

The anti-workslop guide.

How to make sure the AI-assisted work leaving your hands is real work. And what to do, without starting a fight, when someone hands you the other kind.

BY DAVE BROWN

Something strange is happening inside companies that went all-in on AI. Usage is up. Mandates have landed. Everyone is producing more, faster, and it all looks quite good. And the return on the investment is, by most measures, close to nothing: a report out of the MIT Media Lab put 95% of organizations at no measurable return at all.

There are a lot of explanations for that gap. This guide is about one of them, and it's the one you can personally do something about by Thursday.

Researchers at BetterUp Labs and the Stanford Social Media Lab gave it a name: **workslop**. Their definition is precise, and worth sitting with: AI-generated work content that masquerades as good work, but lacks the substance to meaningfully advance the task. It isn't bad work in the obvious sense. Bad work gets caught. Workslop is *polished*. It has headers, structure, an executive summary, a confident tone. It passes the two-second glance. Then whoever receives it has to work out what's missing, what's wrong, and what you actually meant, which was, of course, your job.

You have almost certainly been on the receiving end. You open a document and feel a specific, hard-to-place confusion, followed by irritation. *Wait, what is this?* Then the suspicion that the sender didn't write it so much as commission it. That feeling has a name now. And it has a price tag.

40%

of full-time employees
received workslop in the
past month

1h 56m

average time lost dealing
with a single instance

\$9M

estimated annual cost
per 10,000 employees

32%

don't want to work with
the sender again

Those numbers come from a survey of 1,150 U.S. full-time employees. The last one is the one I'd tape to your monitor. The productivity cost is real but recoverable; the reputational cost isn't. Roughly half of people who received workslop rated the sender as less creative, less capable, and less reliable than they had before. Forty-two percent trusted them less. A third quietly decided they'd rather not work with them again, and never said a word about it.

That's the trap. Workslop feels efficient to send and expensive to receive, and the receiver almost never tells you. You get the time savings immediately and pay the bill six months later, in a meeting you weren't invited to.

AI produces drafts. People produce work. You are not finished until the draft is something you can stand behind in a room.

That single line is the whole guide. Everything below is what it looks like in practice: first as personal discipline, then as self-defense, then as something you can actually ask a team to follow.

WHAT'S IN HERE

- 01 Don't be the source. Three disciplines that keep your own output honest.
- 02 Don't be the victim. Five ways to push back that don't cost you the relationship.
- 03 The slop loop. What happens when two AIs argue through two people.
- 04 If you run a team. Seven changes that actually move the number.
- 05 The drop-in standard. A team norm you can paste into a channel today.

One framing note before we start. None of this is an argument against using AI. The same research that found workslop also found that the heaviest, most enthusiastic AI users get dramatically more out of it than everyone else. The researchers call them *pilots*: high agency, high optimism. Pilots use AI 75% more at work than the passengers do. The difference isn't how much they use it. It's that passengers use AI to *avoid* work and pilots use it to *do* work. Workslop is what passengering looks like from the outside.

Three disciplines, and one test before you hit send.

Almost all self-inflicted workslop traces back to the same three failures: thin context, mistaking length for effort, and handing over the part of the job that was actually yours.

Start with the mindset the rest of this depends on: **your name on it means you vouch for it.** Not the tool's name. Yours. The moment something leaves your hands, you are the author of every claim in it, and "the AI wrote that part" is not a thing you will ever get to say out loud without cost. If you can't defend a line, that line isn't finished.

Here's the honest gut-check that catches most of it, and it takes four seconds: *if I'd had unlimited time, is this the work I would have produced?* If AI let you skip the typing, that's the deal working as intended. If it let you skip the *thinking*, you've made workslop, and the person downstream is about to do that thinking for you.

DISCIPLINE 01

Context is the work. The prompt is where you do it.

The model doesn't know what you know. It doesn't know that this client already rejected the phased option, that the VP asking for this actually wants ammunition for a budget conversation, that the number in the third column is wrong for reasons everyone on the team understands and nobody has written down. None of that is on the internet. All of it is the reason your version of the document should be better than a stranger's.

So here's the test that never fails: **could anyone at any company have typed my prompt?** If yes, you are going to get output that could have gone to anyone at any company, which is precisely the texture people recognize as slop. Generic in, generic out. It's not a model limitation, it's an input problem, and it's yours.

