

Workshop: Examining the prevalence, antecedents and consequences of low-quality AI-generated content at work

Alex Liebscher^{2*}, Angela Y. Lee^{1,3*}, Kristina Rapuano², Gabriella Kellerman⁴, Kate G. Niederhoffer², & Jeffrey T. Hancock¹

Affiliations:

¹Stanford University, Department of Communication

²BetterUp, Inc.

³University of Wisconsin-Madison, School of Journalism and Mass Communication

⁴Boston Consulting Group

*Both first authors contributed equally to this research.

*Corresponding authors. Emails: alex.liebscher@betterup.co and angela8@stanford.edu

Keywords:

Artificial intelligence, workshop, workplace behavior, interpersonal communication

Abstract (218 words)

Artificial intelligence (AI) adoption in the workplace has skyrocketed but produced ambiguous effects, raising questions of when AI use may be beneficial or harmful to productivity and collaboration. To understand the conditions under which AI workplace use is harmful, we examined workers' experiences with AI-generated workshop: low-quality, predominantly AI-generated content that appears to fulfill a given task, but lacks the substance to meaningfully advance the work at hand. We conducted a representative survey of American full-time desk-workers ($n = 962$) that found that 52.7% reported sending some amount of AI-generated workshop to others. This behavior was more common among people who had high trust in AI, low agency over AI use, organizational encouragement to use AI, and low psychological safety at work. A substantial proportion (38%) also reported receiving work from their colleagues or teammates that they would consider to be workshop, and that the effort involved in processing, discussing, and revising - or even redoing - the work cost them around 3.4 hours per month on average. In addition to consuming workers' time, receiving workshop negatively affected their mood and harmed their perceptions of their colleagues as trustworthy, reliable collaborators. Together, our findings paint a cautionary tale regarding AI integration into the workplace by highlighting how using AI under certain circumstances can be counterproductive and undermine interpersonal relationships.

Introduction

Advances in artificial intelligence (AI) appear poised to reshape the future of work (Selenko et al., 2022; Anthony et al., 2023; Brynjolfsson et al., 2025). Workers across nearly every industry increasingly turn to tools powered by large-language models (LLMs) like ChatGPT and Gemini to write emails, generate code, process data, and make decisions (Maslej et al., 2025). A recent national survey found that 28% of employed American adults report using ChatGPT for work (Sidoti & McClain, 2025), with a global study finding that 30% of consumer ChatGPT use is for work-related inquiries (Chatterji et al., 2025). Executives at major companies, schools, and healthcare institutions have also moved quickly to encourage their organizations to broadly adopt AI tools in hopes of augmenting worker productivity (Box, 2025; Oliver & Dabney, 2025; Lenharo, 2023).

Continued research on the effects of AI adoption, however, has tempered this initial excitement by demonstrating how workplace AI use can be both beneficial and harmful. While some AI tools have been shown to improve work outcomes by automating routine tasks or polishing written work (Noy & Zhang, 2023; Wu et al., 2025; Brynjolfsson et al., 2025), many fail to help people make meaningful progress on the tasks at hand. For instance, the MIT Media Lab found that fewer than 5% of public AI initiatives produced a “marked and sustained” impact on productivity and profitability (Challapally et al., 2025). Similarly, a longitudinal study of over 7,000 workplaces in Denmark found that implementing AI chatbots did not affect organizational outcomes (Humlum & Vestergaard, 2025). At the same time, there are growing complaints among everyday workers who are frustrated with receiving poor-quality AI-generated work from their colleagues (Hamilton, 2025).

Understanding the discrepancy between the anticipated benefits of AI and the observed mixed results in the workplace is imperative to informing decisions about engaging with AI tools at both the individual and organizational level. Previous studies have investigated this tension by looking at differences in AI adoption across industries (Maslej et al., 2025), shifting perceptions of AI among the public (Kennedy et al., 2023; Cheng et al., 2025), and frontier models’ performance on common work tasks (Doshi & Hauser, 2024; Chang et al., 2025). However, growing research on human-centered AI emphasizes how the effects of AI use are shaped not only by the tools’ technical abilities, but also by how people choose to use them in the context of their own lives (Shneiderman, 2022; Capel & Brereton, 2023; Hancock et al., 2020).

Specifically, workplace AI use may fail to improve productivity when people use these tools to take shortcuts. While unmanageable workloads and low motivation have long caused people to cut corners at work (Beck et al., 2017; Jonason & O’Connor, 2017; Ansari et al., 2024), the specific affordances of AI now allow people to produce plausible-looking work with less effort (i.e. a single prompt) - though the final product may not contain the substance or context to achieve the goals of the task (Cowgill et al., 2024; Herbold et al., 2023). Colloquially, frustrations with AI-generated products are known as slop, which references “the flood of lazy, derivative, endlessly recycled outputs” produced by LLMs (Madsen & Puyt, 2025). Although often used to describe mild annoyances - such as redundant email templates, sycophantic text, or poorly rendered images - such content could have serious consequences when shared in the workplace.

We propose that one reason why workplace AI use may fail to improve productivity is because people sometimes use these tools to generate *workslop: predominantly AI-generated content that fulfills the appearance of completing a task at work, but is perceived by the recipient as low-effort, low-quality, and/or not meaningfully advancing the task*. We conducted an online survey with a sample of American full-time desk-workers across diverse industries including technology, education, health, and marketing (N = 962). The survey examined the prevalence of encountering AI-generated workslop, the dispositional and situational factors that were associated with sending workslop, and the effects of receiving workslop on their affect, productivity, and perceptions of their colleagues. Furthermore, we elicited and analyzed 362 in-depth descriptions of specific instances where individuals received work from their coworkers that they believed to be workslop. We addressed the following research questions:

- RQ1: What is the prevalence of American desk-workers who send and receive AI-generated workslop?
- RQ2: What makes people perceive work to be AI-generated workslop?
- RQ3: What are the psychological and structural predictors of individuals a) sending AI-generated workslop to their colleagues? and b) of sending greater amounts of workslop to their colleagues?
- RQ4: What is the impact of receiving AI-generated workslop on American workers' affect, productivity, and perceptions of their colleagues?

AI-mediated communication in the workplace

The introduction of new technologies like email, instant messaging, and video-conferencing have repeatedly transformed how people communicate and collaborate in the workplace (Skovholt & Svennevig, 2006; Isaacs et al., 2002; Fauville et al., 2023; Colbert et al., 2016). While these tools enabled more flexible work across time and space, they also introduced many new challenges. Recent research on human-AI interaction suggests that the adoption of AI tools in the workplace is likely to follow a similar trajectory of producing both new efficiencies and new sources of friction (Brynjolfsson et al., 2025).

The use of AI-based tools such as ChatGPT to generate, alter, and share work can be understood through the framework of *AI-mediated communication*, which examines how the dynamics of information exchange are altered when AI systems intervene in human-to-human interaction (Hancock et al., 2020). This framework proposes that AI systems function as interlocutors that “modify, augment, or even generate” messages on behalf of a sender, to be interpreted by receivers who may or may not be aware of the AI’s involvement. This mediation of the traditional communication process alters both the characteristics of the content created and the social dynamics of the people involved (Jakesch et al., 2019; Hohenstein & Jung, 2020). For example, past work has found that “smart replies” that provide personalized recommendations for emails are often biased towards short, positive responses (“That’s great!”). When people rely on these smart replies instead of writing their own, they are perceived as being more cooperative by others, unless they are suspected as using AI, in which case they are evaluated as less trustworthy (Hohenstein & Jung, 2020). Similarly, more recent studies have found that content produced by LLM-based chatbots has some shared, distinctive characteristics. Content written by models like ChatGPT, Gemini, and Claude often contains certain punctuation (e.g., em-dashes), words and phrasings (e.g., “delve”, “not only X, but also Y”; see Juzek & Ward, 2025), and tones that are often perceived as sycophantic or robotic (Chen et al., 2025; Cheng et al., 2025).

Workslop as a harmful form of AI-mediated communication in the workplace. To date, hundreds of studies have examined how co-producing content with AI can affect psychological, interpersonal, and organizational outcomes. A clear theme across this literature is *heterogeneity*: the effects of AI depend on how people take up these systems within the practical constraints of their everyday tasks (Shneiderman, 2022; Capel & Brereton, 2023; Hancock et al., 2020; Markowitz, 2024). Among the many ways in which AI can play a role in the workplace, we operationalize *workslop* as a form of workplace AI use that results in partially or fully AI-generated content that is perceived by the receiver as being low-quality, low-effort, or of insufficient substance to contribute to the work at hand. With current models, it can manifest across diverse forms of media - such as documents, slide decks, emails, code, and spreadsheets. Workslop often appears professionally formatted while being excessively long, dense, repetitive, ornate, or containing unusual or nonsensical elements (Hoffman, 2024; Przystalski et al., 2026).

To identify the circumstances under which AI use is beneficial or harmful, human-centered AI research encourages considering individual-level factors, situational-level factors about the social context, and the affordances of current AI models (Shneiderman, 2022; Capel & Brereton, 2023; Hancock et al., 2020). For instance, individual differences, like how much people trust AI or prefer cognitively effortful work, may affect whether they produce workslop and to what extent (Cheng et al., 2025; Glikson & Woolley, 2020; Petty et al., 2009). Organizational factors, such as workers’ roles (i.e., individual contributor, manager) (Jarrahi et al., 2023), or team and organization norms around AI use (Georganta & Ulfert, 2024), may also influence if and how they use AI tools.

Specific characteristics of AI tools should also affect people's ability to produce work that is seen as meaningful or helpful. For instance, because LLMs have been trained to predict the next word in a sequence (Blank, 2023), they can quickly generate text that closely mirrors the style, shape, and structure of professional communication. This fundamentally alters the relationship between effort and output. What once signaled cognitive investment - length, polish, depth - can now be generated with a brief prompt and minimal sender effort. For example, Cowgill et al. (2024) argue that public access to LLM-based tools like ChatGPT has complicated hiring efforts by lowering the cost of producing signals that convey a candidate's qualities. The signals that have historically been used as indicators of genuine thought and investment can now emerge from statistical pattern-matching, without commensurate cognitive investment or effort (Jakesch et al., 2023).

This same tendency may inhibit LLMs' ability to produce work that is relevant to specific contexts. LLMs are optimized to identify modal responses that maximize human preference (Rudko & Bashirpour Bonab, 2025; Naveed et al., 2025). But whether it is a mundane scheduling email or as complex as a million-dollar grant proposal, the work people do everyday is steeped in context, such as the history of the meeting or project, its overarching objectives, and the contributions of other collaborators. Organizational communication is deeply shaped by norms, shared histories, audience expectations, and tacit constraints that are rarely fully articulated (Clark & Brennan, 1991; Hancock et al., 2020). These contextual features are often difficult to encode in prompts, particularly when users themselves have only partial awareness of what matters most. As a result, without careful oversight and tailoring, workers commit an interpersonal failure by sharing illegitimate AI-generated work that fails to consider context.

The interpersonal costs of workslop

Communication scholars have long emphasized that successful communication is a joint activity that depends on shared understanding and mutual effort (Clark & Brennan, 1991). It is especially one in which participants seek to minimize the work that both must perform. When content is superficially polished but contextually off-target, recipients may need to expend additional cognitive effort to process and evaluate the work, and revise or redo it if necessary. Indeed, research demonstrates that individuals exhibited higher accuracy when answering questions about human-written summaries compared to LLM-written summaries (Guo et al., 2025), suggesting that AI-generated content may be cheap to generate but costly to consume. Unlike obviously bad work that can be rapidly recognized and dismissed, content that looks "almost right" demands more careful attention and interpretive labor, which may thus increase cognitive load and consume more time.

Over time, the accumulation of such burdens may erode trust, slow collaboration, and negatively shape perceptions of colleagues' abilities. The perception of colleagues using AI to take a "shortcut" at work may undermine trust in not only the quality of the work, but also the sender themselves (Beck et al., 2017). People whose work is perceived or known to be using AI often face a steep interpersonal cost. A series of 13 experiments found that people who disclosed using AI in the process of working on diverse tasks, from art to writing to analytics, were consistently trusted less by their peers (Schilke & Reimann, 2025). Furthermore, the mere suspicion that AI is used can reduce trust in the quality of the work. Work that contains signals that are perceived to be indicative of AI use can be appraised more negatively, and even cause the sender to be perceived more negatively by their peers (Jakesch et al., 2019; Wu & Kelly, 2020). Past work suggests that this erosion of interpersonal trust is driven in part by reductions in the perceptions of the legitimacy of the person sending this work, who are seen as less capable due to the reliance on an external resource to complete their work (Schilke & Reimann, 2025). In the present study, we examine the cost of workslop on interpersonal perceptions and relational dynamics.

Method

Pilot study

A pilot study invited full-time American workers from the survey platform CloudResearch's Connect platform ($N = 150$) to describe their thoughts about workplace AI use (see Appendix 1 for details). After being presented with a brief definition of workslop, 40.3% of participants believed they had

at some point received work from colleagues they would consider to fit the definition. Moreover, 57.7% of participants believed that some amount of their work was low-effort or low-quality AI-generated work. Review of open-ended descriptions of examples of workslop they had received indicated that these experiences not only cost time to address or revise, but also affected their perceptions of the colleague who sent workslop to them. These preliminary insights were used to guide the final set of questions asked in the main survey.

Main survey

In September 2025, an additional 1,015 full-time American workers were recruited from CloudResearch's Connect platform. Sample size was determined using the pilot study to achieve a margin of error of ± 5 percentage points for our primary prevalence estimate at 95% confidence, requiring approximately 370 participants. We oversampled for additional precision. To be eligible, participants were required to be over 18 years old, currently employed in a full-time role, fluent in English, and located in the United States. Participants had to have a CloudResearch task approval rating of over 90% on more than 100 tasks. To represent a population of desk workers, participants were only eligible if they classified their work arrangement as white collar, professional, creative, IT, or public sector or government. An equal balance of men and women was requested. All procedures were approved by the Salus Institutional Review Board and participants were compensated with \$3.00 upon completion (\$12/hr). Upon consenting, participants were asked to answer some questions about themselves and their workplace. For a full list of all survey measures used, please see Appendix 2. Data and code are available at https://osf.io/xh2wn/overview?view_only=675432f416474ce592257dfd684343e3.

