



Founder's Manifesto

Keep the Human in the Work

A source-first manifesto for better AI-assisted writing

Daniel Stouffer, Founder

October 5, 2026

Boutique Agencies & Ghostwriters

Contents

1	It reads better. So why does it feel wrong?	p. 3
2	You are not resisting progress. You are protecting the work.	p. 4
3	The real problem with AI slop isn't always how it sounds.	p. 5
4	Fidelity Before Fluency	p. 6
5	Five commitments for keeping the human in the work	p. 7
6	Try the Source Check	p. 9
7	Why I'm building Lucidry	p. 11
8	Let's make this how we work.	p. 12



Use AI. Keep your judgment. Protect what matters.



It reads better. So why does it feel wrong?

You know the moment. You've asked AI to clean up a draft, and the result reads easier. The opening moves faster, the repetition is gone, and a paragraph that needed work suddenly sounds ready to publish.

Then you go back to the source. Something's different, and it isn't just the wording.

Here's an illustrative example, built to make the change visible. It's not a customer result or a product test, just a clean way to show what happens.

The source: *"In our early tests, some clients completed the process in about half the time."*

The rewrite: *"Our process cuts completion time by 50%."*

The second sentence is shorter and more confident. It's also making a different promise. Early testing has quietly become an established result, a limited group has become an implied general outcome, and an approximation now reads like a precise measurement.

Nobody asked for those changes. Someone just asked for better writing.

For an agency editor, that's not an abstract distinction. It's the difference between representing a client accurately and putting a stronger claim into their mouth. For the person reading the finished article, it's the difference between understanding what the evidence suggests and being invited to believe more than it actually supports.

This is the part of AI-assisted writing I want us to talk about. Not just the awkward phrases we can spot right away, but the changes that arrive dressed up as improvements.

If your work involves carrying another person's ideas into the world, you already have a reason to care. The same holds true when those ideas are your own. A draft can get more fluent while drifting farther from what you actually meant to say.

The sentence improved. The responsibility didn't disappear.



You are not resisting progress. You are protecting the work.

There's nothing unreasonable about wanting help. Maybe you're running an agency, turning interviews into executive articles, or working through a manuscript on top of everything else your business demands. A tool that helps organize the material or untangle a rough paragraph can be genuinely welcome.

You don't need a lecture about doing everything the hard way. What you need is a way to accept useful assistance without giving up control of the result.

The tension starts the moment a finished-looking draft leaves you with unfinished questions:

- Did that phrase actually come from the interview?
- Does the link still support the claim?
- Would the client recognize their own position in this version, or only recognize the subject?

Checking those questions is work. When it turns into another round of comparison, repair, and clarification, that time belongs in the project just as surely as the drafting time did.

I think of that work as the **Hidden Revision Tax**: the cleanup, verification, and reconstruction that happens after the apparent finish line. You don't need an industry benchmark to see it. Just look at your last completed assignment and count the work between receiving the first draft and approving the version you could actually deliver.

What needed attention? Who provided it? Which changes improved the work, and which simply repaired something that never should have changed in the first place?

My own work brings together writing, publishing, and building digital products. Those responsibilities shape the question I bring to AI: does the finished result help someone do useful work, and can they understand what they're relying on?

The editor who asks one more question is contributing to that usefulness. So is the ghostwriter who restores an uncomfortable but important qualification, or the author who rejects a smoother sentence because it no longer says what they meant.

That gives us something in common even when our jobs look different. We're all people responsible for words that other people trust. An agency founder might be thinking about a client relationship, while an independent author is thinking about a reader they'll never meet.

Both have a reason to ask what survived the edit. Care isn't the opposite of progress. It's part of deciding what progress should even mean.



The real problem with AI slop isn't always how it sounds.

"AI slop" is a useful phrase for a familiar frustration, but it can turn into a lazy diagnosis. We point at inflated language, repetitive transitions, and paragraphs that circle an idea without ever landing on it, then start treating a pile of writing habits as the whole problem.

