

[Subscribe](#)[Sign In](#)

Generative AI

AI-Generated “Workslop” Is Destroying Productivity

by Kate Niederhoffer, Gabriella Rosen Kellerman, Angela Lee, Alex Liebscher, Kristina Rapuano and Jeffrey T. Hancock

September 22, 2025, Updated September 25, 2025



HBR Staff/AI



Post



Post



Share



Save



Buy Copies



Print

Summary. Despite a surge in generative AI use across workplaces, most companies are seeing little measurable ROI. One possible reason is because AI tools are being used to produce “workslop”—content that appears polished but lacks real substance, offloading cognitive labor onto coworkers. Research from BetterUp Labs and Stanford found that 41% of workers have encountered such AI-generated output, costing nearly two hours of rework per instance and creating downstream productivity, trust, and collaboration issues. Leaders need to consider how they may be encouraging indiscriminate organizational mandates and offering too little guidance on quality standards. To counteract workslop, leaders should model purposeful AI use, establish clear norms, and encourage a “pilot mindset” that combines high agency with optimism—promoting AI as a collaborative tool, not a shortcut. [close](#)

A confusing contradiction is unfolding in companies embracing generative AI tools: while workers are largely following mandates to embrace the technology, few are seeing it create real value. Consider, for instance, that the number of companies with fully AI-led processes [nearly doubled](#) last year, while AI use has likewise [doubled](#) at work since 2023. Yet a [recent report](#) from the MIT Media Lab found that 95% of organizations see no measurable return on their investment in these

technologies. So much activity, so much enthusiasm, so little return. Why?

In collaboration with Stanford Social Media Lab, our research team at BetterUp Labs has identified one possible reason: Employees are using AI tools to create low-effort, passable looking work that ends up creating more work for their coworkers. On social media, which is increasingly clogged with low-quality AI-generated posts, this content is often referred to as “AI slop.” In the context of work, we refer to this phenomenon as “[workslop](#).” We define workslop as *AI generated work content that masquerades as good work, but lacks the substance to meaningfully advance a given task*.

Here’s how this happens. As AI tools become more accessible, workers are increasingly able to quickly produce polished output: well-formatted slides, long, structured reports, seemingly articulate summaries of academic papers by non-experts, and usable code. But while some employees are using this ability to polish good work, others use it to create content that is actually unhelpful, incomplete, or missing crucial context about the project at hand. The insidious effect of workslop is that it shifts the burden of the work downstream, requiring the receiver to interpret, correct, or redo the work. In other words, it transfers the effort from creator to receiver.

If you have ever experienced this, you might recall the feeling of confusion after opening such a document, followed by frustration—*Wait, what is this exactly?*—before you begin to wonder if the sender simply used AI to generate large blocks of text instead of thinking it through. If this sounds familiar, you have been workslopped.

According to [our recent, ongoing survey](#) (which you can take), this is a significant problem. Of 1,150 U.S.-based full-time employees across industries, 40% report having received workslop in the last month. Employees who have encountered workslop estimate that an average of 15.4% of the content they receive at work qualifies. The phenomenon occurs mostly between peers (40%), but workslop is also sent to managers by direct reports (18%). Sixteen percent of the time workslop flows down the ladder, from managers to their teams, or even from higher up than that. Workslop occurs across industries, but we found that professional services and technology are disproportionately impacted.

Here’s what leaders need to know about workslop—and how they can stop it from gumming up the works at their company.

The Workslop Tax

Cognitive offloading to machines is not a novel concept, nor are anxieties about technology hijacking cognitive capacity. In 2006, for instance, the technology journalist Nicolas Carr published a provocative essay in *The Atlantic* that asked “[Is Google Making Us Stupid?](#)” The prevailing mental model for cognitive offloading—going all the way back to Socrates’ concerns about the alphabet—is that we jettison hard mental work to technologies like Google because it’s easier to, for example, search for something online than to remember it.