Three practical moves, in rough order of how much they help:

- **Give it the raw material, not a description of the raw material.** Paste the actual meeting notes, the actual thread, the actual numbers. "Summarize our Q3 issues" produces a horoscope. The transcript plus "what did we actually decide, and what's still open?" produces something useful.
- **Write down the three things only you know** before you write the prompt. The constraint, the history, the politics. That's usually the entire delta between slop and work.

- **Name the decision the document has to serve.** Not “write a proposal” but “this needs to get a skeptical CFO to approve one of two options in a fifteen-minute meeting.” The model is surprisingly good at reverse-engineering a document from a decision. It cannot guess the decision.

And then don't stop at the first output. The first response is raw material; the real work is the second and third pass, where you tell it what's wrong, what to cut, what it assumed that isn't true. People who get consistently good results aren't better at prompting. They're just less willing to accept the first thing they're handed.

THE TELL

Read your draft and hunt for any sentence that would be equally true at a competitor. Every one of those is a place where you supplied no context and the model filled the space with plausible filler. Cut them or replace them with something only your team could have written.

DISCIPLINE 02

Length is not effort. It only used to be.

For all of professional history, volume was an honest signal. Nobody wrote nine pages without caring about the subject, because nine pages cost a day. Length meant effort, effort meant seriousness, and readers learned to respect the thick document.

Generative AI severed that link completely, and our instincts haven't caught up. Which is exactly why long AI output *feels* like it's working: you send twelve paragraphs, it reads as thorough, and for about ninety seconds you look like the most diligent person on the thread. Then the reader gets to paragraph five, notices nothing has been decided, and reprices the whole thing. Now you look like someone who wasted their morning.

The discipline is unglamorous: **the finished thing is usually shorter than the draft.** Not tightened. Shorter. Most of what a model generates is connective tissue that exists to make prose flow, and prose flow is not what your reader needs from a status update. Cut without mercy:

- The opening paragraph explaining that the topic is important and evolving rapidly.
- The paragraph summarizing what you're about to say, followed by you saying it.
- The scrupulously balanced “on the other hand” section that surveys both sides and commits to neither. That's not nuance. That's the absence of a position, dressed as maturity.

- The closing recap of everything above it.
- Any sentence that could appear in any deck at any company. (See discipline one.)

Watch for structure inflation too, because it's the sneakier version. Headers, nested bullets, bold lead-ins, a tidy three-part framework: formatting that makes thin thinking look organized. Real structure comes from having a sequence of ideas that depend on each other. Applied structure is just packaging, and experienced readers can smell it in seconds.

Simplest rule I've got: if your summary is longer than the thing it summarizes, start over. And if you catch yourself hoping the length itself will do some persuading: that's the moment to stop, because you already know you don't have the substance.

DISCIPLINE 03

Don't outsource the deciding.

We've been offloading cognition to machines for a long time, and mostly it's been fine. Nobody memorizes phone numbers anymore and civilization held. The difference here is subtle but total: when you Google a fact, you offload *retrieval* and you can verify the answer in a second. When you ask a model to tell you what your team should do next quarter, you're offloading *judgment*, and judgment can't be spot-checked, because the whole point is that there's no correct answer sitting somewhere to compare against.

That's also, not coincidentally, the part you're paid for. AI is genuinely excellent at everything up to the decision, and it cannot make the decision, because the decision requires knowing which stakeholder can absorb bad news this month and which one can't. If you delete your judgment and ship the scaffolding, you've handed over the only part of the work that was actually yours, and the person receiving it now has to supply it.

The fastest diagnostic I know: **can you say what you think in three sentences, right now, with the laptop closed?** If not, you don't have a position. You have a document. Those are extremely different things and only one of them is worth sending.

THE MEETING TEST

Before you send, imagine someone senior reading one line aloud and asking, "Why this?" If your honest internal answer is "because it came out that way," delete it or earn it. Anything you couldn't defend out loud isn't ready to go out in writing. Writing just delays the question.