For this conversation, I want to look at three failures instead:

- **Language that contributes too little.** A practical observation disappears beneath grand language, or a specific explanation turns into advice vague enough to fit almost anyone. Trimming unnecessary words helps, but it can't hand you an insight the source never contained in the first place.
- **Meaning that changes without permission.** Picture an interviewee saying, "We made the process faster, but people felt less heard." A rewrite turns that into, "We streamlined the customer experience." I want to flag this as an illustrative example, not a client story. The new sentence removes exactly the tension that made the original observation worth discussing. The words have gotten more polished while the idea underneath has lost its edge. A ghostwriter could spend considerable effort restoring that edge, because it belonged to the client's thinking rather than to a stylistic preference.
- **Responsibility that gets hard to locate.** One system rewrites the paragraph, another polishes it, and someone pastes the result into a document. If a consequential change slips through, who can explain where it happened, or whether anyone intended it? That's what I mean by a

Black-Box Rewrite. It offers the convenience of new text without an equally useful account of the transformation. The reviewer inherits the questions the workflow left unanswered.

None of this is a reason to shame people for using assistance. Dictation, translation, editing support, and generative tools can all serve different needs.

The standard should concern the work and the responsibilities around it, not whether someone used the same tools we would choose. Nor should a familiar phrase or a punctuation choice become a verdict on who wrote something.

Those shortcuts pull attention away from the question we can actually examine: does this particular article serve its reader and accurately represent its source?



Fidelity Before Fluency

Making a sentence sound human isn't the same as keeping a person's meaning intact. Fix the first problem and ignore the second, and you're still left with writing you need to question. We don't have to make that bargain.

Here's the belief behind this manifesto: we can use AI to make writing better without handing over the judgment, meaning, and responsibility that make it worth reading. That's not a promise that a tool will never make a mistake. It's a different expectation for the process.

We can decide what matters before asking for changes, check what happened afterward, and keep the important decisions with the people who understand the work.

I call this principle **Fidelity Before Fluency**. Fluency is about how well the writing reads. Fidelity is about whether the meaning, evidence, and intent we needed to preserve actually survived.

This isn't an argument for keeping every original sentence. Good editing can reorganize a section, challenge an argument, or cut something the writer has grown attached to. The real question is whether those changes are understood and intentional, not quietly accepted just because the prose sounds smoother.

Research gives us one reason to take both the benefits and the tradeoffs seriously. In a 2024 short-story experiment, Anil Doshi and Oliver Hauser found that access to AI-generated ideas improved evaluations of individual stories while making the resulting stories more similar to one another.^{[1](#)}

That study doesn't predict what will happen to your client article. It does show that improvement and increasing similarity can show up at the same time.

Our response doesn't have to be rejection. We can ask more of how assistance gets designed and used.

Think about refinement. Its purpose is to make valuable material more usable, not to replace everything distinctive with something easier to process. In writing, that valuable material includes the evidence, the perspective, and sometimes the unresolved question the author is brave enough to leave open.

You can see the difference in an ordinary handoff. Instead of receiving a clean document and an instruction to approve it, the editor receives the source, the revised document, and an explanation of the changes that deserve attention.

The review begins with context rather than guesswork. The point isn't to slow every task down. It's to stop treating unseen change as an acceptable cost of moving faster.



Five commitments for keeping the human in the work

A shared belief becomes useful when it changes a decision. These are five commitments I propose for people who want the benefits of AI assistance and still take responsibility for what reaches the reader. They're working standards, not a certification and not a claim to have invented editorial care.

1. Start with something worth saying

Before refining a draft, name the reader, the purpose, and the useful idea. What should someone understand, decide, or be able to do after reading it? An article that can't answer that question needs more thinking, not another polishing pass.

For a ghostwriter, that might mean going back to the interview for the example behind a broad claim. For an agency, it might mean asking whether the article serves the client or just fills a slot in the calendar. Refinement can't substitute for having something to contribute.

2. Protect the substance

Identify what must survive:

- The claim, and its limits
- The supporting evidence
- The distinctive perspective
- The relationships between them

A quotation isn't interchangeable with a paraphrase. A cautious conclusion shouldn't become a promise just because promises read more cleanly.