Participants

Eleven participants were excluded for failing one or more of three attention checks. A manual review of open-ended responses further excluded 42 participants for misunderstanding the primary workslop description question (i.e., wrote nonsense, described an instance of high-quality AI-generated work). The final analytical sample included 962 participants. There were approximately equal numbers of men ($n = 480$, 49.9%) and women ($n = 482$, 50.1%). The majority were White ($n = 737$, 76.6%), 7 were American Indian or Alaskan Native (0.7%), 72 were Asian (7.5%), 122 were Black or African American (12.7%), and 24 were another ethnicity or preferred not to say (2.5%). The mean age was 40 years old ($SD = 10.3$ years). A subset of 71 were Hispanic or Latino (7.4%), with the majority not being Hispanic or Latino ($n = 891$, 92.6%). The sample was relatively highly educated, with 756 indicating they held a Bachelor's degree or higher (78.6%).

Approximately half of the participants (55.6%) were individual contributors ($n = 535$), with 337 who were frontline managers (35.0%) and 90 who were managers-of-managers (9.4%). For the analyses below, these two categories of managers were aggregated into one. Mean work experience was 18.4 years ($SD = 10.2$ years) and participants worked regularly with a mean of 12 other people (*median* = 8, $SD = 19.2$). Participants varied in the degree in which they worked in-person vs. remotely, with 300 working in-person all of the time (31.2%), 201 working fully remotely (20.9%), and 461 working some degree of hybrid (47.9%). The mean individual income was \$74,655 (*median* = \$50,000, $SD = \$48,633$, *range* = \$10,000 – \$300,000). Participants came from a variety of industries, with the top three being Technology (18.2%), Government / Public Sector (17.9%), and Healthcare (16.2%) (see Appendix 3 for a full table).

Measures

See Appendix 2 for detailed descriptions of all measures administered, including response options and the full set of items. Validated scales were used whenever possible and Cronbach's α 's were used to assess reliability. The term "workslop" was not used anywhere in the survey.

Individual differences in psychological traits. Participants completed the *Ten Item Personality Inventory* (Gosling et al., 2003), a 10-item scale assessing openness to new experiences (Cronbach's $\alpha = .53$), conscientiousness ($\alpha = .64$), extraversion ($\alpha = .77$), agreeableness ($\alpha = .48$), and neuroticism ($\alpha =$

.76). Responses were rated on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree). They completed the 5-item *Need for Cognition* scale (Cacioppo & Petty, 1982) to assess intrinsic motivation for effortful cognitive activities (e.g. “I like to have the responsibility of handling a situation that requires a lot of thinking”). Items were rated on a 5-point scale (1 = extremely uncharacteristic of me, 5 = extremely characteristic of me) ($\alpha = .82$). Finally, they completed the 3-item *Pride in Work* scale (Schaufeli et al., 2006) assessing their dedication to their work (e.g. “How often do you find your work to be full of meaning and purpose?”). Items were rated on a 5-point scale (1 = never, 5 = always) ($\alpha = .80$).

Work characteristics. Each participant described their *job title*, *years of work experience*, the *number of people they work with regularly*, their *role* (individual contributor, manager, manager-of-managers), and *work arrangement* (1 = in-person all the time to 5 = remote all the time). They selected the *industry* they worked in from a drop-down list of industries matched to the 2020 United States Census occupational categories. *Annual individual income* was selected in U.S. dollars as a bracket from less than \$25,000 to more than \$300,000 in increments of \$25,000; they could also prefer not to say. To understand workplace norms around AI, they also completed a single-item developed by the research team about *AI use encouragement* (“My supervisor or organization encourages us to use AI”) (1 = strongly disagree, 5 = strongly agree).

Psychological experiences at work. To understand how participants felt at work, they completed the 5-item *Quantitative Workload Inventory* which assesses the amount and intensity of job obligations (e.g. “How often does your job require you to work very fast?”, “How often do you have to do more work than you can do well?”) (Spector & Jex, 1998). Items were rated on a 5-point scale (1 = never, 5 = always) ($\alpha = .86$). They also completed the 5-item *Edmonson Psychological Safety* scale (Edmondson, 1999) which measures how comfortable individuals feel speaking up with ideas, mistakes, and requests for help within their team (e.g. “Members of this team are able to bring up problems and tough issues”). Items were rated on a 7-point scale (1 = very inaccurate, 7 = very accurate) ($\alpha = .83$). Finally, they completed two subscales from the *Team Climate Inventory* (TCI) (Anderson & West, 1998). Three items were used to assess alignment with the team’s *vision* (e.g. “I believe [my team’s] objectives are worthwhile to the organization”) and *task orientation* to high-quality work (e.g. “My colleagues and I monitor each other so as to maintain a high standard of work”). Items were rated on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree) and both subscales had good reliability ($\alpha = .86, .77$).

Workplace AI use. Participants indicated *how frequently they used AI* in their own work, and how often they *shared work generated with AI* to colleagues. Both items were rated on an 8-point scale (1 = never, 2 = once a month, 3 = several times a month, 4 = once a week, 5 = several times a week, 6 = daily, 7 = several times a day, 8 = more than several times a day). “Work” was specified to include a broad range of products including “documents, emails, slide decks, code, agendas, newsletters, budgets, multimedia content, etc.”

Attitudes towards AI. Next, they completed a 3-item measure of *general trust in AI* from Jakesch et al. (2019) which builds on the Organizational Model of Trust (Mayer et al., 1995) to assess AI’s perceived 1) ability to achieve core tasks effectively (“I believe that AI will produce output that is accurate.”), 2) benevolence to produce positive outcomes for the public (“I believe that AI will have a positive effect on most people.”), and 3) integrity in facilitating user safety (“I believe that AI will be safe and secure to use.”). Items were rated on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree). In line with past work (Jakesch et al., 2019; Ma et al., 2017), the measure was the mean of its items ($\alpha = .89$).

Participants also completed a 6-item measure of *AI mindsets* adapted from Lee and Hancock (2024), which are the core assumptions that people hold about the nature of their relationship to AI. Three items assessed *agency mindsets* about the power dynamic between themselves and their AI use, such as whether they feel in control of this use (“I’m good at using AI to do things that are important to me”) or if

they feel dependent on their use (“I feel like using AI makes me less capable”) ($\alpha = .64$). Three items assessed *valence* mindsets about whether they felt using AI is generally personally beneficial (“Using AI facilitates my learning and growth”) or harmful (“Using AI is harmful to me”) ($\alpha = .83$). Items were rated on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree).

Sending AI-generated workslop. Each participant indicated how often, if at all, *they personally shared AI-generated workslop* to their colleagues (“How much of the AI-generated work that you send to colleagues do you think is actually unhelpful, low effort, or low quality?”). Response options were listed on a 0% to 100% scale in increments of 10%. For individuals who indicated they had shared workslop (i.e., responses greater than 0%), they reported how much the following *motivations* influenced their behavior: 1) the urgency of the task, 2) the importance of the task, 3) the ambiguity of the task, 4) the competing demands of this task relative to other needs, 5) their pressure to perform at work, 6) their desire to buy more time, 7) their lack of expertise to complete the task themselves, 8) their need for a thought partner, and 9) their organization encouraging them to use AI “no matter what.” Items were rated on a 5-point scale (1 = not at all, 5 = definitely). These items were generated based on pilot survey data.

Receiving workslop. Participants were asked if they had received workslop in their jobs in the last month, which was defined as “work content that you believe is AI-generated that looks like it completes a task at work, but is actually unhelpful, low quality, and/or seems like the sender did not put in enough effort” (1 = Yes, 0 = No). Participants were asked *how much of the work* they received in the last month they would consider to fit this description on a 0% to 100% scale (in increments of 10%), *how many people* they received this kind of work from, and *how many hours this month* they had spent dealing with this kind of work. They also indicated how much of their colleagues’ AI-generated work they would consider to be helpful or high-quality (0% to 100% scale, in increments of 10%).

To obtain a richer understanding of their experiences, they were asked to think of a *real example* of workslop they encountered in the last month. A series of open-ended questions asked them to describe the workslop, why they believed it was AI-generated, and the impact this had on their ability to complete and advance their own work, if any. They also indicated who the work came from (e.g. a teammate, a direct report), and if they had taken any of the following actions to address their belief that the work was AI-generated: 1) spoke to the colleague directly, 2) spoke to the colleague’s manager, 3) spoke with their manager, 4) spoke with a peer or teammate, 5) spoke to human resources, 6) documented the occurrence, or 7) did nothing and moved on.

Consequences of receiving workslop. Following best practices for eliciting *time cost* calculations (Ernala et al., 2020), participants indicated the time they spent engaging with their example of workslop. They then broke down the percentage of time that was spent reading or processing the work, evaluating the work, discussing the work, revising or fixing the work, or redoing the work.

The *emotional cost* of receiving workslop was assessed with an adapted 10-item version of the Positive and Negative Affect Scale (PANAS) from Watson et al. (1988) (e.g., how happy, content, frustrated they felt after receiving workslop). Items were rated on a 5-point scale (1 = not at all, 5 = extremely).

The *cognitive cost* of engaging with their example of workslop was assessed with the 1-item Paas mental effort scale (Paas et al., 1992) (1 = very low mental effort, 9 = very high mental effort).

To examine the *interpersonal costs* of receiving workslop, participants completed a 9-item scale based on Mayer et al.’s (1995) Organizational Trust measure to indicate how the experience changed their perception of the colleague’s ability (e.g. “They are very capable of performing their job”), benevolence (e.g. “They really look out for what is important to me”), and integrity (e.g. “Sound principles seem to guide their behavior”). Items were rated on a 5-point scale to capture change (1 = I thought much less of them, 3 = Didn’t change how I saw them, 5 = I thought much more highly of them). They also completed an 8-item measure of *interpersonal impressions* based on the Stereotype Content Model (Cuddy et al.,

2008) to indicate if receiving workslop affected perceptions of their coworkers' warmth (e.g., friendliness, pleasantness) and competence (e.g., diligent, capable). Items were rated on a 5-point scale to capture change (1 = Less, 3 = No change, 5 = More) and both subscales had acceptable reliability ($\alpha = .79, .89$). Finally, they indicated if this experience changed *how likely they were to want to work* with this colleague in the future (1 = would strongly avoid working with them again, 5 = would strongly want to work with them again).

Analytical approach

The statistical approach to understanding the predictors of sending workslop (RQ 3a-b) used a hurdle Poisson analysis to examine who sent any workslop, and who sent the greater amounts of workslop (Mabire-Yon, 2025). The first dependent variable was a binary indicator of whether the respondent reported that more than 0% of the AI-generated work they send is workslop. Given potential model instability with many exploratory predictors, we conducted a robustness check against a bootstrap LASSO model in line with Lima et al. (2020) (see Appendix 6). The second dependent variable is treated as a positive count truncated at 10 (i.e. 100% of their AI-generated work is workslop). Out of concern that treating the outcome as a count variable provides an inaccurate likelihood function, we also compare this to a Beta regression model. All alternative model details may be found in Appendix 6. Gender (1 = Woman) and manager status (1 = is manager) were dummy-coded. Respondents were excluded for not providing their level of education ($n = 1$). All other predictors are standardized to a mean of 0 and standard deviation of 1.

Thematic content analysis was used to identify the signals people use to detect workslop (RQ2) by qualitatively coding participants' descriptions of a specific instance of workslop they received in the last month (see Appendix 9 for details and codebook). Good inter-rater reliability was obtained for all codes (Cohen's kappa = .78 - .95). The full set of anonymized responses can be found at https://osf.io/xh2wn/overview?view_only=675432f416474ce592257dfd684343e3.


To assess the impact of receiving workslop on affect, productivity, and perceptions of their colleagues (RQ4), we conducted four linear regressions on the outcome variables of amount of time spent per month dealing with workslop, mean negative emotional response to workslop, cognitive effort of dealing with workslop, and desire to work with the sender again in the future. The same set of predictors from the hurdle analysis were used, in addition to a dummy-coded variable indicating the recipient's relationship to the workslop sender. A standardized estimate of time spent dealing with the instance of workslop was included as a predictor for the three latter models. As an exploratory analysis, we compare these four consequence models to identify potentially meaningful predictors using a false discovery rate (FDR) adjustment ($q < .05$) (Benjamini & Hochberg, 1995), and discuss the directionality of effects.

Results

Sending AI-generated workslop

How many people send workslop? Many people reported using AI at work (83.2%) and shared AI-generated work with their colleagues (69.1%). Around half of workers believed at least some of their AI-generated work is unhelpful, low-effort, or low-quality (52.5% of total sample; 63.1% of workplace AI users) (RQ 1). Figure 1 illustrates the notable variance in how much of their AI-generated work they would consider to be workslop ($M = 17.9\%$, $SD = 22.4\%$, $range = 0\%$ to 100%). There is a moderate correlation between the amount of workslop one sends and the amount they report receiving (Spearman's $\rho = 0.45$, $p < .001$). A small set of individuals were *prolific producers* of workslop (9.6%) who say that more than 25% of their AI-generated work was workslop and that they shared AI-generated work at least once a week (see Appendix 4 for more descriptives).

Figure 1. Prevalence of American desk-workers sharing AI-generated workslop with colleagues




Note. Visualization of how much of their own AI-generated work American desk workers consider to be “unhelpful, low-effort, or low-quality.” The x-axis indicates what percentage of the AI-generated work they sent they would consider to be workslop. The gray bar on the left indicates the 16.8% of desk workers who did not use AI at work.

What were people’s self-selected motivations for sending workslop? Participants rated the degree to which multiple motivations for sending workslop resonated with them. Results of an exploratory factor analysis on nine statements describing potential motivations for sending workslop indicated two primary reasons why people send workslop (see Appendix 5 for details). The first factor was interpreted as *work pressure*: people produce workslop because the task was urgent and important, they were under pressure to do many things at work, and they wanted to buy time ($M = 2.97$, $SD = 1.15$). The second factor related to *need for support* where people produced workslop because they did not feel they had expertise in the task and wanted a thought partner to help them ($M = 2.47$, $SD = 1.14$).

