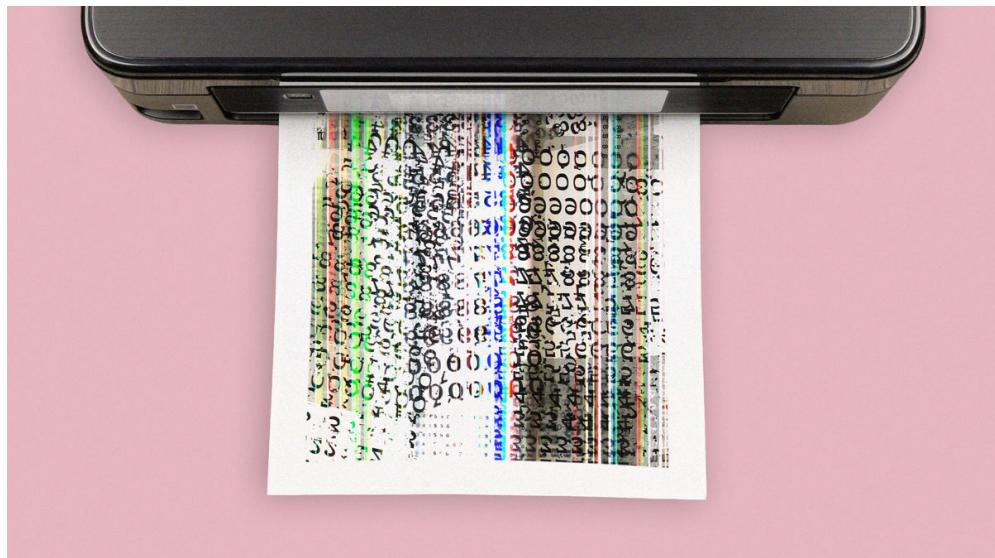

Digital Article / Generative AI

Why People Create AI “Workslop”—and How to Stop It

It isn't a failure of the technology. It's a failure of management.
by *Kate Niederhoffer, Alexi Robichaux, and Jeffrey T. Hancock*

Published on HBR.org / January 16, 2026 / Reprint H090VT



HBR Staff; Printer: Alexisaj/dp3010/Getty Images

As AI tools have proliferated in workplaces and pressure to use them has mounted, employees have had to contend with the scourge of workslop, or low-effort, AI-generated work that looks plausibly polished, but ends up wasting time and effort as it offloads cognitive work onto the recipient. For the person on the receiving end, it can be a confusing and infuriating experience.

When we coined the term workslop in an HBR article this past fall, we explained that it can have a toxic effect on workplace dynamics, breeding mistrust and leading team members to think less of the sender’s intelligence and trustworthiness, among other traits. In our ongoing research, we have heard a number of examples of workslop seeding ill will, eroding trust, and generally having a corrosive effect on workplace morale:

- At one company, employee use of vibe coding created scores of critical bugs in the codebase, causing one engineer to “resent his team so much” that he left with two days’ notice.
- At another company, a qualitative researcher described feeling “gaslit” and “angry” after their manager put their research findings into ChatGPT to generate tables and a discussion section. “The results were incorrect and the discussion section was jargon-y nonsense,” they told us. They were upset “not only because it was unhelpful,” but because their work was put into ChatGPT without their permission. “It felt violating.”
- One employee at a technology company noticed the tone in their performance review was unlike their manager and that the document recycled content from their self-evaluation. The experience made them feel “unvalued and unappreciated,” and they “gave up all hope” they would ever be promoted.

While many of the responses to our article focused on the productivity costs—the time people wasted dealing with each instance of workslop that crossed their desks—what should really worry leaders is the impact workslop can have on human relationships.

It’s tempting to respond to workslop with disdain for those who produce it. Our research points to an uncomfortable answer: The proliferation of workslop is a management failure. Specifically, it is the result of unclear AI mandates and overwhelmed teams. Leaders are issuing

vague directives for employees to start using extremely powerful tools, while many of those employees are overburdened, psychologically depleted, and operating in environments where it doesn't feel safe to admit uncertainty or ask for help.

Workslop is not inevitable. Here's what leaders need to know, and what they can do to prevent it.

The Workslop Pressure Cooker

It often takes time to understand how a new technology will actually be used and the effects it will have on the people who use it. As generative AI tools gain footholds at work, we're starting to see how they change the workplace. The three of us have watched different versions of this process play out in different contexts and with different technologies.