There's a compounding version of this worth naming. Every judgment call you hand to a model is a rep you don't take. One skipped rep is nothing. A year of skipped reps is a real erosion in the specific muscle that makes you valuable, and the erosion is invisible right up until the moment you need to make a hard call fast and notice you're reaching for a tool instead of an opinion. Use AI to *pressure-test* your thinking, argue with it, find your blind spots. Just don't let it do the thinking and then sign the result.

The four questions, right before you send.

This takes thirty seconds and catches nearly everything. Run it on anything AI helped you produce, including, and especially, the things you're confident about.

QUESTION	WHAT YOU'RE ACTUALLY CHECKING
Ownership	Can I defend every claim, number, and recommendation here without pointing at a tool? If I'd be embarrassed to have written it myself, I can't send it.
Substance	Did AI add length, or did it add value? If I cut this in half, what would actually be lost? If the answer is "nothing," cut it in half.
Context	Does this contain the specifics of <i>our</i> situation (our constraints, our people, our decision), or could it have been sent by anyone, to anyone?
Receiver	What will the person opening this have to do before they can use it? If they have to reconstruct my thinking, I haven't done my job. I've transferred it.

The fourth is the one that matters most and gets skipped most. Workslop is fundamentally a transfer of effort, so the only reliable test is to look at the transfer directly: what is now on their plate that was on mine ten minutes ago?

Five habits that keep you clean.

- **Read every line before it leaves you.** Pasting unread output is the single most common way slop gets made. There is no version of this where you skip the read.
- **Use AI where you can verify, not where you can't.** First drafts, structure, rewrites, counterarguments: excellent. Facts you have no way to check: dangerous, and eventually career-limiting.

- **Add the part AI can't.** The priority call, the recommendation, the thing to ignore, the political read. If that's absent, you sent scaffolding.
- **Treat shorter as a win.** Cutting is the highest-leverage editing you'll do, and it's the one thing the model won't volunteer.
- **Disclose plainly when it matters.** "I drafted this with AI and checked the numbers against the source" costs you nothing and buys you enormous credibility. Concealment is what gets punished, not use.

PART TWO · DON'T BE THE VICTIM

Five ways to push back without making an enemy.

Everyone writes the guide about not sending slop. Almost nobody writes the one you actually need at 4pm on a Tuesday, when it's already in your inbox and the sender outranks you.

Here's why this half matters more than it looks. When people receive workslop, the overwhelmingly common response is to absorb it: quietly rewrite it, quietly do the missing thinking, quietly downgrade their opinion of the sender, and say nothing. That's the worst possible outcome for everyone involved. The sender learns that it worked and does it again. You eat two hours and some resentment. And the actual standard on your team drops by one notch, permanently, because nobody defended it.

The reason people absorb it is simple: naming it feels like an accusation. So the whole trick (and it is the entire trick) is to **make the conversation about the missing thing, never about the tool.** "This feels AI-generated" is a character allegation and it will go badly whether or not you're right. "What are you recommending?" is a normal professional question that happens to be unanswerable by anyone who hasn't done the work.

Judge the work, not the method. It's the same rule that protects you when you're the one using AI.

MOVE 01

Ask for the position, not for more document.

The single highest-return question you can ask, and it works in every direction on the org chart: *"What are you recommending, in your own words?"*

It's genuinely neutral. You're asking someone to summarize their own work, which is the most ordinary request in professional life. But it's unfakeable. A person who thought it through answers instantly and usually with relief. A person who generated it either has to go think, or hands you another document, which tells you everything. Either way the work goes back where it belongs and nobody has been accused of anything.

Do not ask for a shorter version of the document. That's a formatting request and you'll get a shorter document with the same hole in it. Ask for the *position*. And ask for it in a medium that's hostile to volume: Slack, a hallway, a two-minute call.

SAY THIS

"Before I dig in, what are you recommending, and what's the main reason?"

Variants: "What's the decision you need from me here?" · "If you had to pick one, which and why?" · "Give me the three-sentence version."

MOVE 02

Fix it upstream: ask in a shape that's hard to slop.

Most workslop is invited. "Can you put together something on the vendor situation?" is an open invitation to produce volume, because you haven't said what would count as success, so the safest move available to the person you asked is to send a lot and let you sort it out. That's not laziness, it's rational under ambiguity. AI just made the lot-of-stuff option free.