What are the predictors of individuals sending workslop (RQ 3a-b)? As shown in Figure 2, both individual-level characteristics and structural factors of the workplace were associated with sending workslop. The logistic model indicates that people were more likely to send some amount of workslop when they had low agency over their AI use ($\beta = -0.32$, $SE = 0.10$, $p = .001$) and were more neurotic ($\beta = 0.25$, $SE = 0.09$, $p = .006$). At the structural level, people were more likely to send some amount of workslop the more they believed their organization encouraged AI use ($\beta = 0.47$, $SE = 0.08$, $p < .001$). The Poisson model indicates that the people who greater amounts of workslop were those who trusted AI more ($\beta = 0.14$, $SE = 0.05$, $p = .005$), had less agency over AI ($\beta = -0.13$, $SE = 0.04$, $p = .005$), perceived AI as negative ($\beta = -0.37$, $SE = 0.06$, $p < .001$), were younger ($\beta = -0.10$, $SE = 0.03$, $p = .004$), less conscientious ($\beta = -0.19$, $SE = 0.03$, $p < .001$), and less neurotic ($\beta = -0.17$, $SE = 0.04$, $p < .001$). At the structural level, individuals who sent more workslop also felt lower psychological safety ($\beta = -0.12$, $SE = 0.04$, $p = .003$). Additional analyses in Appendix 6 corroborate these results with alternative model specifications.

Figure 2. Dispositional and situational predictors of sending AI-generated workslop




Note. All predictors were mean-centered and standardized to 1 SD. 99% CIs shown. Significant predictors ($p < .01$) are red. Gender was dichotomized to 0 = man and 1 = woman. Manager status was dichotomized to FALSE = individual contributor and TRUE = manager.

Receiving AI-generated workslop

How many people report receiving workslop? As shown in Figure 3a, 37.6% of workers said they received work from colleagues that they believed to be AI-generated workslop ($SD = 48.5\%$) (RQ 1). On average, they estimated 14.4% ($SD = 19.7\%$) of all work they received from colleagues fit the description of workslop (Figure 3b), but that 43.9% ($SD = 36.8\%$) of the work they received from colleagues was helpful and high-quality AI-generated work.

We also examined participants' perceptions of how many individuals were responsible for sending AI-generated workslop, in line with past work on incidence rates (Spronk et al., 2019). On average, they said that 2.2 colleagues had sent them workslop in the last month - suggesting that such behavior is not widespread and instead comes from a small group of individuals. To understand this in greater depth, we compared participants' reports of how many people sent them workslop against the number of people they worked with on a regular basis ($M = 12.5$, $SD = 19.2$). The average incidence rate of workslop (responsible colleagues divided by number of close colleagues) was 0.18 ($median = 0.07$, $SD = 0.27$) (Figure 3c), suggesting that about 1 out of 5 close colleagues was responsible for sending workslop.

Figure 3. American workers’ experience with receiving AI-generated workslop

Note. (a) Visualization of how many full-time American desk-workers reported receiving AI-generated work that is “unhelpful, low-effort, or low-quality” in the last month (37.6%) (95% bootstrap CIs). (b) Workers’ estimations of how much of the work they receive from their colleagues they would consider to be workslop (14.4%). (c) Workers’ estimations of how many people sent them workslop in the last month relative to the number of people they work closely with on a regular basis (0.18).

What are people’s experiences with receiving workslop? We asked participants who indicated they received workslop to describe a concrete incident that occurred in the last month, using a combination of survey questions and open-ended responses elaborating on the details of the incident (e.g. why they believed it was AI-generated, who it came from, and the impact on their work). In total, the dataset contains 362 examples. Qualitative analysis revealed that workslop took on many different forms including emails, code, presentations, project deliverables, data analyses, and financial documents. Good inter-rater reliability was obtained for all codes (Cohen’s kappa = .78 - .95) (See Appendix 7 for details and codebook). Most workslop came from participants’ teammates (40.6%), followed by their direct reports (18.8%), peers on other teams (18.8%), and their supervisors (8.6%) (see Appendix 8). This descriptive data demonstrates the breadth of work tasks in which participants encounter workslop.

As shown in Table 1, workers described five primary signals for detecting workslop (RQ 2) (see Appendix 9 for the codebook). First, certain *stylistic or linguistic features* were perceived as indicative, such as flowery, verbose, robotic, or sycophantic language; bullet-points, boldface, em-dashes and emojis used as formatting choices; and blurry features in images or graphics. Second, some felt the work *did not sound like their colleague* based on their personal knowledge of the sender. Third, some noticed the work *missed important context or contained errors* that did not meaningfully advance the work at hand by being too generic. For example, plans for a work proposal that referenced products their company did not sell, or the use of nonstandard metrics to assess specialized project outcomes. Fourth, the work was perceived as being *completed too fast* for realistic human work. Finally, AI use may have also been disclosed proactively by the sender, explicitly or accidental through residual information (“This was generated with ChatGPT”), or revealed retroactively after questioning.

To better understand how workers dealt with receiving workslop, we asked them to indicate how they reacted, if at all (see Appendix 10). Almost half (44.5%) of workers said they spoke to the colleague directly. A much smaller proportion (14.6%) spoke to a different peer or teammate, and 9.9% spoke to their manager about the incident. Still, 39% of workers did nothing after receiving workslop and moved on.

Table 1: Reasons for suspecting work to be AI-generated workslop

Reason for suspecting workslop	Prevalence (%)	Examples
Stylistic or linguistic cues	33.7%	“The content was accurate, but stylistically it was stereotypical AI writing, with bulleted points, em-dashes, boldface headings, etc.” “I believe you can do it! I would absolutely love to be a part of this exciting meeting and delve into the topics with you further. It's not annoying - it's motivating! Keep it up. I could continue, but I'm making myself sick. Generative AI has a disgustingly syrupy and syncophantic writing style that is suited only for psychopaths.” “It was obvious. Especially with AI generated humans, you can always tell by the hands.”
Missing context or contains errors	26.9%	“These type of letters are sometimes hard to draft because you have to know the whole history of the account and why it was not paid. I could tell right away that the letter was done with the AI help. It was too long and it mentioned some codes and regulations that didn't apply to that specific letter and problem.” “It included products that we no longer or have never offered at our specific location. It also drew incorrect conclusions that made no sense in the context of our workflow.” “The presentation used a number of metrics we don't typically consider at work and weren't really useful to the overall discussion. The sales director couldn't really speak to why he'd included them either.”
Did not sound like their colleague	18.5%	“An email from a coworker that seemed like the use of ChatGPT was likely. It communicated relatively effectively but seemed unlike their usual style of writing. Certain word choices and a change in sentence structure seemed uncharacteristic.” “The email did not sound like her at all. She always sounds friendly in her emails when she speaks with me and this email was different. She sounded stiff and stilted in the message. Some of the words in the message are not ones she would have used.” “It is very easy to identify when you know the person. You recognize words and terms they don't usually use. The words are highly technical and come from a knowledgeable person. You know they didn't write it.”
AI use was disclosed, proactively or	1.4%	“The email title always starts with " AI translation.”” “The employee told me that they used AI to generate the report.”


retroactively		“I ask him more about it and he told me that he put the prompts in AI to write him a report based on his prompts and information that he provided to the AI.”
Speed of task completion	5.3%	“It's obvious from the phrasing, especially with certain individuals that I've worked with for a while, that they did not write that themselves. It also often comes back to me at speed at which I know they're not capable of writing that level of content themselves.” “It was sloppy and out of order. It contained info that would take a human days to complete”

Note. Categories are not mutually exclusive.

How does workslop affect people's reported productivity (RQ 4)? The people who received workslop reported spending an average of 3.3 hours dealing with such work in the last month ($SD = 5.2$ hours) (Figure 4a). There was notable variance (*range*: 0 hours to 36 hours), with a third (29.2%) indicating they spent 0 hours per month dealing with workslop. We also calculated the time that participants said they had spent dealing with the specific instance of workslop they encountered. On average, an instance of workslop cost people around 1.7 hours ($SD = 7.5$ hours).

More granular items revealed how this time was spent. On average, people said they spent the largest proportion of time reading and processing the work (35%, $SD = 0.26$), followed by evaluating the work (21.4%, $SD = 0.14$) and discussing the work with others (16.6%, $SD = 0.16$). They also spent a portion of time revising or fixing the work (15.7%, $SD = 0.16$), or redoing it entirely (11.4%, $SD = 0.15$) (Figure 4a).

Figure 4. Consequences of receiving workslop



Note. (a) Visualizations of the consequences of workslop on productivity, including self-reported hours spent dealing with low-quality AI-generated work and proportion of time dedicated to different responses to the work. (b) Quantification of the emotional responses to receiving workslop. (c) Quantification of the interpersonal impact of receiving workslop, including change in factors of interpersonal trust and change in factors of the Stereotype Content Model.

What are the cognitive consequences of receiving workslop (RQ 4)? On average, participants indicated on the Paas mental effort scale that it took considerable effort to understand, process, or revise the workslop they received ($M = 5.4$, $SD = 2.1$). On a scale rated from 1 (Very low mental effort) to 9 (Very high mental effort), this mean value is similar to the cognitive demands of, for example, engineering students completing a computer-simulation task (Aubteen Darabi et al., 2007) or physicians answering difficult medical exam questions (Szulewski et al., 2017). In other words, desk workers dealing with workslop demands cognitive load typical of domain experts engaging in challenging work.

What are the emotional consequences of receiving workslop (RQ 4)? Recipients of workslop feel on average annoyed ($M = 2.82$, $SD = 1.2$; 31.2% feel very or extremely) and frustrated ($M = 2.64$, $SD = 1.3$; 29.3% feel very or extremely). The full set of emotional reactions may be found in Appendix 11.

What are the interpersonal consequences of receiving workslop (RQ 4)? We assessed the interpersonal cost of receiving workslop through one item (future collaboration intention) and two measures: organizational trust dimensions of ability, benevolence and integrity; and the Stereotype Content Model dimensions of warmth and competence. First, we found that receiving workslop made approximately a third of participants (36%) want to avoid working with that colleague in the future, 50.3% indicating no change, and 13.8% indicating greater intentions to work together in the future. Second, one-sample t-tests against no-change ($M = 0$) suggested that receiving workslop reduced perceptions of the colleague's ability ($M = 2.64$, $t(364) = -8.4$, $p < .001$, Cohen's $d = -0.89$), benevolence ($M = 2.72$, $t(364) = -6.5$, $p < .001$, Cohen's $d = -0.69$), and integrity ($M = 2.80$, $t(362) = -5.6$, $p < .001$, Cohen's $d = -0.59$). Similarly, after receiving workslop, workers perceived their colleagues as less competent ($M = -0.58$, $t(366) = -14.2$, $p < .001$, Cohen's $d = -1.50$) and warm ($M = -0.28$, $t(366) = -8.5$, $p < .001$, Cohen's $d = -0.89$) (Figure 4c). See Appendix 11 for more details.

What factors predict weaker or stronger perceived consequences? To understand the factors that influence the perceived consequences of receiving workslop, we fit linear regression models for four consequence outcomes and identify significant predictors after adjusting for the false discovery rate (see Appendix 12). The clearest pattern emerged for cognitive effort: managers reported expending significantly more effort than non-managers, as did those higher in work pride. Higher education and spending less time on the workslop instance were both associated with lower effort. For willingness to collaborate with the sender again, trust in AI was the sole significant predictor, with higher trust associated with greater willingness. Hours spent dealing with workslop and negative emotional responses had no significant predictors after correction.

Discussion

As organizations rapidly integrate AI into their workforces, there is an urgent need to understand the conditions under which workplace AI use supports - or undermines - productive collaboration. Answering this question requires going beyond analyses of the technical capabilities of frontier models, to consider the perspectives of the everyday workers who are on the frontlines of technological change. To meet this need, we surveyed full-time American desk-workers ($N = 962$) to understand how they contended with the day-to-day implications of sending and receiving AI-generated work.

Examination of these experiences revealed that many workers have encountered a specific, problematic form of workplace AI use, which we term *workslop*: low-quality, predominantly AI-generated content that appears to fulfill a task but lacks the substance needed to meaningfully advance the work at hand. We advance efforts to understand the effects of AI on organizations by (1) identifying workslop as a widespread but counterproductive application of AI, (2) characterizing the individual and organizational conditions that make sending workslop more likely, and (3) demonstrating that receiving workslop not only impedes productivity but also undermines workplace relationships.

Conceptualizing AI-generated workslop

Relative to past studies of workplace AI use that focus on documenting the technical limitations of AI (i.e., producing inaccurate information like hallucinations; Farquhar et al., 2024), workslop can be understood as a failure mode of AI-mediated communication (Hancock et al., 2020) because it emerges

from how the AI mediates the sender-receiver relationship: senders offload cognitive work to AI systems that lack full context, creating work with manipulated quality signals that impose disproportionate burdens on the recipient. Analysis of workers' experiences with workslop indicates that it comes in many forms, ranging from interpersonal communications (e.g., emails, Slack messages) and internal drafts to public-facing work and presentations. Below, we discuss several distinctive features that characterize workslop across modalities.

First, it is important to note that workslop is an inherently perceptual phenomenon. Whether a given artifact is experienced as workslop depends on the recipient's interpretation of whether it is low-quality, the product of low-effort, or largely AI-generated. These judgments are important because workers generally have limited visibility into the steps their colleagues take when writing emails, developing slides, or producing other work artifacts. Similarly, they often lack information about if, when, or how AI tools were used unless it is explicitly disclosed or accidentally revealed (i.e., leaving in text that says "Claude can make mistakes"). As a result, even work produced with substantial prompting and thoughtful revision may therefore still be perceived as workslop if that effort is not legible to the recipient. Our results suggest that the mere perception of AI being used to cut corners at work can carry significant consequences. This aligns with past work finding that even when individuals could not confirm if AI was used or not, they rated those suspected of using AI as significantly less trustworthy (Schilke & Reimann, 2025). In this context, the perception of AI-enabled shortcutting - rather than confirmed AI use or technical error - can be sufficient to trigger negative interpersonal judgments.

Second, a feature that differentiates workslop from merely low-quality work is that it often appears legitimate on the surface, but lacks the substance needed to meaningfully advance the work at hand. For example, workers described well-formatted emails "with broken links and misinformation about our clinics", code that looked functional but was "more complicated than necessary", and passable-looking data analyses that "did not have anything to do with our products." The affordances of AI allow workslop to appear complete, extensive, and aligned with general genre expectations (Blank, 2023). In contrast, low-quality work produced by humans often contains signals (e.g. fragmented phrasing, grammatical errors, or inclusion of placeholder images) that indicate the work is unfinished and should be treated as provisional (Beck et al., 2017; Cowgill et al., 2024). This discrepancy can accrue rhetorical debt (MacArthur, 2025): long-term burdens that must be shouldered when people take shortcuts by producing "not quite right [work] that we postpone making right." Like the concept of technical debt in software engineering, where reliance on short-term patches causes infrastructural issues, sharing workslop can let people appear productive while simultaneously incurring downstream costs.