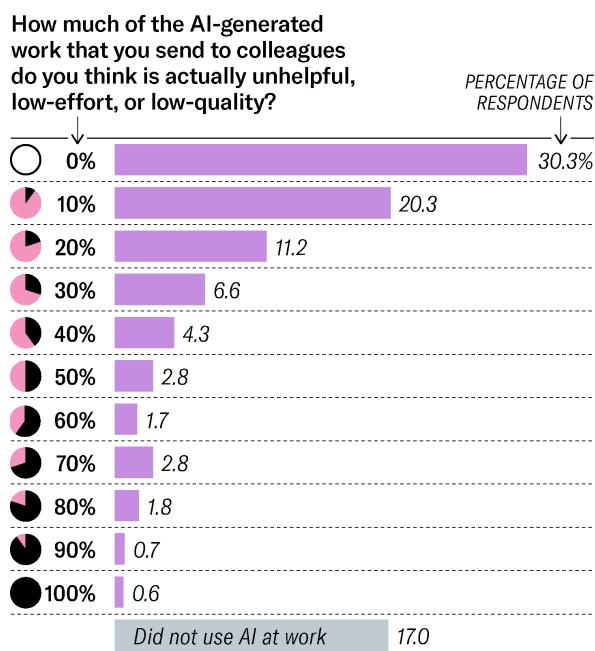
As researchers, two of us (Hancock and Niederhoffer) have spent decades tracking how technology shapes communication—from early work on communication patterns in internet chat rooms to the psychological dimensions of social media usage. Most recently, through our two-year collaboration between BetterUp Labs and the Stanford Social Media Lab we've tracked how people think, feel, and perform with AI in the workplace.

As CEO of BetterUp, another of us (Robichaux) has had a front-row seat to how AI is reshaping the way we work. In almost every conversation, the same questions come up: How do we make sure our AI investments pay off? How do we move fast without breaking things? Executives feel real pressure to mandate AI use, especially as boards push for leaner teams and larger spans of control. The implicit expectation is that AI will help managers “do more with less,” whether for growth, innovation, or both.

As we study the effects of generative AI, a troubling picture is starting to emerge. Our [initial research](#) revealed the presence of workslop. Of the 1,150 U.S.-based full-time employees across industries who participated in our survey, 41% of participants could recall receiving a specific instance of workslop that affected their work. Moreover, more than half of respondents admitted to *sending* workslop to colleagues. One in 10 said that 50% or more of the AI-generated work they sent colleagues was “actually unhelpful, low effort, or low quality.”

Do You Send Colleagues Workslop?

When asked how much of the AI-generated work they sent to colleagues was “actually unhelpful, low-effort, or low-quality,” half of survey respondents admitted to sending at least some AI-generated workslop.



Online survey of 1,150 full-time U.S. desk workers conducted in August and September of 2026 by BetterUp in partnership with the Stanford Social Media Lab.

Source: Kate Niederhoffer et al.



Now, we're starting to understand *where* workslop comes from: With further analysis on predictors and antecedents, the data we collected reveals that workslop emerges from pressures being felt by people at opposite ends of the organization.

At the top, many leaders are facing pressure to make responsible investment decisions about AI in the face of uncertainty and macroeconomic pressures. Boards increasingly view AI as a hopeful lever to compensate for slowing productivity and to maintain competitiveness in a rapidly digitizing global economy.

In response, leaders are using a blunt strategy, mandating that employees use AI broadly and quickly. We saw this in our data: 41% of employees in our survey reported that leadership encouraged them to use AI without detailed instructions or the deep, contextual understanding needed to meaningfully apply AI to specific tasks or workflows.

- One company president shared: “As a manager of managers, I feel that I am being pushed to drive AI usage through the organization with no clear vision, which leads to constant frustration at the board level. They want the use of AI, but other than saying ‘use it everywhere every day,’ there is no place our team can use it to actually be viewed as successful.”
- An executive vice president at another company said: “My org is still very afraid that we are not using AI ‘enough,’ so I don’t think there’s any way out but through. The examples that my fellow executives are setting is really pitiful. Quantity and ‘use of AI’ is prized more than quality and effectiveness. I’m looking for another job.”
- One vice president in our survey explained: “I think organizations (including mine) can be much more clear on how AI should be used specifically to the mission and objectives of the org and team. There’s such a huge mandate to ‘use AI!’ but no specification on what quality AI output looks like specific to work.”

Why People Create Workslop

In our survey, we asked employees how much of the AI-generated work they sent to colleagues is unhelpful, low effort, or low quality. Fifty-three percent of respondents admitted to doing this at least some of the time.

Why are people sending colleagues work they know to be subpar? Primarily because they feel stretched too thin, respondents say. As companies have

tightened budgets, consolidated roles, and asked employees to take on more tasks without formal role redesign, individual contributors and frontline managers are stretched more than ever. This has left employees psychologically depleted and juggling heavier workloads.

