the data means, which evidence is credible, what your stakeholders care about, and what conclusions are defensible.

AI can't replace that judgment because it lacks your professional context, your organizational knowledge, and your accountability. What AI *can* do is free you from tasks that consume time without requiring your unique expertise.

Healthcare productivity gains from AI are estimated at 5 to 10% within five years, but only when AI handles execution tasks while humans retain oversight and decision-making. The productivity comes from giving professionals back hours per day to focus on high-judgment work.

Research on delegation confirms this: effective delegators focus their time on tasks requiring specific expertise and delegate tasks where others can achieve 70% of their performance. If AI can structure, draft, and format at 70% of your level, that's good enough for a starting point. You refine it with your 30%, and the overall process is faster.

The Real Goal

The point of AI delegation isn't to remove yourself from the work. It's to remove yourself from the parts that don't require your judgment.

When you spend two hours formatting a slide deck, that's two hours you're not spending interpreting data or advising stakeholders. When AI handles the formatting in ten minutes, you have two hours back for the work only you can do.

The goal isn't to automate your thinking. It's to amplify it.

You wouldn't personally compile every reference if you had an assistant who could do it reliably. You'd review their work for accuracy but delegate the compilation. The same logic applies to AI.

Knowing When to Step In

Here's a practical test: **If a mistake would require your professional judgment to catch, don't delegate the final decision to AI.**

- Can AI draft a cost-effectiveness summary? Yes. Should you review it for accuracy? Absolutely.

- Can AI format references? Yes. Should you verify citations match claims? Yes.

- Can AI structure your slide deck? Yes. Should you decide which structure best serves your audience? Yes.

AI accelerates execution. You provide judgment, context, and accountability. That partnership is where the value lives.

Good delegation has always been about knowing what to hand off and what to keep. You already make these decisions with colleagues every day. Now you're applying the same instinct to AI.