Third, workslop is characterized by effort asymmetry. When people use AI to create low-quality work, the person receiving it must spend time and effort reading, evaluating, and revising it. For example, one freelance writer shared how a collaborator's decision to send an AI-generated draft "saved the [collaborator] 20 minutes but cost [the writer] 45 minutes" because they had to fact-check the content and fix errors. In this vein, workslop violates the principle of least collaborative effort (Clark & Brennan, 1991) which proposes that good communication should reduce effort for both the person sending the message and the person receiving it - not just one or the other. This effort asymmetry is further compounded by expectancy violations. Workers described noticing stylistic and contextual cues that deviated from what they expected from particular colleagues, which triggered heightened scrutiny of the work. For instance, one worker wrote that "[their] attention kept being drawn to whatever words [they] noticed that [they] had seen used often by AI." As expectancy violations theory suggests (Burgoon, 1993; Burgoon & Walther, 1990), such deviations call for increased attention and evaluative processing, leading recipients to expend additional effort on work they might otherwise have accepted at face value.

Prevalence and predictors of sending workslop

Both sending and receiving workslop appears to be prevalent. More than half (52.5%) of full-time American desk workers have sent to their colleagues work that they considered to be workslop, and 38% have received such work in the last month. Closer examination of the incidence rate of receiving workslop (0.18) suggests that most workslop is produced by a small proportion of workers rather than broadly

distributed. This pattern echoes recent findings that a small proportion of online users generate an outsized share of harmful online content (Lee et al., 2025). Additional research may help quantify the extent which these prolific producers of workslop send a disproportionate amount of workslop compared to other producers (cf. Smith et al., 2014).

Who are the individuals who create workslop, and what factors amplify their behavior? Workers reported that their primary motivation for sharing workslop revolved around *work pressure*, such as having too many tasks to complete and not enough time. We identified the antecedents of workslop with greater precision by using a hurdle Poisson model to examine the individual-level and workplace-level predictors that were associated with 1) sending any workslop at all, and 2) sending greater magnitudes of workslop. Results from the initial hurdle indicated that people who were neurotic, perceived themselves as having little agency over AI, and have been encouraged to use AI by their organization were more likely to have sent workslop. Individuals who sent the greatest *amount* of workslop tended to be younger, less neurotic, less conscientious, and in organizations with low psychological safety where they did not feel comfortable seeking help or having conversations about capacity. They also had distinct attitudes towards AI, having high general trust in AI's capabilities but low agency and optimism about their own ability to use AI in meaningful ways. Together, these findings highlight how individual differences interact with workers' attitudes towards AI and workplace dynamics to shape whether people use AI in problematic ways.

These findings have important implications for approaches to workplace AI use. As organizations increasingly encourage or even mandate AI adoption (Stewart, 2025), such approaches run the risk of backfiring by increasing workslop instead of productivity. Workers facing intense workloads and pressure to adopt AI may use these tools as a shortcut (Beck et al., 2017), saving themselves time in the short-term but creating more work for others down the line. Such negative applications of AI may be particularly likely when people have disproportionate trust in AI and low confidence in themselves. Past work on automation complacency indicates that over-reliance on automated systems can increase errors when workers pay less attention to reviewing and correcting their outputs (Parasuraman & Manzey, 2010; Klingbeil et al., 2024; Spatola, 2024). When AI mandates are introduced without commensurate training, clear guidelines, or psychologically safe environments where employees feel comfortable questioning or resisting inappropriate use, workers may be left ill-equipped to deploy AI in ways that avoid shifting effort onto recipients.

The contrasting effects of neuroticism between the two models is worth noting. Although neuroticism is generally considered to negatively impact work performance (Nitu et al., 2023), it appears to reduce the overall amount of workslop one may produce. This finding was unexpected and deserves future investigations to better understand the limits of our model. One possible explanation is that a sufficiently neurotic individual may act carelessly when under stress (producing some amount of workslop), but even more neurotic individuals may be characterized by resignation or fear, thus sending less workslop.

Workslop undermines productivity and interpersonal relationships at work

The consequences of workslop identified in this study provide an important counterweight to dominant narratives that frame workplace AI use as uniformly beneficial (Mendy et al., 2025). Although individual workers can save time and effort by using AI to automate tasks, sending workslop redistributes this work to coworkers who must then interpret, correct, or redo poor-quality work. Workers in our study estimated spending 3.4 hours on average each month dealing with workslop. This additional labor required moderate cognitive effort and frequently made them feel annoyed and frustrated. Notably, these cognitive costs are greater for managers than for individual contributors. Although managers may possess the domain experience needed to evaluate such work, they often lack the day-to-day engagement required to efficiently interpret and resolve issues, amplifying the burden of workslop at higher organizational levels. Taken together, these dynamics may help explain why organizational investments in AI have not consistently yielded the productivity and collaboration gains they were expected to deliver (MIT, 2025).

Furthermore, workslop may have more pernicious effects by undermining trust between colleagues. Past work under the AI-MC framework has found that people who bring AI-generated content into interpersonal exchanges are often judged differently by their peers (Jakesch et al., 2019; Hohenstein et al., 2020; Fu et al., 2024). Similarly, we found that people who sent workslop to their colleagues were perceived as significantly less trustworthy, competent, and capable of doing good work. In addition, one-third of recipients reported less desire to work with the colleague who sent them workslop. These interpersonal costs are consistent with workslop functioning as a violation of the principle of least collaborative effort: recipients bear unexpected burden, and respond with reduced willingness to collaborate (Clark & Brennan, 1991). Workslop may also harm workplace relationships by signalling deception. When individuals are seen as using AI to generate poor-quality drafts, proposals, or presentations, they may be viewed as allowing the AI to assume an outsized degree of autonomy over work that is meant to reflect their own contributions. Because work products typically convey information about effort, competence, and investment in a collaboration (Spence, 1973), AI-generated work can disrupt this inference process, making it harder for recipients to assess how much care or capability a sender has actually invested (Galdin & Silbert, 2025; Cowgill et al., 2024; Jakesch et al., 2023). For workers who are encouraged to use AI, this creates a consequential paradox: disclosing AI use may invite negative evaluations from colleagues (Jakesch et al., 2019; Wu & Kelly, 2020), while nondisclosure - or being perceived as relying on AI - may be interpreted as deceptive and burdensome, leaving workers with no clearly safe relational path.

Limitations and Future Directions

There are several important limitations to this study that should be considered when interpreting the results. First, the study relies on participants' self-reported, recalled experiences of AI-generated workslop, and their subjective judgments of its quality. Such examples may be more intense than the average, or more commonplace interactions with workslop, in line with past work indicating that people tend to remember more emotionally intense situations (Schone et al., 2025). The lack of an objective assessment of work quality or effort could mean participants' judgment of workslop might be misattributed (Jakesch et al., 2019). However, subjective assessment helps anchor this work in inherently contextually-dependent situations. These limitations may be overcome through additional efforts to understand individuals' ranged experiences with workslop over a longer time period (i.e., employing an ecological momentary assessment), or by asking about the most recent form of workslop they encountered.

Similarly, an important limitation is that participants described situations where they believed, to the best of their knowledge, their colleague had shared AI-generated work. However, these beliefs may be faulty. We were not able to obtain ground truth about whether this work actually *did* consist of AI-generated content, or attain greater insight into what this suspected workslop looked like. Broadly speaking, past research suggests that it may be difficult to identify clear markers for AI-generated content as people often struggle to discern human-written materials from AI-written materials (Jakesch et al., 2023; Rae, 2024). For example, Yan et al. (2025) found that people often mistake authentic images for AI-generated images, a tendency that could reflect increased suspicion regarding content authenticity as AI tools become more common. Moreover, "false alarms" where people mistake human-generated content as workslop may have their own sets of consequences for interpersonal trust in the workplace. However, because of the contextual and perceptual nature of workslop, detecting it naturally would be infeasible. Future research will need to overcome this in order to experimentally investigate the consequences of workslop in controlled settings.

Another limitation revolves around asking participants to estimate the hours of work they wasted due to needing to interpret, evaluate, discuss, or revise low-quality AI-generated work they received. Numerous studies have highlighted the difficulty of obtaining accurate time estimations from individuals through survey data, without using intensive behavior-logging methodologies (i.e., such as tracking the time they spend on their work computers) (Parry et al., 2021). Recognizing these limitations, we chose to follow best practices for eliciting time estimates from survey respondents in Ernala et al. (2020) given that

it would not be feasible to obtain accurate behavioral data on participants' time spent at work without obtaining company or team-level data.

The cross-sectional nature of this survey means that causal conclusions cannot be drawn from this work. While we can investigate the associations between individuals' work experiences, individual differences, and their likelihood of producing AI-generated workslop, these findings do not speak to the factors that cause individuals to be more likely to generate AI-generated workslop. Future research employing longitudinal or experimental designs are needed to assess the causal impact of individual and environmental factors on the likelihood of producing workslop, as well as the effects of receiving workslop on individuals' perceptions of their colleagues or their own productivity. Longitudinal studies may also reveal the evolving nature of workslop production and detection.

In the future, as LLMs become increasingly more fluent, and possibly use more context, we may be tempted to believe that workslop will not continue to be consequential. Just as vision models in late 2025 have nearly eliminated the visual cues for detecting AI-generated images (Kramer et al., 2025), the same technical advancements may make communications data (i.e. documents, messages, code) indistinguishable in form and style from human-written content. However, regardless of improved model fluency or increased context, workslop senders will still, by definition, be taking a shortcut to accomplish work tasks. Recipients may perceive this as intentionally lower effort, imagination, or critical thinking. There is also implicit cultural and social context that may never be able to be integrated by such models. Therefore, workslop is likely to remain prevalent despite technical improvements.

Finally, we suggest that not all workslop is counterproductive. There are some situations where the positive payoff of workslop is greater than its consequences. From a distributed cognition perspective, sharing low-quality work may be better than sending nothing at all, or could be an implicit request for collaboration (Hutchins, 2001). Imperfect artifacts can invite refinement and innovation, and the line between workslop and low-quality yet productive work may hinge on factors such as disclosure, organizational norms, and the relationship between sender and recipient. Taking shortcuts, like relying on AI and sending workslop, may be risky (Beck et al., 2017), but given structural factors – organizational pressure, low psychological safety, etc. – could be an adaptive, rational response to dysfunctional systems. Our results show that many individuals may send workslop because they are under pressure to work quickly, yet in some cases the speed of AI-assisted work may ultimately serve broader goals.

Conclusion

Many American workers in diverse occupations report both sending and receiving AI-generated workslop: work that they believe to be low-quality and unable to meaningfully progress their projects and initiatives. We conceptualize this phenomenon within the theoretical framework of AI-mediated communication, quantify its prevalence, and describe characteristics about its form and signals used for detection. We also identify specific individual-level and organization-level predictors of sending workslop. Finally, we describe the negative impact of receiving workslop on productivity and interpersonal relationships. While it is important to recognize that not all AI-generated work is workslop, and can often be beneficial, our results offer a counterpoint by highlighting a specific case in which workplace AI use at best is ignored, or worse, harms productivity, collaboration, and trust in teammates. Understanding the nuances of how specific people make decisions about if and how they use AI in their jobs can help us identify the contexts in which AI use may be most beneficial, or most detrimental.

Code and Data Availability

All code and data may be found at
https://osf.io/xh2wn/overview?view_only=675432f416474ce592257dfd684343e3

Acknowledgements

We would like to sincerely thank Ben Lira Luttges, Dave Markowitz, and Cassie Brandes for constructive feedback on this research. A.Y.L. is supported by the Mark & Mary Stevens Stanford Interdisciplinary Graduate Fellowship and the Stanford Social Impact Labs.

Author Contributions

A.L. and A.Y.L. designed the survey; collected, curated, and analyzed the data; developed the methodology; designed visualizations; and wrote the manuscript. K.R. designed visualizations, interpreted results, and reviewed and edited the manuscript. G.K. conceptualized the research. K.N. and J.H. conceptualized the research, supervised the work, and reviewed and edited the manuscript.

Funding / Conflicts of Interest

Data collection was paid for by BetterUp, Inc. The funder provided support for the research and the conceptualization of the data collection, but did not have input into data analyses, reporting, or manuscript-writing. Authors with BetterUp affiliations are paid employees of BetterUp, Inc.; A. Lee and J. Hancock are paid consultants within BetterUp Labs at BetterUp, Inc. G. Kellerman is a paid employee of Boston Consulting Group.

References

- Anderson, N. R., & West, M. A. (1998). Measuring climate for work group innovation: development and validation of the team climate inventory. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, 19(3), 235-258. [https://doi.org/10.1002/\(SICI\)1099-1379\(199805\)19:3<235::AID-JOB837>3.0.CO;2-C](https://doi.org/10.1002/(SICI)1099-1379(199805)19:3<235::AID-JOB837>3.0.CO;2-C)
- Anthony, C., Bechky, B. A., & Fayard, A. L. (2023). “Collaborating” with AI: Taking a system view to explore the future of work. *Organization Science*, 34(5), 1672-1694. <https://doi.org/10.1287/orsc.2022.1651>
- Ansari, K., Ghasemaghaei, M., & Turel, O. (2024). Cutting corners as a coping strategy in information technology use: Unraveling the mind's dilemma. *Information & Management*, 61(8), 104057. <https://doi.org/10.1016/j.im.2024.104057>
- Aubteen Darabi, A., Nelson, D. W., & Paas, F. (2007). Learner involvement in instruction on a complex cognitive task: Application of a composite measure of performance and mental effort. *Journal of Research on Technology in Education*, 40(1), 39-48. <https://doi.org/10.1080/15391523.2007.10782495>
- Beck, J. W., Scholer, A. A., & Schmidt, A. M. (2017). Workload, risks, and goal framing as antecedents of shortcut behaviors. *Journal of Business and Psychology*, 32(4), 421-440. <https://doi.org/10.1007/s10869-016-9450-0>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1), 289-300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Blank, I. A. (2023). What are large language models supposed to model?. *Trends in Cognitive Sciences*, 27(11), 987-989.
- Box. (April 29 2025). Box and IBM Partner to Bring New Enterprise-Level AI Models to Support Content Generation and Productivity.