Protecting the source doesn't mean declaring it correct. The source may contain an outdated figure, a mistaken attribution, or an unsupported assertion. Investigate those problems and correct them explicitly instead of letting them slide. Carrying an error forward faithfully isn't integrity, and fixing it silently robs the next reviewer of information they may need.

3. Make important changes visible

Not every repaired comma needs a meeting called about it. A changed percentage, a removed qualification, or a restructured argument deserves a different level of attention than a typo fix.

The reviewer should be able to tell routine cleanup apart from edits that call for a real decision. Keep the original accessible, and record the consequential changes in a form the next person can actually use. That might be tracked changes, a side-by-side comparison, or a short editorial note.

The point isn't to generate paperwork. It's to keep someone from having to reconstruct the same decisions from scratch later.

4. Keep judgment accountable

Agree ahead of time who can approve the result and what that approval actually means. A tool can flag a possible inconsistency, but it can't take personal responsibility for how a client gets represented or what a reader is encouraged to believe.

Accountability also means giving the reviewer real context and real time. Sticking a human approval button at the end of an opaque process isn't the same as supporting human judgment. The person signing off needs an honest way to understand what they're approving.

5. Respect the people around the document

Work within the permissions you've actually been given:

- Honor client expectations about AI use
- Protect confidential material
- Credit contributions properly
- Remember the reader who will encounter the finished work without ever seeing how it was made

Respect belongs inside the team too. Welcome the colleague who catches a subtle problem instead of treating their question as a delay. Someone willing to say "I don't think this is what they meant" is protecting the work everyone else has already invested in.

The tools can change. Our responsibility to the reader doesn't.



Try the Source Check

Try it on one passage before you change your whole workflow. You don't need new software to start.

Pick one short passage from a work you're already handling, something you're authorized to use. Keep confidential client information out of any tool or conversation that hasn't been cleared for it.

I call this the **Source Check**. Its job isn't to catch every possible error. It's to make the relationship between the original and the revision easier to see.

1. **Write down what the passage has to keep.** Note its main claim, any limits on that claim, the evidence behind it, and whatever's distinctive about the voice or viewpoint. That list becomes your reference point. Have it in hand before a smoother-sounding version starts quietly reshaping your judgment.

2. **Request a bounded edit.** Use whatever tool or process you already use. Ask for the specific improvement the passage actually needs, maybe clearer sentence structure, maybe less repetition. Keep that request separate from any permission to introduce new claims or shift the argument.
3. **Put the original and the revision side by side.** Read for meaning before you judge the polish.
4. **Leave yourself three short notes.**

What improved? Name a specific gain: a clearer explanation, a tighter paragraph, a better sequence of ideas. Give the tool credit where it earned it. This exercise isn't a hunt for reasons to reject every change.

What changed materially? Check names, numbers, quotations, citations, qualifications, and how they connect. Ask whether the same position is being stated with the same degree of certainty. Notice if an unusual but meaningful phrase got smoothed into something generic.

What still needs a decision? Accept what serves the work, restore what got lost, and flag what's still uncertain. An open question should stay visible, not vanish inside a paragraph that now reads as finished.

Use the opening example as a practice case. A faithful revision might read, "In early tests, some clients finished the process in roughly half the time." The wording's cleaner, but the early-stage finding, the limited group, and the approximation are all still there. Whether the finding itself holds up is a separate question, one that still needs evidence.

Then talk through the exercise with a colleague, using an authorized or sanitized example. Ask them, "What did the rewrite change that we might have missed without the source beside it?" Let them push back on your read, including any change you dismissed too fast.

One passage is enough to start a real conversation. Run the practice across actual work and it'll help your team figure out what deserves protection and what kinds of changes actually need an explanation. It's a way to develop shared judgment, not another score to display.



Why I'm building Lucidry

I'm Daniel Stouffer, an author, publisher, and product builder. Lucidry brings those parts of my work together around one question I care about: how do we make useful assistance more dependable without making the people behind the work less visible?

As an author, I want clearer writing without losing the thought that made me start. As a publisher, I care about the reader's experience and the credibility of what reaches them. As a builder, I want a process that makes important decisions visible instead of leaving the next person to guess.

That's the purpose I'm building Lucidry around. I call it an agent-native content refinery, but the plain-English idea matters more than the category: improve the work while protecting what matters in the source.