Unlike this mental outsourcing to a machine, however, workslop uniquely uses machines to offload cognitive work to another human being. When coworkers receive workslop, they are often required to take on the burden of decoding the content, inferring missed or false context. A cascade of effortful and complex decision-making processes may follow, including rework and uncomfortable exchanges with colleagues.

Consider a few examples.

When asked about their experience with workslop, one individual contributor in finance described the impact of receiving work that was AI-generated: “It created a situation where I had to decide whether I would rewrite it myself, make him rewrite it, or just call it good enough. It is furthering the agenda of creating a mentally lazy, slow-thinking society that will become wholly dependant [sic] upon outside forces.”

In another case, a frontline manager in the tech sector described their reaction: “It was just a little confusing to understand what was actually going on in the email and what he actually meant to say. It probably took an hour or two of time just to congregate [sic] everybody and repeat the information in a clear and concise way.”

A director in retail said: “I had to waste more time following up on the information and checking it with my own research. I then had to waste even more time setting up meetings with other supervisors to address the issue. Then I continued to waste my own time having to redo the work myself.”

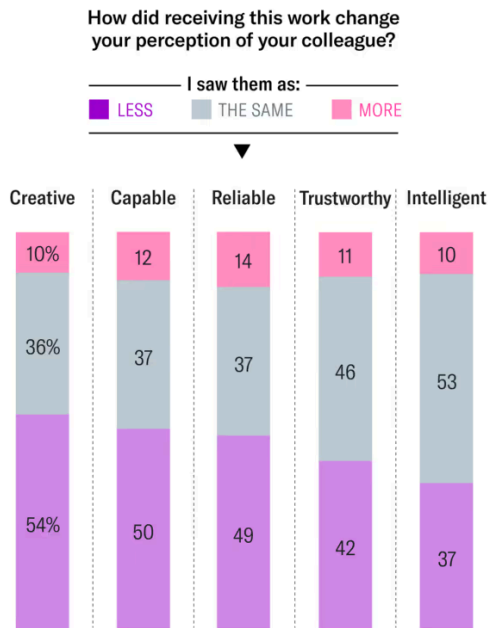
Each incidence of workslop carries real costs for companies. Employees reported spending an average of one hour and 56 minutes dealing with each instance of workslop. Based on participants’ estimates of time spent, as well as on their self-reported salary, we find that these workslop incidents carry an invisible tax of \$186 per month. For an organization of 10,000 workers, given the estimated prevalence of workslop (41%), this yields over \$9 million per year in lost productivity.

Respondents also reported social and emotional costs of workslop, including the problem of navigating how to diplomatically respond to receiving it, particularly in hierarchical relationships. When we asked participants in our study how it feels to receive workslop, 53% report being annoyed, 38% confused, and 22% offended.

The most alarming cost may be interpersonal. Low effort, unhelpful AI generated work is having a significant impact on collaboration at work. Approximately half of the people we surveyed viewed colleagues who sent workslop as less creative, capable, and reliable than they did before receiving the output. Forty-two percent saw them as less trustworthy, and 37% saw that colleague as less intelligent. This may well echo [recent research](#) on the competence penalty for AI use at work, where engineers who allegedly used AI to write a code snippet were perceived as less competent than those who didn't (and female engineers were disproportionately penalized).

Workslop Makes Colleagues Think Less of Each Other

Respondents report that receiving low-effort, AI-generated work lowers their opinion of colleagues across five traits.



Source: Kate Niederhoffer et al.



[See more HBR charts in Data & Visuals >](#)

What's more, 34% of people who receive workslop are notifying teammates or managers of these incidents, potentially eroding trust between sender and receiver. One third of people (32%) who have received workslop report being less likely to want to work with the sender again in the future.

Over time, this interpersonal workslop tax threatens to erode critical elements of collaboration that are essential for successful workplace AI adoption efforts and change management.