Constraint is the best slop filter ever invented, and it's the only one that operates *before* the work happens. Specify three things when you delegate: the length, the decision it has to serve, and the requirement of a recommendation. All three are hard to fake:

- **A length cap forces prioritization.** "One page, max" means someone has to decide what matters. That decision is the work.
- **A required recommendation forces a position.** "Three options, your pick, and why" can't be produced without a human choosing.
- **A named decision forces relevance.** "This has to get us to a go/no-go by Friday" kills every paragraph that doesn't move toward it.

Two more asks I've found weirdly effective: *"include the two strongest objections to your recommendation"* and *"tell me what you decided to leave out."* Both require somebody to have actually held the material in their head. Neither reads as distrust. They read as rigor, which is exactly what they are.

MOVE 03

Return it once, with one specific gap.

When something lands short, there are three options and most people pick the worst one. You can rewrite it silently (costs you two hours, teaches nobody, breeds resentment). You can do a full teardown (reads as an indictment, wrecks the relationship, and takes longer than rewriting). Or you can send it back once, with exactly one specific ask.

One ask reads as collaboration. Five reads as a performance review. So pick the single most load-bearing hole (usually the missing recommendation, the unverified numbers, or the absent context about your actual situation) and ask for that alone. Include the deadline when you hand it back, so it's a workflow step rather than a verdict.

Keep the language about the artifact, never the author or the tool. "This section needs X" is a note about a document. "You clearly didn't read this" is a note about a person, and people defend themselves instead of fixing documents.

SAY THIS

"This is a good start. The piece I need is your actual recommendation and the reasoning behind it. Can you add that and send it back by Thursday?"

Variants: "Can you flag which of these numbers you've verified against the source?" · "What's the part that's specific to our situation?" · "Which of these three would you actually do?"

MOVE 04

Stop paying the tax in silence.

The two-hour rework cost is invisible for exactly one reason: receivers eat it privately. Nobody logs it, nobody mentions it, and so the organization genuinely does not know it's happening. Meanwhile a third of people who receive workslop *are* telling someone, just not the sender. They're telling their manager, or each other. That's the worst of both worlds: the cost stays hidden and the reputational damage still lands.

Make the cost visible, but route it to the standard rather than the person. Procedural language does this well because it's about process and future work, not blame and past work:

SAY THIS

“That took about two hours to reconcile on my end. Can we agree on what these should include so it doesn’t next time?”

Notice what it does: states a fact, proposes a fix, assigns no motive, and ends pointed forward.

When it comes from above you (and 16% of the time it does), the calculus changes but the move doesn’t. You’re not going to send your VP’s document back for revisions. You’re going to ask for the takeaway, which is a completely normal thing for a direct report to want: *“Just so I get this right, what’s the one thing you want me to take away from this?”* That’s deferential on its face and it does the same job. It surfaces whether there’s a position underneath.

If it’s a pattern rather than an incident, take it to norms instead of people. “Could we agree that anything going to the client gets a human pass first?” is a proposal about how the team works. It costs no one their dignity, and it’s much easier to say yes to than an apology.

MOVE 05

Never answer slop with slop.

This is the one people fail, and it’s the one that does structural damage. You receive four pages you suspect were generated. You’re annoyed, you’re busy, and there’s an extremely available response: paste it into a model, ask for a rebuttal, send four pages back. It feels like defending yourself. It feels, briefly, like winning.

What you’ve actually done is double the volume and remove the last human from the conversation. Now there are two “positions” bouncing back and forth, neither of which anyone holds, and the thread can never resolve, because resolution requires two people converging, and there aren’t two people in it anymore.

Break it by changing the medium and the direction. Get on a call. Compress instead of expanding. Put a decision in writing in plain language, in your own words, with your name on it. It feels like unilateral disarmament. It isn’t. It’s the only move on the board that can actually end the exchange, and the person who makes it is visibly the adult.

ONE THING NOT TO DO

Don’t become the AI police. Don’t run detectors, don’t hunt for tells, don’t open with “be honest, did you write this?” Detection is unreliable, the accusation is unrecoverable if you’re wrong, and the whole framing is a trap: it makes AI use the offense, when the offense is unfinished work.