<https://newsroom.ibm.com/2025-04-29-box-and-ibm-partner-to-bring-new-enterprise-level-ai-models-to-support-content-generation-and-productivity>

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889-942. <https://doi.org/10.1093/qje/qjae044>
- Burgoon, J. K. (1993). Interpersonal expectations, expectancy violations, and emotional communication. *Journal of language and social psychology*, 12(1-2), 30-48. <https://doi.org/10.1177/0261927X93121003>
- Burgoon, J. K., & Walther, J. B. (1990). Nonverbal expectancies and the evaluative consequences of violations. *Human Communication Research*, 17(2), 232-265. <https://doi.org/10.1111/j.1468-2958.1990.tb00232.x>
- Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of personality and social psychology*, 42(1), 116. <https://doi.org/10.1037/0022-3514.42.1.116>
- Capel, T., & Brereton, M. (2023, April). What is human-centered about human-centered AI? A map of the research landscape. In *Proceedings of the 2023 CHI conference on human factors in computing systems* (pp. 1-23). <https://doi.org/10.1145/3544548.3580959>
- Challapally, A., Pease, C., Raskar, R., & Chari, P. (July 2025). The GenAI Divide: The state of AI in Business. *MIT NANDA*. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
- Chang, C. T., Farah, H., Gui, H., Rezaei, S. J., Bou-Khalil, C., Park, Y. J., ... & Daneshjou, R. (2025). Red teaming ChatGPT in medicine to yield real-world insights on model behavior. *npj Digital Medicine*, 8(1), 149. <https://doi.org/10.1038/s41746-025-01542-0>
- Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). How people use chatgpt (No. w34255). *National Bureau of Economic Research*. <https://doi.org/10.3386/w34255>
- Chen, S., Gao, M., Sasse, K., Hartvigsen, T., Anthony, B., Fan, L., ... & Bitterman, D. S. (2025). When helpfulness backfires: LLMs and the risk of false medical information due to sycophantic behavior. *npj Digital Medicine*, 8(1), 605. <https://doi.org/10.1038/s41746-025-02008-z>
- Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., & Jurafsky, D. (2025). Sycophantic AI decreases prosocial intentions and promotes dependence. *arXiv preprint arXiv:2510.01395*.
- Cheng, M., Lee, A.Y., Rapuano, K., Liebscher, A., Niederhoffer, K., & Hancock, J. (2025). Metaphors of AI indicate that people increasingly perceive AI as warm and human-like. *Nature Communications Psychology*, 4(8). <https://doi.org/10.1038/s44271-025-00376-6>
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127-149). American Psychological Association. <https://doi.org/10.1037/10096-006>
- Colbert, A., Yee, N., & George, G. (2016). The digital workforce and the workplace of the future. *Academy of management journal*, 59(3), 731-739. <https://doi.org/10.5465/amj.2016.4003>

- Cowgill, B., Hernandez-Lagos, P., & Wright, N. L. (2024). Does AI cheapen talk? Theory and evidence from global entrepreneurship and hiring. Theory and Evidence From Global Entrepreneurship and Hiring. *Columbia Business School Research Paper*, (4702114).
<https://dx.doi.org/10.2139/ssrn.4702114>
- Cribari-Neto, F., & Zeileis, A. (2010). Beta regression in R. *Journal of statistical software*, 34(1), 1-24.
<https://doi.org/10.18637/jss.v034.i02>
- Cuddy, A. J., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in experimental social psychology*, 40, 61-149. [https://doi.org/10.1016/S0065-2601\(07\)00002-0](https://doi.org/10.1016/S0065-2601(07)00002-0)
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science advances*, 10(28), eadn5290.
<https://doi.org/10.1126/sciadv.adn5290>
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative science quarterly*, 44(2), 350-383. <https://doi.org/10.2307/2666999>
- Ernala, S. K., Burke, M., Leavitt, A., & Ellison, N. B. (2020, April). How well do people report time spent on Facebook? An evaluation of established survey questions with recommendations. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1-14).
<https://doi.org/10.1145/3313831.3376435>
- Farquhar, S., Kossen, J., Kuhn, L., & Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017), 625-630.
- Fauville, G., Luo, M., Queiroz, A. C. M., Lee, A., Bailenson, J. N., & Hancock, J. (2023). Video-conferencing usage dynamics and nonverbal mechanisms exacerbate Zoom Fatigue, particularly for women. *Computers in Human Behavior Reports*, 10, 100271.
<https://doi.org/10.1016/j.chbr.2023.100271>
- Friedman, J. H., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33, 1-22.
<https://doi.org/10.18637/jss.v033.i01>
- Galdin A., & Silbert J. (2025). Making Talk Cheap: Generative AI and Labor Market Signaling. arXiv preprint arXiv:2511.08785
- Georganta, E., & Ulfert, A. S. (2024). Would you trust an AI team member? Team trust in human-AI teams. *Journal of occupational and organizational psychology*, 97(3), 1212-1241.
<https://doi.org/10.1111/joop.12504>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of management annals*, 14(2), 627-660.
<https://doi.org/10.5465/annals.2018.0057>
- Gosling, S. D., Rentfrow, P. J., & Swann Jr, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in personality*, 37(6), 504-528.
[https://doi.org/10.1016/S0092-6566\(03\)00046-1](https://doi.org/10.1016/S0092-6566(03)00046-1)

- Guo, Y., Sohn, J. H., Leroy, G., & Cohen, T. (2025). Are LLM-generated plain language summaries truly understandable? A large-scale crowdsourced evaluation. arXiv preprint arXiv:2505.10409
- Hamilton, D. (2025, February 3). The rise of AI resentment at work: Why employees are pushing back. *Forbes*.
<https://www.forbes.com/sites/dianehamilton/2025/02/03/the-rise-of-ai-resentment-at-work-why-employees-are-pushing-back/>
- Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication*, 25(1), 89-100. <https://doi.org/10.1093/jcmc/zmz022>
- Heinze, G., Wallisch, C., & Dunkler, D. (2018). Variable selection—a review and recommendations for the practicing statistician. *Biometrical journal*, 60(3), 431-449.
<https://doi.org/10.1002/bimj.201700067>
- Herbold, S., Hautli-Janisz, A., Heuer, U., Kikteva, Z., & Trautsch, A. (2023). A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific reports*, 13(1), 18617.
<https://doi.org/10.1038/s41598-023-45644-9>
- Hoffman, B. (2024, July 11). First came “spam.” Now, with AI, we’ve got “slop.” *The New York Times*.
<https://www.nytimes.com/2024/06/11/style/ai-search-slop.html>
- Hohenstein, J., & Jung, M. (2020). AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust. *Computers in Human Behavior*, 106, 106190.
<https://doi.org/10.1016/j.chb.2019.106190>
- Humlum, A., & Vestergaard, E. (2025). Large language models, small labor market effects (No. w33777). *National Bureau of Economic Research*. <https://doi.org/10.3386/w33777>
- Hutchins, E. (2001). Distributed Cognition. In N. J. Smelser & P. B. Baltes, *International Encyclopedia of the Social & Behavioral Sciences* (pp. 2068-2072). Pergamon.
<https://doi.org/10.1016/B0-08-043076-7/01636-3>
- Kennedy, B., Tyson, A., & Saks, E. (2023, February 15). Public awareness of artificial intelligence in everyday activities. *Pew Research Center*.
<https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/>
- Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and reliance on AI—An experimental study on the extent and costs of overreliance on AI. *Computers in Human Behavior*, 160, 108352.
<https://doi.org/10.1016/j.chb.2024.108352>
- Krumpal, I. (2013). Determinants of social desirability bias in sensitive surveys: a literature review. *Quality & quantity*, 47(4), 2025-2047. <https://doi.org/10.1007/s11135-011-9640-9>
- Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D. J., & Kamm, C. (2002, November). The character, functions, and styles of instant messaging in the workplace. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work* (pp. 11-20).
<https://doi.org/10.1145/587078.587081>

- Jakesch, M., French, M., Ma, X., Hancock, J. T., & Naaman, M. (2019, May). AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness. In *Proceedings of the 2019 CHI conference on human factors in computing systems* (pp. 1-13). <https://doi.org/10.1145/3290605.3300469>
- Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences*, 120(11), e2208839120. <https://doi.org/10.1073/pnas.2208839120>
- Jarrahi, M. H., Möhlmann, M., & Lee, M. K. (2023). Algorithmic management: The role of AI in managing workforces. *MIT Sloan Management Review*, 64(3), 1-5.
- Jonason, P. K., & O'Connor, P. J. (2017). Cutting corners at work: An individual differences perspective. *Personality and Individual Differences*, 107, 146-153. <https://doi.org/10.1016/j.paid.2016.11.045>
- Juzek, T. S., & Ward, Z. B. (2025, January). Why does ChatGPT “Delve” so much? Exploring the sources of lexical overrepresentation in Large Language Models. In *Proceedings of the 31st international conference on computational linguistics* (pp. 6397-6411).
- Kosmidis, I., & Zeileis, A. (2026). Extended-support beta regression for [0, 1] responses. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 75(1), 139-157. <https://doi.org/10.1093/jrssc/qlaf039>
- Kramer, R. S., Jones, A. L., Fitousi, D., & Tree, J. J. (2025). AI-generated images of familiar faces are indistinguishable from real photographs. *Cognitive Research: Principles and Implications*, 10(1), 70. <https://doi.org/10.1186/s41235-025-00683-w>
- Lee, A. Y., & Hancock, J. T. (2024). Social media mindsets: A new approach to understanding social media use and psychological well-being. *Journal of Computer-Mediated Communication*, 29(1), zmad048. <https://doi.org/10.1093/jcmc/zmad048>
- Lee, A. Y., Neumann, E., Zaki, J., & Hancock, J. (2025). Americans overestimate how many social media users post harmful content. *PNAS Nexus*, 4(12), pgaf310. <https://doi.org/10.1093/pnasnexus/pgaf310>
- Lenharo, M. (2023). An AI revolution is brewing in medicine. What will it look like?. *Nature*, 622(7984), 686-688. <https://doi.org/10.1038/d41586-023-03302-0>
- Lima, E., Davies, P., Kaler, J., Lovatt, F., & Green, M. (2020). Variable selection for inferential models with relatively high-dimensional data: Between method heterogeneity and covariate stability as adjuncts to robust selection. *Scientific Reports*, 10(1), 8002. <https://doi.org/10.1038/s41598-020-64829-0>
- Mabire-Yon, R. (2025). Hurdle Models in Psychology—A Practical Guide for Inflated Data. *International Journal of Psychology*, 60(3), e70042. <https://doi.org/10.1002/ijop.70042>
- MacArthur, M. (2025). Large language models and the problem of rhetorical debt. *AI & SOCIETY*, 40, 6425–6438. <https://doi.org/10.1007/s00146-025-02403-w>
- Madsen, D. Ø., & Puyt, R. W. (2025). When AI turns culture into slop. *AI & SOCIETY*. <https://doi.org/10.1007/s00146-025-02630-1>

- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., ... & Oak, S. (2025). Artificial intelligence index report 2025. arXiv preprint arXiv:2504.07139.
- Markowitz, D. M. (2024). From complexity to clarity: How AI enhances perceptions of scientists and the public's understanding of science. *PNAS nexus*, 3(9), page 387. <https://doi.org/10.1093/pnasnexus/pgae387>
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of management review*, 20(3), 709-734. <https://doi.org/10.2307/258792>
- Mendy, J., Jain, A., & Thomas, A. (2025). Artificial intelligence in the workplace—challenges, opportunities and HRM framework: a critical review and research agenda for change. *Journal of Managerial Psychology*, 40(5), 517-538.
- Meinshausen, N., & Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 72(4), 417-473. <https://doi.org/10.1111/j.1467-9868.2010.00740.x>
- Naveed, H., Khan, A. U., Qiu, S., Saqib, M., Anwar, S., Usman, M., ... & Mian, A. (2025). A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology*, 16(5), 1-72. <https://doi.org/10.1145/3744746>
- Nițu, E. M., Sălceanu, C., & Chendrean, S. T. (2023) Exploring the dynamics of neuroticism: A comprehensive analysis of two different workplace settings. *The "Black Sea" Journal of Psychology*, 14(2). <https://doi.org/10.47577/bspsychology.bsjop.v14i2.237>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192. <https://doi.org/10.1126/science.adh2586>
- Oliver, M. & Dabney, J. (2025, September 17). Ohio State University launches initiative requiring that students study AI. *CBS News*. <https://www.cbsnews.com/news/ohio-state-university-initiative-requires-students-study-artificial-intelligence/>
- Paas, F. G. (1992). Training strategies for attaining transfer of problem-solving skill in statistics: a cognitive-load approach. *Journal of educational psychology*, 84(4), 429. <https://psycnet.apa.org/doi/10.1037/0022-0663.84.4.429>
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human factors*, 52(3), 381-410. <https://doi.org/10.1177/0018720810376055>
- Parry, D. A., Davidson, B. I., Sewall, C. J., Fisher, J. T., Mieczkowski, H., & Quintana, D. S. (2021). A systematic review and meta-analysis of discrepancies between logged and self-reported digital media use. *Nature Human Behaviour*, 5(11), 1535-1547. <https://doi.org/10.1038/s41562-021-01117-5>
- Petty, R. E., Briñol, P., Loersch, C., & McCaslin, M. J. (2009). The need for cognition. In M. R. Leary & R. H. Hoyle (Eds.), *Handbook of individual differences in social behavior*, 318-329. The Guilford Press.