In this context, blanket mandates to use AI—often without the training, agency, or cultural trust to thoughtfully experiment with these powerful new tools—end up encouraging people to use AI performatively. These low-effort, low-value uses demonstrate compliance with directives to experiment, even as they shift the burden of the work onto the receiver. Hence, workslop.

To understand what's happening here, leaders need to see the bigger picture. Consider two datapoints. First, in BetterUp's dataset of more than 400,000 employees, surveys of employees have shown a decline of 2–6% since 2020 in the foundational mindsets that enable leadership and performance (e.g. focus, agility, strategic planning). Second, Edelman's 2025 Global Trust Barometer tells a parallel story: Employee trust in employers declined by three points, a shift they referred to as “unprecedented.”

These declines in both performance and trust are likely caused by convergence of pressures including: prolonged post-pandemic stress, economic and geopolitical uncertainty, hybrid work fragmentation, rising burnout, and generational shifts in employee expectations. The cumulative strain of these pressures results in depleted cognitive and relational capacities necessary for effective collaboration: trust, creative problem-solving, managers' capacity to recognize teammates, and goal achievement.

Workslop is both a symptom and potential accelerant of these deeper issues. A symptom because of the dynamics that incentivize it, and an accelerant because, left unchecked, it will damage relationships, harm teamwork, and undermine trust.

AI *will* increase productivity; we've seen this in our own research, and in both academic publications, and industry reports. But the pervasiveness of workslop shows that the mere presence of AI is no guarantee that organizations will see those gains. Rather, it suggests that many workplaces need to address serious underlying issues in order to see broad benefits.

Relieving the Pressure: How Leaders Can Reduce Workslop

Our research reveals several key factors that can make organizations more resistant to workslop. In addition to dialing back unclear blanket AI mandates, leaders should look inward and invest in a systemic response tailored to their organizations.

First, we found that people with a sense of competence and control over AI tools are half as likely to create workslop. There are many ways companies can build these mindsets for their teams, including increasing AI literacy investments, getting employees to share and model the AI practices that are currently hidden in the shadows, or forward deploying AI engineers to help teams effectively integrate AI into their workflows. These kinds of investments will produce a healthy return.

Our research also reveals that building a culture of trust within teams is an important protective factor against workslop. When people believe that they can admit to using AI, can raise concerns about how AI might affect quality, and can ask for feedback without fear of being stigmatized, they produce better work. For our participants, trust in their team reduced workslop by 61%.

Finally, we need to reframe how we think about workslop. Each instance of workslop is a signal of an organization under strain—pressed from above to prove AI's ROI, and from below by the mounting pressure from employee overload. It offers leaders important intelligence, because it exposes cracks in organizational design that need to be addressed for companies to successfully navigate the changes ahead.

Eye-rolling and casting judgment when we receive workslop is the easy way out. It's a classic example of the fundamental attribution error: We blame laziness and incompetence in individuals while discounting the influence of the situation. Addressing it requires a system-level response that operates on at least three levels:

- **Culture:** Leaders should start by rebuilding a sense of trust with a return to everyday practices of collaboration like giving feedback, asking questions, and making space for dialogue.
- **Practice:** Leaders need to create agency with AI by building clear expectations and norms for when and how to use AI, and explicit review processes that reinforce, not offload, human judgment.

- Accountability: Organizations need someone fluent in both technology and relationships. We can imagine a new role for forward deployed AI collaboration architects that can locate friction, tailor AI integrations to employee motivations and workflows, and connect AI strategies to specific outcomes.

The greatest irony of all is that to make AI work at work, we need to get better at being human. Leaders need to make space for the unpolished, slower-but-more-rewarding work of *human* collaboration. Without organizational changes that enable agency and trust, rather than AI mandates for overburdened teams, we'll all drown in the sludge of workslop.

See more HBR charts in **Data & Visuals** on HBR.org.

This article was originally published online on January 16, 2026.



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



Alexi Robichaux is the CEO and co-founder of BetterUp.



Jeffrey T. Hancock is the Harry and Norman Chandler Professor of Communication at Stanford University. He is also the Founding Director of the Stanford Social Media Lab and Director of the Cyber Policy Center at Stanford. His research focuses on psychological aspects of technology use, including AI and social media.

Copyright 2026 Harvard Business Publishing. All Rights Reserved. Additional restrictions may apply including the use of this content as assigned course material. Please consult your institution's librarian about any restrictions that might apply under the license with your institution. For more information and teaching resources from Harvard Business Publishing including Harvard Business School Cases, eLearning products, and business simulations please visit hbsp.harvard.edu.