Plenty of excellent work is AI-assisted. Plenty of terrible work is entirely human. Hold the standard, stay out of the forensics.

Two AIs arguing through two people who stopped reading.

It has the exact shape of rigorous debate, which is why it can run for weeks before anyone notices nothing is being decided.

The fastest way a team drowns.

The pattern: Person A generates a long proposal. Person B, rather than reading it and thinking, feeds it to a model and generates a long rebuttal. A responds with another generation. Round three arrives. The thread is now enormous, the tone is professional, and something that looks unmistakably like rigorous disagreement is taking place.

Except neither position was ever held by a human being. It's two text generators taking turns, with people acting as message routers. And because it has the *form* of debate (claims, counterclaims, evidence, structure), everyone involved feels productive. That's what makes it so corrosive. Ordinary conflict at work resolves because two humans eventually get tired and converge. This can't converge. There's nothing in it to converge on.

Four ways to break it:

- **A human holds the pen.** Every artifact in a decision loop needs a named owner who states their actual position in their own words. AI can inform a position. It cannot be one. If you can't say what you think in three sentences without the tool, you don't have a position yet.
- **Compress, don't expand.** The diagnostic is simple and reliable: each round is producing more text and less clarity. Real progress makes a document *shorter*. The moment round two is longer than round one, stop generating and start deciding.
- **No relays.** If you're forwarding AI output, you add your take: what you agree with, what to cut, what the call is. Forwarding raw output with "thoughts?" is handing someone your homework with a smile.
- **Move to a decision, not another draft.** If two rounds haven't converged, a third won't. That's not an information problem, it's a decision problem. Get on a call, make the call, write it down in plain language.

Seven changes that actually move the number.

Individual discipline has a ceiling. Slop propagates because of what leadership rewards, measures, and mandates, and those are all things you can change this quarter.

01

Kill the blanket mandate.

"Use AI everywhere, all the time" is the single instruction most likely to manufacture slop. It signals that *volume of usage* is the goal, so people paste AI into everything, including the tasks it's worst at. Worse, an indiscriminate mandate models exactly the behavior you're trying to prevent: passing the buck instead of doing the thinking. Replace it with specifics: here's where this genuinely helps us, here's where it doesn't, here's how we check it. Discernment is the whole game, and it has to be demonstrated from the top or it doesn't exist.

02

Publish a quality bar, not an adoption target.

If the only number leadership talks about is "percentage of employees using AI," you are measuring activity and you will get activity, including a great deal of activity that costs you money. Publish the quality standard instead: what good AI-assisted work looks like here, and what disqualifies something from being sent at all.

03

Judge the work, not the tool.

Adopt one line as policy: **human-AI work is held to the exact same standard as human-only work.** This quietly removes the largest available excuse, because "the AI wrote it" stops being a defense the instant the standard is identical either way. It also protects the good-faith users, who currently face a real and well-documented competence penalty just for admitting they used a tool. Same bar, no discount, no stigma.

04

Build pilots, not passengers.

The research splits users cleanly. *Pilots* (high agency, high optimism) use AI to sharpen their own output and hit specific goals. *Passengers* use it to avoid work. Passengers are the slop factory, and you don't fix that with a rule. You fix it by making purposeful use visible: have someone walk the team through how they used AI on a real deliverable, *including how they checked it*. "Here's what I generated, here's what I threw out, here's what I verified" needs to become a normal and respected thing to say out loud.

05

Make pushback safe and boring.

Right now, questioning a colleague's deliverable feels like an accusation, so people absorb the rework instead and the standard erodes silently. Give the team neutral, shared language (the scripts in part two exist for this), and then use it yourself, publicly, on your own team's work. "Can you give me the three-sentence version with your actual recommendation?" should be as unremarkable as asking for a status update.

06

Measure outcomes, not prompts.

Track whether the work got better, faster, or clearer. Not how many prompts were typed, not seat utilization, not tokens. If the dashboard rewards AI activity, you are directly subsidizing slop, and your people will correctly optimize for what you're counting.

07

Name the failure mode out loud.

Give it a word your team actually uses. "Workslop" works fine. Naming it makes it visible, and visible things get avoided: people are markedly less likely to send something they know has a label and a reputation. It also gives receivers a way to describe what happened without describing a person, which is the whole reason the term is useful.