- Przystalski, K., Argasiński, J. K., Grabska-Gradzińska, I., & Ochab, J. (2026). Stylometry recognizes human and LLM-generated texts in short samples. *Expert Systems with Applications*, 296(B), 129001. <https://doi.org/10.1016/j.eswa.2025.129001>
- Rae, I. (2024, May). The effects of perceived AI use on content perceptions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1-14). <https://doi.org/10.1145/3613904.3642076>
- Robinson, S. L., & Bennett, R. J. (1995). A typology of deviant workplace behaviors: A multidimensional scaling study. *Academy of management journal*, 38(2), 555-572. <https://doi.org/10.2307/256693>
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of statistical software*, 48(1), 1-36. <https://doi.org/10.18637/jss.v048.i02>
- Rudko, I., & Bashirpour Bonab, A. (2025). ChatGPT is incredible (at being average). *Ethics and Information Technology*, 27(3), 36. <https://doi.org/10.1007/s10676-025-09845-2>
- Schöne, J. P., Rocklage, M. D., Parkinson, B., & Goldenberg, A. (2025). Overestimation in the aggregation of emotional intensity of social media content. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspa0000458>
- Selenko, E., Banks, S., Shoss, M., Warburton, J., & Restubog, S. L. D. (2022). Artificial intelligence and the future of work: A functional-identity perspective. *Current Directions in Psychological Science*, 31(3), 272-279. <https://doi.org/10.1177/09637214221091823>
- Sidoti, O., & McClain, C. (2025). 34% of US adults have used ChatGPT, about double the share in 2023. <https://www.pewresearch.org/short-reads/2025/06/25/34-of-us-adults-have-used-chatgpt-about-double-the-share-in-2023>
- Schaufeli, W. B., Bakker, A. B., & Salanova, M. (2006). Utrecht work engagement scale-9 (UWES-9). *Educational and Psychological Measurement*. <https://doi.org/10.1037/t05561-000>
- Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes*, 188, 104405. <https://doi.org/10.1016/j.obhdp.2025.104405>
- Shneiderman, B. (2022). Human-centered AI. *Oxford University Press*. <https://doi.org/10.1093/oso/9780192845290.001.0001>
- Skovholt, K., & Svennevig, J. (2006). Email copies in workplace interaction. *Journal of Computer-Mediated Communication*, 12(1), 42-65. <https://doi.org/10.1111/j.1083-6101.2006.00314.x>
- Smith, M. E., Hancock, J. T., Reynolds, L., & Birnholtz, J. (2014). Everyday deception or a few prolific liars? The prevalence of lies in text messaging. *Computers in Human Behavior*, 41, 220-227. <http://dx.doi.org/10.1016/j.chb.2014.05.032>
- Spatola, N. (2024). The efficiency-accountability tradeoff in AI integration: Effects on human performance and over-reliance. *Computers in Human Behavior: Artificial Humans*, 2(2), 100099. <https://doi.org/10.1016/j.chbah.2024.100099>

- Spector, P. E., & Jex, S. M. (1998). Development of four self-report measures of job stressors and strain: interpersonal conflict at work scale, organizational constraints scale, quantitative workload inventory, and physical symptoms inventory. *Journal of occupational health psychology*, 3(4), 356. <https://doi.org/10.1037//1076-8998.3.4.356>
- Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355-374. <https://doi.org/10.2307/1882010>
- Spronk, I., Korevaar, J. C., Poos, R., Davids, R., Hilderink, H., Schellevis, F. G., ... & Nielen, M. M. (2019). Calculating incidence rates and prevalence proportions: not as simple as it seems. *BMC public health*, 19(1), 512. <https://doi.org/10.1186/s12889-019-6820-3>
- Stewart, A. (2025). Microsoft pushes staff to use internal AI tools more, and may consider this in reviews. 'Using AI is no longer optional.' *Business Insider*. <https://www.businessinsider.com/microsoft-internal-memo-using-ai-no-longer-optional-github-co-pilot-2025-6>
- Szulewski, A., Gegenfurtner, A., Howes, D. W., Sivilotti, M. L., & Van Merriënboer, J. J. (2017). Measuring physician cognitive load: validity evidence for a physiologic and a psychometric tool. *Advances in Health Sciences Education*, 22(4), 951-968. <https://doi.org/10.1007/s10459-016-9725-2>
- Tay, J. K., Narasimhan, B., & Hastie, T. (2023). Elastic net regularization paths for all generalized linear models. *Journal of statistical software*, 106, 1-31. <https://doi.org/10.18637/jss.v106.i01>
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1), 267-288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology*, 54(6), 1063. <https://doi.org/10.1037//0022-3514.54.6.1063>
- Wu, S., Liu, Y., Ruan, M., Chen, S., & Xie, X. Y. (2025). Human-generative AI collaboration enhances task performance but undermines human's intrinsic motivation. *Scientific Reports*, 15(1), 15105. <https://doi.org/10.1038/s41598-025-98385-2>
- Wu, Y., & Kelly, R. M. (2020, December). Online dating meets artificial intelligence: How the perception of algorithmically generated profile text impacts attractiveness and trust. In *Proceedings of the 32nd Australian conference on human-computer interaction* (pp. 444-453). <https://doi.org/10.1145/3441000.3441074>
- Yan, H. Y., Moore, R. C., Tu, F., & Hancock, J. T. (2025). Detecting Synthetic, Doubting Authentic: AI Attribution Bias for Political Imagery. https://doi.org/10.31219/osf.io/w6bzx_v1

Appendices

Appendix 1: Description of Pilot Study

A pilot study administered in August 2025 asked full-time American workers from the survey platform CloudResearch's Connect platform ($N = 150$) to describe their thoughts about workplace AI use. The survey items included can be found below. Upon consenting, participants were asked to complete sociodemographic items about themselves, questions about their work environment, and questions about their attitudes towards AI and use of AI.

If they indicated that they shared or sent AI-generated work with colleagues, they were asked "How much of the work that you send to colleagues is AI-generated content that you think was actually unhelpful, low effort, or low quality?" (from 0% to 100%, in 10% intervals). Next, they answered if they had received work: "In your job, have you ever received work content that you believe is AI-generated that looks like it completes a task at work, but is actually unhelpful, low quality, and/or seems like the sender didn't put in enough effort? It could have appeared in many different forms, including documents, slide decks, emails, and code. It may have looked good, but was overly long, hard to read, fancy, or sounded different than normal." (1 = yes, 0 = no). They also estimated the percentage of work they received from colleagues that fit this description (0% to 100%, in 10% intervals).

Participants that said they had received any work were asked to recall an "actual example of this that you've experienced in the last month." They answered several open-ended questions providing specific details about who sent them this work, what the work was for, what made them believe this work was AI-generated, and if this had any impact on their own work. Finally, they indicated their "best estimate of how many additional hours this kind of work has taken for you to deal with in the last month, beyond what it would have taken if the work was higher quality."

At the conclusion of the pilot, they were debriefed and compensated for their time: "This is part of an ongoing research project to understand the prevalence of AI-generated content in the workplace, and to better understand how people feel about the nature of work as AI tools become more prevalent. In particular, we are interested in understanding AI work slop - which is AI-generated work that may incur costs when sent to others. You can opt to remove your data if you choose by contacting the research team on CloudResearch."

Appendix 2: Main study survey

Below are the primary instructions and items presented to participants in the main study survey.

Instructions

Thank you for taking part in this study. You will be asked to answer some questions about your personal experiences at work. There will be both multiple choice questions as well as open-ended questions.

There are several attention check questions. This survey should take about 15 minutes to complete.

Please do not use any AI, Large Language Model (LLM), or automated system to complete this survey. We are interested in authentic human beliefs and behavior, even if the AI helps you write better, complete the survey faster, etc. We are monitoring a number of signals to identify AI use, and we may reject anyone believed to have completed this survey using AI.

Demographics & Work Experience

First, we have some questions about your job.

Which of the following categories best describes your role at work?

- Individual contributor

- Front-line manager (a manager of individual contributors)
- Manager of managers

What is your job title?

Which industry best describes the one you currently work in? [From the United States Census Bureau]

[\[link\]](#)

- Agriculture, Forestry, and Fishing
- Arts and Entertainment
- Automotive
- Construction
- Healthcare
- Insurance
- Energy
- Financial Services
- Government and Public Sector
- Leisure & Hospitality
- Manufacturing
- Pharma/Biotech
- Professional and Business Services
- Real Estate
- Retail
- Technology
- Telecommunications
- Transportation
- Utilities

In the last month, about how many different people have you worked with consistently? If you have worked alone over the last month, please enter 0. [numeric response]

How many total years of full-time work experience do you have? [numeric response]

What is your total individual income in U.S. dollars for the past 12 months from all wages, salary, commissions, or tips?

- No income
- Less than \$25,000
- \$25,000-49,999
- \$50,000-74,999
- \$75,000-99,999
- \$100,000-124,999
- \$125,000-149,999
- \$150,000-174,999
- \$175,000-199,999
- \$200,000-224,999
- \$225,000-249,999
- \$250,000-274,999
- \$275,000-\$299,999

- More than \$300,000
- Prefer not to say

Individual Differences.

Ten-Item Personality Measure (Gosling et al., 2003)

Here are a number of personality traits that may or may not apply to you. Please indicate the extent to which you agree or disagree with that statement. You should rate the extent to which the pair of traits applies to you, even if one characteristic applies more strongly than the other. I see myself as... (7-point scale, 1 = strongly disagree, 7 = strongly agree)

- Extraverted, enthusiastic (extraversion)
- Critical, quarrelsome (agreeableness, reverse-coded)
- Dependable, self-disciplined (conscientiousness)
- Anxious, easily upset (neuroticism)
- Open to new experiences, complex (openness to experiences)
- Reserved, quiet (extraversion, reverse-coded)
- Sympathetic, warm (agreeableness)
- Disorganized, careless (conscientiousness, reverse-coded)
- Calm, emotionally stable (neuroticism, reverse-coded)
- Conventional, uncreative (openness to experiences, reverse-coded)

Need for Cognition Scale (Cacioppo & Petty, 1982)

For each sentence below, please select how uncharacteristic or characteristic (5-point scale) this is for you personally. (5-point scale, 1 = extremely uncharacteristic, 5 = extremely characteristic)

- It's enough for me to get the job done; I don't care how or why it works. (reverse-coded)
- I would prefer complex to simple problems.
- I would rather do something that requires little thought than something that is sure to challenge my thinking abilities. (reverse-coded)
- I like to have the responsibility of handling a situation that requires a lot of thinking.
- I prefer to think about small, daily projects to long-term ones. (reverse-coded)

Quantitative Workload Inventory (Spector & Jex, 1998)

Please rate the following questions about your job. (5-point scale, 1 = never, 5 = always)

- How often does your job require you to work very fast?
- How often does your job require you to work very hard?
- How often does your job leave you with little time to get things done?
- How often is there a great deal to be done?
- How often do you have to do more work than you can do well?

Pride in Work (Schaufeli et al., 2006) (5-point scale, 1 = Never, 5 = Always)

- How often do you find your work to be full of meaning and purpose?
- How often do you feel proud of the work that you do?
- How often do you feel it is important for you to do a good job at work?

Perceptions of AI.

Next, we'd like to know about your experience with Artificial Intelligence (AI).

Self-reported knowledge of AI. (Cheng et al., 2025)

How much do you know about AI? (5-point scale, 1 = not knowledgeable at all, 5 = extremely knowledgeable)

Trust in AI (adapted from Jakesch et al., 2019 and Mayer, 1995)

How much do you agree or disagree with the following statements? (5-point scale, 1 = strongly disagree, 5 = strongly agree)

- I believe that AI will produce output that is accurate. (ability)
- I believe that AI will have a positive effect on most people. (benevolence)
- I believe that AI will be safe and secure to use. (integrity)

AI Mindsets: Valence and agency towards AI (adapted from Lee & Hancock, 2024)

Below are some statements about AI. We want you to think about the ways that you personally think about and use AI. Please rate the extent to which you agree or disagree with each statement. (5-point scale, 1 = strongly disagree, 5 = strongly agree)

- Using AI is fun and enjoyable for me. (valence)
- Using AI facilitates my learning and growth. (valence)
- Using AI is harmful to me. (valence, reverse-coded)
- I'm good at using AI to do things that are important to me. (agency)
- Using AI to do things makes me feel dependent on it. (agency, reverse-coded)
- I feel like using AI makes me less capable. (agency, reverse-coded)

Experiences using AI

How often do you use AI in your work? (8-point scale, 1 = Never, 8 = More than several times a day)

How often do you **send or share** work with colleagues after using AI to help produce it? "Work" may include documents, emails, slide decks, code, agendas, newsletters, budgets, multimedia content, etc. (8-point scale, 1 = Never, 8 = More than several times a day)

How much of the AI-generated work that you send to colleagues do you think is actually unhelpful, low effort, or low quality? (11-point scale, 0 = None at all - 0%, 10 = All of it - 100%)

How much of the AI-generated work that you send to colleagues do you think is helpful and high quality? (11-point scale, 0 = None at all - 0%, 10 = All of it - 100%)

You said that about [% of my work considered workslop] of the AI-generated work that you send to colleagues is actually unhelpful, low effort, or low quality. In general, before you send this kind of work, how much do you think each of the following applies to you/the task? (5-point scale, 1 = Not at all, 5 = Definitely)

- The task is very urgent
- The task is very important
- The task is one of many things I'm responsible for completing
- The task is very ambiguous
- I am under great pressure to perform at work
- I am looking to buy myself more time
- I lack the expertise to complete the task myself
- I need a thought partner to complete the task

- I am encouraged to use AI no matter what

My supervisor or organization encourages us to use AI. (5-point scale, 1 = Strongly disagree, 5 = Strongly agree)

Experiences with Workstop

In your job, have you received work content in the last month that you believe is AI-generated that looks like it completes a task at work, but is actually unhelpful, low quality, and/or seems like the sender didn't put in enough effort? It could have appeared in many different forms, including documents, slide decks, emails, and code. It may have looked good, but was overly long, hard to read, fancy, or sounded different than normal. (1 = Yes, 0 = No)

In the last month, in general, how much of the work you received from colleagues fits this description? (11-point scale, 0 = None at all - 0%, 10 = All of it - 100%)

Earlier in the survey, you said that you regularly work with around [number of regular colleagues] people. In the last month, how many of these people have sent you work that fits this description?

In the last month, in general, how much of the work you received from colleagues is AI-generated work that is helpful and high quality? (11-point scale, 0 = None at all - 0%, 10 = All of it - 100%)

You indicated that you've **received** AI-generated content that looks like it completes a task at work, but is actually unhelpful or low-quality. **Please think about an actual example of this that you've experienced in the last month.** We're interested in learning more about what this was like for you. First, who sent you this work?

- My teammate (1)
- My direct report (2)
- My supervisor (3)
- A peer on another team (4)
- Someone higher-up on another team (5)
- Someone lower level on another team (6)
- My client (7)

Please describe what they sent you and what it was for. Please write 2 to 4 sentences. Do not use AI to complete this question.

What made you believe this work was AI-generated? Please write 1 to 3 sentences. Do not use AI to complete this question.

Considering that your colleague's work may have been AI-generated, what impact did this have your own work? For example, did it save you time, or waste it? Did you finish your work sooner, or later? Was it easier, or harder, to make more progress on the task or project? Please write 1 to 3 sentences. Do not use AI to complete this question.