What Leaders Can Do

In some ways, this is not a new story. There has always been sloppy work. We are prone to procrastination, to shortcuts, to leaning into busywork instead of careful thinking when we are tired. Gen AI gives us a new technology with which to lean into the same old bad habits—but now with the added cost of creating more work for our colleagues and undermining collaboration, at scale.

So, how can organizations avoid this outcome? How can they instead optimize for employee usage of AI that delivers measurable return on their significant investment in this transformative technology? Here we offer a few key principles from our [original research](#) and experience helping Fortune 500 companies with successful AI adoption across their workforce:

Indiscriminate imperatives yield indiscriminate usage.

When organizational leaders advocate for AI everywhere all the time, they model a lack of discernment in how to apply the technology. It's easy to see how this translates into employees thoughtlessly copying and pasting AI responses into documents, even when AI isn't suited to the job at hand. Gen AI is not appropriate for all tasks, nor can it read minds. To be sure, AI can positively transform some aspects of work, but it still requires thoughtful guidance and feedback from workers in order to produce useful outputs on complex or ambiguous work.

What's more, indiscriminate mandates model the behavior of passing the buck. While employees should be invested with autonomy around AI, the organization should be working through its own careful policies and recommendations around best practices, top tools, and norms. If AI is everyone's job, it is also—and foremost—the job of organizational leaders to develop guidance for employees to help them use this new technology in ways that best align to the organization's [strategy](#), values, and vision.

Mindsets matter.

Our labs have been tracking predictors of gen AI adoption [across the workforce since 2023](#) and we found that workers with a combination of high agency and high optimism are much more likely to adopt gen AI than those with low agency and low optimism. We call these workers “pilots,” as opposed to “passengers.” Pilots [use gen AI](#) 75% more often at work than passengers, and 95% more often outside of work.

Perhaps even more importantly, though, given these findings on workslop, is [how pilots use gen AI](#). Pilots are much more likely to use AI to enhance their own creativity, for example, than passengers. Passengers, in turn, are much more likely to use AI in order to avoid doing work than pilots. Pilots use AI purposefully to achieve their goals.

Recommit to collaboration.

So many of the tasks required to work well with AI—giving prompts, offering feedback, describing context—are collaborative. Today’s work requires more and more collaboration, not only with humans but also, now, with AI. The complexity of collaboration has only deepened. Workslop is an excellent example of new collaborative dynamics introduced by AI that can drain productivity rather than enhance it. Our interactions with AI have implications for our colleagues, and leaders need to promote human-AI dynamics that support collaboration.

Seamless collaboration in 2025 must include the ways we incorporate AI work products into our common workflows, in service of shared outcomes, rather than as a vehicle for subversively dodging responsibility. This is a new, critical frontier of organizational citizenship behaviors that will differentiate between companies that maximize the value of AI and those who churn through AI activity without impact.

Workslop may feel effortless to create but exacts a toll on the organization. What a sender perceives as a loophole becomes a hole the recipient needs to dig out of. Leaders will do best to model thoughtful AI use that has purpose and intention. Set clear guardrails for your teams around norms and acceptable use. Frame AI as a collaborative tool, not a shortcut. Embody a pilot mindset, with high agency and optimism, using AI to accelerate specific outcomes with specific usage. And uphold the same standards of excellence for work done by bionic human-AI duos as by humans alone.

More Resources

- [HBR’s Most-Read Gen AI Articles of 2025 So Far](#)
- [Addressing Gen AI’s Quality-Control Problem](#)
- [The Perils of Using AI to Replace Entry-Level Jobs](#)
- [Why Companies with Successful AI Pilots Do Differently](#)
- [Don’t Let an AI Failure Harm Your Brand](#)

KN

Kate Niederhoffer is VP of BetterUp Labs and a social psychologist. Her research explores the role of AI in workplace adoption, human development, and interpersonal communication.