Paste this into a channel and you're done.

Seven lines. This is the “how do we use AI here” document people keep asking for, minus forty pages of policy nobody reads.

—— HOW WE USE AI ON THIS TEAM

01 · AI DRAFTS, YOU FINISH

Anything AI helps produce is a draft until a human has read every line, verified the claims, and can defend it. Your name on it means you vouch for it.

02 · SAME BAR, EVERY TIME

We judge the work, not the method. AI-assisted work is held to the identical standard as work done without it. “The AI wrote it” is not a defense, and using AI is not a mark against you.

03 · NO RAW PASTE

Don't send output you haven't read and shaped. If you wouldn't have written it yourself given the time, don't send it.

04 · ADD THE THINKING

Your job is the judgment: the priorities, the recommendation, the context only you have. Don't delete that and ship the scaffolding.

05 · SHORTER IS A WIN

If it's longer than it needs to be, it isn't finished. Length is not effort.

06 · OWN YOUR POSITION

In any decision, state what you think, in your own words. AI can inform your view; it can't be your view. No AI-arguing-with-AI loops.

07 · ASKING FOR THE REAL VERSION IS NORMAL

Asking a teammate for “the short version with your actual recommendation” is a welcome request, not an insult. Nobody has to absorb rework in silence here.

THE SCRIPTS

Naming it without starting a fight.

None of these accuse anyone of anything. They just quietly re-establish the standard: somebody has to have done the thinking, and it isn't going to be the person receiving the message.

WHEN	WHAT TO SAY
It's long and you can't find the point	"Before I dig in, what are you recommending, and what's the main reason?"
There's no decision in it	"What's the decision you need from me here?"
It reads generic	"This is solid background. What's the part that's specific to our situation?"
The numbers feel unverified	"Can you flag which of these you've checked against the source?"
You're about to rewrite it yourself	"I think this needs another pass before it goes out. Want to take it, or should I?"
It's the second time this month	"That took about two hours to reconcile. Can we agree on what these should include so it doesn't next time?"
It came from your manager	"Just so I get this right, what's the one thing you want me to take away?"
You're delegating and want to prevent it	"One page, three options, your pick, and the two strongest objections to it."

THE WHOLE THING, COMPRESSED

Workslop is effort transfer disguised as productivity.

You beat it personally by refusing to send anything you can't defend out loud, which means bringing real context to the prompt, cutting instead of padding, and keeping the judgment for yourself. You beat it interpersonally by asking for the position instead of the document, constraining the ask before the work starts, returning things once with a single specific gap, making the rework cost visible in neutral terms, and never answering a generated argument with another one. You beat it organizationally by holding AI-assisted work to precisely the same bar as everything else, killing the blanket mandate, and measuring whether the work got better rather than whether the tool got used.

The people getting real value out of these tools aren't using them less than everyone else. They're using them as a collaborator on work they're still doing, not as a way to hand their colleagues their job.

ONE MORE THING

If you want a second set of eyes.

Reading a guide and changing how a team actually works are different exercises.

I put this together because the pattern is real and I keep watching it play out, in my own inbox and in teams that are genuinely trying to do this well. The individual discipline is the easy part. Rolling a standard out across a team without it landing as an accusation, or as one more policy nobody follows, is considerably harder, and it's difficult to get right entirely on your own.

So I'm doing a small number of informal reviews for people working on exactly this. Send me what you've got (a draft team standard, a mandate you've been handed and aren't sure about, or just a description of how AI work is currently landing on your team) and I'll give you honest, specific feedback on what I'd change and what I'd leave alone.

This is early and informal. No pressure either way; the guide above is yours to run with regardless.

GET IN TOUCH

Email hi@davecto.com with the subject line **"Workslop Review"** and a couple of sentences about where it's showing up for you.

More guides like this one, for people trying to use AI without embarrassing themselves.

Weekly, plain-language breakdowns on Instagram.

[@davecto](#)

Figures cited throughout are from "AI-Generated 'Workslop' Is Destroying Productivity" (Harvard Business Review, 2025), reporting research by BetterUp Labs and the Stanford Social Media Lab across 1,150 U.S. full-time employees, and from the MIT Media Lab's report on enterprise AI returns.