Consequences: Time Cost

Overall, how much time did it take for you to deal with this work? This includes time you spent reading or processing it, evaluating it, discussing it with others, revising it, or redoing it. Please estimate as precisely as you can. Enter a 0 in each box that you'd like to leave empty. (based on Ernala et al., 2020)

- Minutes
- Hours

What percentage of the time it took to deal with this work did you spend on each of the following? Please enter your answer using only numbers. The sum of these numbers should be 100. If you did not do a particular step, please enter 0.

- Reading or processing the work (1)
- Evaluating the work (2)
- Discussing the work with others (i.e., the person who sent it to you, a different colleague, a manager) (3)
- Revising or fixing the work (4)
- Redoing the work (5)

Consequences: Emotional Cost

To what extent did you feel each of the following emotions when you received this work? (adapted version of the Positive and Negative Affect Scale (PANAS); Watson et al., 1988) (5-point scale, 1 = Not at all, 5 = Extremely)

- Frustrated (negative)
- Annoyed (negative)
- Angry (negative)
- Confused (negative)
- Content (positive)
- Happy (positive)
- Interested (positive)
- Impressed (positive)
- Offended (negative)

Consequences: Cognitive Cost

Thinking about this work you received, how much mental effort did it require to understand, revise, or integrate it into your work? (Paas et al., 1992) (9-point scale, 1 = Very low mental effort, 9 = Very high mental effort)

Consequences: Interpersonal Cost

How did receiving this work change your perception of your colleague? I saw them as... (based on the Stereotype Content Model; Cuddy et al., 2008) (5-point scale, 3 = No change)

- Less pleasant ↔ More pleasant (warmth)
- Less capable ↔ More capable (competence)
- Less intelligent ↔ More intelligent (competence)
- Less friendly ↔ More friendly (warmth)
- Less reliable ↔ More reliable (competence)
- Less diligent ↔ More diligent (competence)
- Less creative ↔ More creative

After receiving this work from your colleague, how much more or less would you want to work with them again in the future? (5-point scale, 1 = I would really like to avoid working with them again, 3 = No change in how much I want to work with them, 5 = I would really like to work with them again)

To what extent has receiving this work from your colleague changed how you view them in the following ways? (Mayer et al., 1995) (5-point scale, 1 = Much less true of them, 5 = Much more true of them)

- They are very capable of performing their job.
- They have knowledge about the work that needs to be done.
- They are well-qualified.
- They really look out for what is important to me.
- They have a strong sense of justice.
- They will go out of their way to help me.
- Their behavior is very consistent.
- They have good values.
- Sound principles seem to guide their behavior.

Did you do any of the following to address your belief that the work was AI-generated?

- Spoke to the colleague directly (1)
- Spoke with the colleague's manager (2)
- Spoke with my manager (3)
- Spoke with a peer/teammate (4)
- Spoke to HR (5)
- Documented the occurrence (6)
- Did nothing and moved on (7)
- I did something else

Thank you for sharing your experience. We're interested in understanding how much additional time it may have taken for you to engage with this kind of work over the last month. Please provide your best estimate of how many additional hours this kind of work has taken for you to deal with in the last month, beyond what it would have taken if the work was higher quality. If it did not affect your work or it saved you time, please enter 0.

Workplace Culture and Norms

Edmonson Psychological Safety (Edmondson, 1999) (7-point scale, 1 = Very inaccurate, 7 = Very accurate)

- If you make a mistake on this team, it is often held against you. (reverse)
- Members of this team are able to bring up problems and tough issues.
- In my team, it is the norm to see mistakes as learning opportunities.
- It is difficult to ask other members of this team for help. (reverse)
- Working with members of this team, my unique skills and talents are valued and utilized.

Team Climate Inventory (Vision) (Anderson & West, 1998) (5-point scale, 1 = Strongly disagree, 5 = Strongly agree)

- I am clear what our team's objectives are.
- I believe these objectives are worthwhile to the organization.
- I believe the members of my team are committed to these objectives.

Team Climate Inventory (Task Orientation) (Anderson & West, 1998) (5-point scale, 1 = Strongly disagree, 5 = Strongly agree)

- My colleagues and I monitor each other so as to maintain a high standard of work.
- The team has clear criteria which members try to meet in order to achieve excellence as a team.
- Team members are prepared to question the basis of what the team is doing.
- There is real concern among team members that the team should achieve the highest standard of performance.
- The team critically appraises potential weaknesses in what it is doing to achieve the best possible outcome.

Appendix 3: Number of participants by industry

Table A1: Number of participants by industry

Field	Number	Percent
Technology	175	18.2%
Government and Public Sector	172	17.9%
Healthcare	156	16.2%
Professional and Business Services	109	11.3%
Financial Services	89	9.3%
Manufacturing	47	4.9%
Retail	41	4.3%
Arts and Entertainment	39	4.1%
Insurance	27	2.8%
Transportation	18	1.9%
Construction	17	1.8%
Pharma/Biotech	17	1.8%
Telecommunications	16	1.7%
Leisure and Hospitality	12	1.2%
Real Estate	9	0.9%
Agriculture, Forestry, and Fishing	8	0.8%
Energy	4	0.4%
Automotive	3	0.3%
Utilities	3	0.3%

The list of industries was adapted from the 2020 United States Census occupational categories.

Appendix 4: Descriptive data of workplace AI use and reported workslop production

Table A2. Cross-table of workplace AI usage and workslop production

	Sends workslop	Does not send workslop
Shares AI content often	9.6% (n = 92) (i.e. prolific producers of workslop)	26.6% (n = 256)
Does not share AI content often	11.6% (n = 112)	52.2% (n = 502)

People who “send workslop” indicates that participants reported that at least some of the AI-generated work they shared would be considered workslop (i.e., 10% or greater). Participants that “do not send workslop” include those who do not use AI at work, do not share AI-generated content at work, or do not share any AI-generated content that they would consider to be workslop (i.e., 0%). People who “share AI content often” indicates people who said they share AI-generated content at work at least once a week, or more frequently. People who “do not share AI content often” includes those who do not use AI at work, never not share AI-generated content at work, or do so infrequently (i.e., once a month).

Appendix 5: Self-reported motivations for sending workslop

Analysis of motivations at the item-level were conducted with the subset of participants ($n = 510$) that indicated they received workslop and completed the motivation items. The reasons why people said they shared workslop (on a 1 = not at all to 5 = definitely scale) were because 1) the task was one of many they were responsible for completing ($M = 3.35$, $SD = 1.04$), 2) the task was very urgent ($M = 2.90$, $SD = 1.11$), 3) the task was very important ($M = 2.88$, $SD = 1.13$), 4) they felt they were under great pressure to perform at work ($M = 2.86$, $SD = 1.17$), and 5) they were looking to buy themselves more time ($M = 2.85$, $SD = 1.21$).

Reason	Mean (SD)	Somewhat, A lot, or Definitely	A lot or Definitely
The task is one of many things I'm responsible for completing	3.35 (1.04)	80.0%	47.5%
The task is very urgent	2.9 (1.11)	65.0%	31.7%
The task is very important	2.88 (1.13)	63.0%	31.3%
I am under great pressure to perform at work	2.86 (1.17)	61.0%	31.3%
I am looking to buy myself more time	2.85 (1.21)	60.6%	32.7%
The task is very ambiguous	2.71 (1.09)	58.0%	24.4%
I need a thought partner to complete the task	2.63 (1.14)	54.3%	25.0%
I lack the expertise to complete the task myself	2.07 (1.1)	33.1%	12.5%
I am encouraged to use AI no matter what	1.98 (1.14)	30.1%	12.9%

An exploratory factor analysis using a geomin oblique rotation in the *lavaan* package in R (Rosseel, 2012) indicated that reported motivations for sending workslop were mostly organized around two latent variables ($\chi^2 = 75.55$, $p < .001$; $CFI = .94$, $TLI = .88$; $AIC = 13,125.64$; $RMSEA = .08$, $SRMR = .04$). As shown in A3, the two-factor model indicated that people were driven by two primary kinds of motivations. The first factor, which we interpreted as “work pressure”, included participants saying they generated workslop because the task was urgent and important, they were under pressure to do many things at work, and they wanted to buy time. The second factor, which we interpreted as “wanting external support”, included people saying they generated workslop because the task was ambiguous, they did not feel they had expertise, and they wanted a thought partner to help them. The final item on being encouraged to use AI did not load on any factor.

Table A3. Results of exploratory factor analysis for self-reported motivations for sharing AI-generated workslop

	Item Means & SDs	Factor 1	Factor 2
The task is one of many things I'm responsible for completing	3.35 (1.04)	.59	
The task is very urgent	2.90 (1.11)	.72	
I am under great pressure to perform at work	2.86 (1.17)	.69	
The task is very important	2.88 (1.13)	.56	
I am looking to buy myself more time	2.85 (1.21)	.48	
The task is very ambiguous	2.71 (1.09)		.53
I need a thought partner to complete the task	2.63 (1.14)		.38
I lack the expertise to complete the task myself	2.07 (1.10)		.70
I am encouraged to use AI no matter what	1.98 (1.14)	–	–

Appendix 6: Analysis of predictors of sending workslop

Given the large number of predictors relative to our sample size and the cross-sectional nature of our data, we conducted bootstrap LASSO analyses as a complementary variable selection approach to identify robust predictors in the first portion of the hurdle model (Heinze, Wallisch, and Dunkler, 2018). LASSO (Least Absolute Shrinkage and Selection Operator) regression applies an L1 penalty that shrinks coefficients toward zero, with some coefficients shrinking to exactly zero and thus being excluded from the model (Tibshirani, 1996). LASSO is data-driven and identifies predictors based on their unique

contribution to prediction accuracy while accounting for correlations among predictors. By combining LASSO with bootstrap resampling, we assessed the stability of predictor selection across repeated samples.

For the binary outcome of produced workslop or not, we generated 500 bootstrap resamples by sampling responses with replacement from the original dataset. For each resample, we fit a LASSO logistic regression model with 10-fold cross-validation using the glmnet package in R (Friedman, Hastie, and Tibshirani, 2010; Tay, Narasimhan, and Hastie, 2023). We selected the regularization parameter (λ) that minimized cross-validation error and extracted coefficient estimates at this optimal λ (Friedman, Hastie, and Tibshirani, 2010). All predictors were standardized prior to analysis.

For each predictor, we calculated selection frequency (proportion of 500 bootstrap samples in which the predictor had a non-zero coefficient), mean coefficient value across samples where selected, and 95% bootstrap confidence intervals of coefficients. We defined "robust predictors" as those selected in $\geq 95\%$ of bootstrap samples, indicating the predictor consistently emerged as important across virtually all resamples regardless of which specific observations were included (Meinshausen and Bühlmann, 2010).

The unregularized logistic regression and the bootstrapped LASSO logistic regression results largely align. There are minor differences which do not influence our conclusions. Specifically, extraversion, quantitative workload index, AI valence, and trust in AI are significant predictors in the regularized model but not the unregularized model. These are marginally significant in the unregularized model and thus not reported in the main text.

Table A4. Dual model of predictors of sending any amount of workslop

term	Unregularized					LASSO Regularized			
	β	Std. Error	p-value	99% CI	LRT	β	Std. Error	Selection %	99% CI
Intercept	0.20	0.13	0.116	[-0.13 - 0.54]		0.18	0.13	100.0%	[-0.14 - 0.53]
Sociodemographics					$\chi^2(3) = 14.5^{**}$				
Gender [Woman]	-0.37	0.15	0.015	[-0.77 - 0.02]		-0.32	0.15	97.0%	[-0.75 - 0]
Age	-0.03	0.07	0.636	[-0.22 - 0.16]		-0.05	0.06	77.0%	[-0.23 - 0.13]
Education	0.10	0.07	0.169	[-0.09 - 0.28]		0.09	0.06	86.8%	[-0.05 - 0.27]
Psychological Traits					$\chi^2(7) = 32.9^{***}$				
Extraversion	0.16	0.08	0.045	[-0.04 - 0.36]		0.14	0.07	97.0%	[0 - 0.33]
Openness	0.08	0.08	0.325	[-0.13 - 0.29]		0.07	0.07	73.8%	[-0.1 - 0.27]
Neuroticism	0.25	0.09	0.006	[0.02 - 0.48]		0.22	0.08	99.6%	[0.01 - 0.45]
Conscientiousness	-0.13	0.08	0.103	[-0.35 - 0.08]		-0.13	0.08	94.2%	[-0.39 - 0.02]
Agreeableness	-0.01	0.08	0.860	[-0.23 - 0.2]		-0.04	0.07	66.0%	[-0.22 - 0.16]
Need for Cognition	-0.03	0.08	0.682	[-0.24 - 0.17]		-0.04	0.07	67.6%	[-0.19 - 0.15]
Work Pride	-0.11	0.09	0.216	[-0.35 - 0.12]		-0.10	0.07	80.8%	[-0.32 - 0.09]
Workplace Characteristics					$\chi^2(3) = 71.5^{***}$				
Work Arrangement	-0.10	0.07	0.195	[-0.29 - 0.09]		-0.09	0.06	81.2%	[-0.29 - 0.03]
Workplace AI Encouragement	0.47	0.08	0.000	[0.27 - 0.68]		0.45	0.08	100.0%	[0.23 - 0.64]
Manager [True]	0.21	0.15	0.173	[-0.19 - 0.6]		0.20	0.13	90.4%	[-0.15 - 0.56]
Psychological Experience at Work					$\chi^2(4) = 19.8^{***}$				
Psychological Safety	-0.22	0.10	0.025	[-0.49 - 0.03]		-0.22	0.09	100.0%	[-0.45 - -0.02]
Team Climate (Shared Vision)	-0.08	0.10	0.436	[-0.35 - 0.19]		-0.09	0.09	68.0%	[-0.31 - 0.17]
Team Climate (Task Orientation)	0.18	0.09	0.047	[-0.05 - 0.42]		0.15	0.08	95.0%	[-0.01 - 0.39]
Quantitative Workload Index	0.13	0.07	0.092	[-0.07 - 0.32]		0.12	0.07	95.2%	[0 - 0.31]
Attitudes toward AI					$\chi^2(3) = 25.7^{***}$				
AI Agency	-0.32	0.10	0.001	[-0.58 - -0.07]		-0.25	0.10	98.2%	[-0.52 - -0.02]
AI Valence	0.23	0.13	0.072	[-0.1 - 0.57]		0.20	0.11	87.8%	[0 - 0.47]
Trust in AI	0.29	0.12	0.016	[-0.02 - 0.59]		0.28	0.11	99.8%	[0.03 - 0.53]

Signif. codes: '***' = <0.001, '**' = <0.01, '*' = <0.05

Similarly, as an added robustness test of the second portion of the hurdle model, we augment the results of the Poisson model with those of a Beta regression model. This model is appropriate in this context because by nature of the hurdle model, the lower bound of 0 is excluded and the data generating process has a ceiling at 100%, which can be scaled to 1. The very few individuals ($n = 6$) who report that 100% of their work is workslop are excluded from this analysis. Recent methods have explored alternatives to such augmented or censored approaches (Kosmidis & Zeileis, 2025). We use the betareg in R to fit the Beta regression (Cribari-Neto & Zeileis, 2010).