I learned that lesson the hard way, while I was actually building this thing. I ran a whitepaper I'd written through an early version of the refine engine, just to see what it would do with something real.

It came back clean. Professional. Grammatically flawless. And it wasn't mine anymore.

My contractions had been straightened out. Plain talk I'd chosen on purpose got upgraded into business prose nobody actually talks in. And a two-word sentence I liked, one I'd written because it hit hard, "Margin. Burned.", got quietly deleted. No flag, no note, just gone.

The quality report scored that rewrite a 98. By the numbers, it was a success. By every other measure that mattered to me, it was a loss.

I've spent more than 25 years building digital products, and before that, years working as a career firefighter. Different jobs, same lesson: polished isn't the same as proven. A clean result can still be the wrong result, and a high score can still miss the point entirely.

That's the design goal now. A controlled passage from source material through refinement and validation to a usable result. The original stays available, important changes stay inspectable, and uncertainty reaches the person responsible for deciding what happens next.

An application or authorized agent should have to respect this boundary, not route around it because the output looked fine.

I'm not offering a promise of perfect preservation, and I'm not offering a way to disguise authorship. I want Lucidry to earn your confidence through visible behavior, including an honest account of where a result still needs attention.

Lucidry is how I'm putting this philosophy to work. The philosophy should hold up whether you use Lucidry or not, and it should stay open to criticism from people who understand editorial work better than I do, in contexts different from mine.

A serious objection can improve the product. A difficult source can expose a weakness. Both are worth more to me than agreement that never changes how the work actually gets done.



Let's make this how we work.

Go back to the sentence we started with. The goal was never to preserve awkward wording forever. It was to make the writing clearer while leaving someone able to explain and stand behind what it says.

That's a modest standard, but it has real consequences for how we work together. It asks us to make room for a careful question, keep the original close enough to consult, and treat a surfaced uncertainty as useful information rather than an embarrassment.

We don't need identical tools or identical opinions about AI to practice this. We need a willingness to compare, question, explain, and improve. You can show up as an enthusiastic builder, a cautious editor, or someone still deciding where assistance belongs in your work.

This is the conversation I want to develop through Refinery Notes. I want to look at real passages, share practical methods, and talk honestly about what I'm learning while building Lucidry. I also want to invite other practitioners to change my mind when the evidence calls for it.

Join the source-first conversation. Start by giving the Source Check to someone whose judgment you trust, then think about what you'd add to the five commitments. Which one gets hard to keep under deadline pressure? What would make it easier to actually practice, not just nod along with?

Bring an example where assistance worked beautifully, too. A worthwhile conversation should help us recognize good decisions, not just collect failures. And when someone disagrees with an edit, the job is to understand the difference in purpose or interpretation, not to figure out who's on the right

side of an argument.

You shouldn't have to buy anything, wear a badge, or hand over a client document just to take part. A good question, a permitted example, or a correction offered with care, that's a contribution.

The kind of writing I want more of is writing that helps a reader understand something and earns the confidence it asks them to place in it. That takes more than generating words, and it gives us something worthwhile to build together.

Use AI. Keep your judgment. Protect what matters.

[Explore Lucidry](#)

Research note

1. Anil R. Doshi and Oliver P. Hauser. "Generative AI enhances individual creativity but reduces the collective diversity of novel content." *Science Advances* 10, no. 28 (2024): eadn5290. DOI: 10.1126/sciadv.adn5290. The cited experiment concerns short-story ideation, not Lucidry performance or professional editing outcomes.

About the Author

Daniel Stouffer, Founder

Refined through the Lucidry pipeline, with editorial review & human-in-the-loop guidance by Daniel Stouffer.

Read the origin story: lucidry.io/origin-story.

Try It Free

lucidry.io/try-free

This manifesto was built to be passed along. Share it at lucidry.io/help/resources/agencies-ghostwriters/keep-the-human-in-the-work Share it at lucidry.io/help/resources/agencies-ghostwriters/keep-the-human-in-the-work

Read the web edition: lucidry.io/help/resources/agencies-ghostwriters/keep-the-human-in-the-work