The results between the Poisson and Beta regression models largely align. The Poisson regression identifies age as a significant predictor, whereas the Beta regression does not.

Table A5. Dual model of predictors of the magnitude of workslop sent

term	Zero-truncated Poisson				Beta Regression			
	β	Std. Error	p-value	99% CI	β	Std. Error	p-value	99% CI
Intercept	0.75	0.06	0.000	[0.6 - 0.9]	-1.04	0.07	0.000	[-1.22 - -0.87]
Sociodemographics								
Gender [Woman]	0.14	0.07	0.041	[-0.04 - 0.31]	0.12	0.08	0.147	[-0.09 - 0.32]
Age	-0.10	0.03	0.004	[-0.18 - -0.01]	-0.08	0.04	0.035	[-0.18 - 0.02]
Education	-0.06	0.03	0.054	[-0.14 - 0.02]	-0.02	0.04	0.532	[-0.12 - 0.07]
Psychological Traits								
Extraversion	-0.02	0.04	0.564	[-0.11 - 0.07]	-0.01	0.04	0.792	[-0.12 - 0.1]
Openness	-0.01	0.04	0.790	[-0.11 - 0.09]	-0.01	0.04	0.757	[-0.12 - 0.1]
Neuroticism	-0.17	0.04	0.000	[-0.27 - -0.07]	-0.12	0.05	0.009	[-0.24 - 0]
Conscientiousness	-0.19	0.03	0.000	[-0.27 - -0.11]	-0.19	0.04	0.000	[-0.29 - -0.08]
Agreeableness	-0.04	0.03	0.293	[-0.13 - 0.05]	-0.06	0.04	0.148	[-0.17 - 0.05]
Need for Cognition	-0.02	0.04	0.672	[-0.12 - 0.08]	0.00	0.04	0.997	[-0.11 - 0.11]
Work Pride	0.00	0.04	0.914	[-0.1 - 0.11]	0.01	0.05	0.805	[-0.11 - 0.14]
Workplace Characteristics								
Work Arrangement	-0.02	0.03	0.611	[-0.1 - 0.07]	-0.01	0.04	0.881	[-0.11 - 0.1]
Workplace AI Encouragement	-0.08	0.04	0.034	[-0.17 - 0.02]	-0.03	0.04	0.562	[-0.14 - 0.09]
Manager [True]	0.12	0.06	0.068	[-0.05 - 0.28]	0.10	0.08	0.196	[-0.1 - 0.3]
Psychological Experience at Work								
Psychological Safety	-0.12	0.04	0.003	[-0.23 - -0.02]	-0.12	0.05	0.018	[-0.25 - 0.01]
Team Climate (Shared Vision)	0.10	0.05	0.031	[-0.02 - 0.22]	0.13	0.06	0.016	[-0.01 - 0.28]
Team Climate (Task Orientation)	0.04	0.04	0.370	[-0.07 - 0.15]	0.02	0.05	0.763	[-0.11 - 0.14]
Quantitative Workload Index	-0.01	0.03	0.830	[-0.09 - 0.08]	0.02	0.04	0.693	[-0.09 - 0.12]
Attitudes toward AI								
AI Agency	-0.13	0.04	0.005	[-0.24 - -0.01]	-0.14	0.05	0.012	[-0.27 - 0]
AI Valence	-0.37	0.06	0.000	[-0.52 - -0.22]	-0.41	0.07	0.000	[-0.6 - -0.22]
Trust in AI	0.14	0.05	0.005	[0.01 - 0.28]	0.22	0.06	0.001	[0.05 - 0.39]

Appendix 7: Qualitative codebook for types of workslop

We asked participants who indicated they received workslop to describe a concrete incident that occurred in the last month using an open-ended question. Qualitative data were coded by two members of the research team using thematic content analysis guided by grounded theory principles. *Types of workslop* were coded as one of nine categories, such as financial documents, emails, and code. The first coder manually evaluated all responses, developed a detailed codebook with examples (see below), which was given to the second coder, who evaluated 20% of the responses. Discrepancies were discussed

iteratively between the coders until resolution was reached. Inter-rater reliability was assessed with Cohen's kappa (.86).

Table A6. Qualitative codebook for assessing the prevalence of types of reported workslop and prevalence

	Code	Description	Examples	Prevalence (%)
1	Financial documents	Written work relating to processing, documenting, or planning financial materials such as receipts, budgets, and spreadsheets of financial information. This <i>does not</i> include presentations about financial information or graphical renderings of finance-related data.	“Breakdown of clients transactions that didn't match up with our records. Had to go back and check up with client to double check for completeness.” “I work with creating govt contracts. I often get project budget forecasts suggesting how much a project might cost to complete. These are often ai generated and not necessarily accurate estimates.”	1.1%
2	Emails, instant messages, and interpersonal communications	Emails, work-related direct messages, Slacks, and interpersonal communications that are on a person-to-person basis, including summaries of meeting notes Potentially one to many	“It was an email about work schedules. It had Sunday instead of Saturday with I knowing we don't work on Sunday. It was a minor inconvenience. I had to call my teammate up just to clarify between Saturday and Sunday. “It was an email talking about a new doctor for the clinic. It had broken links and misinformation about our clinics.	33.4%
3	Internal documents	Written documents or materials intended to be shared within or across teams inside of an organization, such as a proposal, project report, meeting notes, and schedules	“A brochure draft with AI written sections. The AI pieces were overly long and not totally coherent. They also made up quotes.” “I was sent a regulatory report from a close peer. We were working together on research. The information didn't seem accurate so I fact checked the citations and it was bad information.” “It was a news story with localized information on housing price trends that was prepared with AI assistance. Some of the data clearly seemed inaccurate, and an editor had to seek correction	18.5%

			before it could be published.”	
4	Work that is shared externally to the organization, or intended to be shared externally	Written work that is shared out beyond an interpersonal context and intended to be shared to a public-facing context, such as memos going to an organization, written chapters, research articles, blogs, and brochures	“A colleague sent in an article written for the companies health blog. It was a fully AI generated article with 1 photo in the email to add to the post.” “It was a lesson plan. I believe they used ChatGPT to streamline the lesson planning process, basically providing the AI with the learning objectives and letting the AI plug all that data into a plan template.” “It was a report on a customer asking for a business loan. The analysis was off by a little bit on if we should actually do the loan. The explanation was too wordy and didn't sound like a human wrote it.”	26.9%
5	Code	Code, code comments, script, comments in an online database, or the development of a technical feature	“Code snippets that looked like they weren’t written by someone with this colleague’s seniority. Come to think of it, I don’t think anyone in their team could write bad code like that.” “I was sent some code to review that AI generated. It was supposed to be used to solve a problem but actually made the code more complicated than necessary.”	5.3%
6	Graphics and images	A visual, image, or graphic that is <i>not</i> data-based	“It was a mock-up for a potential ad. It had blurring and artifact issues.”	3.1%
7	Presentations	A slide deck or presentation, which goes beyond simple written work and includes visualizations or expectation of a verbal presentation	“A powerpoint on the current project with the formatting incorrect. Looked like glitches through out” “I received an AI generated presentation that had too many gimmicks and catch-phrases. This did not match the style of the intended audience. I sent the work back to be reworked.”	5.9%
8	Data analysis	Statistical analysis, visualization, or representation of quantitative data	“I received an analysis of a data set that looked extremely like it was coming from a machine because it had very generic answers and did not did into the details. we needed more granularity for analyses” “It was data on sales history but it was pretty pointless. It really didn't have anything to do with our products, so it was unhelpful, I don't	5.3%

			think they were paying attention to what it spit out and just sent it to be 'done' with their task.”	
--	--	--	--	--

Appendix 8: Senders of workslop

Table A7. Proportion of relationships between sender and recipient.

Workslop Incident Sender	N	Percent
My teammate	147	40.6%
A peer on another team	68	18.8%
My direct report	68	18.8%
My supervisor	31	8.6%
Someone higher-up on another team	22	6.1%
My client	13	3.6%
Someone lower level on another team	13	3.6%

Note. These data are from the subset of individuals who described a specific instance of receiving workslop ($n = 362$).

Appendix 9: Qualitative codebook for signals of workslop

We asked participants who indicated they received workslop to describe a concrete incident that occurred in the last month, using a combination of survey questions and open-ended responses. Qualitative data were coded by two members of the research team using thematic content analysis guided by grounded theory principles. *Perceived signals of workslop* were coded where participants' responses were binarized if they referenced each of the following five reasons: stylistic or linguistic cues, work not sounding like their colleague, missing context or containing errors, the speed of work completion, and disclosure of AI use. These five codes were not mutually exclusive because initial review of the responses indicated that the majority mentioned multiple signals, and enforcing a categorical classification would result in losing some of the richness of the data. The first coder manually evaluated all responses, developed a detailed codebook with examples (see below), which was given to the second coder, who evaluated 20% of the responses. Discrepancies were discussed iteratively between the coders until resolution was reached. Inter-rater reliability was assessed with Cohen's kappa (see Tables A6 and A9). All codes achieved substantial agreement or higher ($\text{kappa} > .70$).

When asked to describe why they felt this work was AI-generated, participants provided four main explanations for their perceptions. Beyond a few instances where AI use was disclosed proactively or retroactively by the colleague, one of the most common reasons why people believed they had received AI-generated workslop was because of the presence of linguistic or stylistic features they associated with AI-produced content. This included “overly flowery” wording that involved unusually long, verbose, or sophisticated language, “robotic” or “technical” language choices that were not perceived as being human-like, as well as a “sycophantic” tone that frequently used buzzwords. For text-based content like emails, memos, and meeting notes, the presence of formatting choices like bullet-points, boldfaced or italicized text, emojis, and certain punctuation markers - specifically the long em-dash - were also commonly reported as indicators of AI-generated content. The presence of blurry features such as human

hands missing fingers, or statistical visualizations missing key axes, were perceived as indicators of AI-generated content for visual forms of work.

The second reason was feeling that the tone of the work “did not sound” like their colleague (i.e. low consistency over time), where comparisons of the content or style of the workslop were made against their personal knowledge of the colleague who sent them this work. Third, people believed content to be workslop when it missed crucial context that was relevant to the work task at hand, which increased their perceptions that the work was created at least in part by AI rather than by the colleague themselves. This included situations where colleagues sent plans for a work proposal that referenced products their company did not sell, or the use of nonstandard metrics to assess a specific project’s outcomes. Finally, people also perceived work to be workslop when it was completed in a timeframe that was too fast for a human to complete.

Table A8. Qualitative codebook for assessing the prevalence of perceived signals of workslop.

	Code	Description	Examples	Inter-rater reliability (Cohen’s Kappa)	Prevalence (%)
1	AI disclosure	Mentions that the AI use was disclosed proactively by the sender, retroactively by the sender, after a conversation or investigation, or implicitly / accidentally with information indicating AI use (i.e., “This was generated by AI” being left in an email)	“The email title always starts with “AI translation.”” “The employee told me that they used AI to generate the report.” “I ask him more about it and he told me that he put the prompts in AI to write him a report based on his prompts and information that he provided to the AI.”	.92	1.4%
2	Stylistic and linguistic cues	Mentions formatting, aesthetic, or linguistic indicators that are perceived as being associated with AI-generated content, including “sounding like AI” or “looks like what AI would put out”	“The content was accurate, but stylistically it was stereotypical AI writing, with bulleted points, em-dashes, boldface headings, etc.” “I believe you can do it! I would absolutely love to be a part of this exciting meeting and delve into the topics with you further. It’s not annoying -- it’s motivating! Keep it up. I could continue, but I’m making myself sick. Generative AI has a disgustingly syrupy and syncophantic writing style that is suited only for psychopaths.”	.86	33.7%

3	Did not sound like their colleague	Mentions that the content, tone, style, or quality of the work was perceived as being different than that of the sender's typical work; leveraging personal knowledge of the sender	"The email did not sound like her at all. She always sounds friendly in her emails when she speaks with me and this email was different. She sounded stiff and stilted in the message. Some of the words in the message are not ones she would have used." "It is very easy to identify when you know the person. You recognize words and terms they don't usually use. The words are highly technical and come from a knowledgeable person. You know they didn't write it."	.92	18.5%
4	Missing context or contains errors	Mentions that the work was highly generic or missing crucial context necessary to meaningfully advance the work task at hand, or that it contained errors or inaccuracies that did not support the work	"It included products that we no longer or have never offered at our specific location. It also drew incorrect conclusions that made no sense in the context of our workflow." "It lacked details about the location and the products that we were selling."	.75	26.9%
5	Speed of task completion	Mentions that the work was completed at a speed that was perceived as being dissimilar to a human (i.e., "too fast for a person to do")	"It was sloppy and out of order. It contained info that would take a human days to complete"	.74	5.3%

Appendix 10. Responses to receiving workslop

Table A9. Responses to receiving workslop.

Option	N	% of participants
Spoke to the colleague directly	161	44.5%
Did nothing and moved on	141	39.0%
Spoke with a peer/teammate	53	14.6%
Documented the occurrence	36	9.9%
Spoke with my manager	36	9.9%
Spoke with the colleague's manager	12	3.3%
Something else	7	1.9%
Spoke to HR	5	1.4%

Note. These data are from the subset of individuals who described a specific instance of receiving workslop ($n = 362$).

Appendix 11: Consequences of workslop

The results presented here are from the subset of individuals who described a specific instance of receiving workslop ($n = 362$).

Table A10. Impact of receiving workslop on emotions.


Emotional Response	Mean (SD)	Indicated “Very” or “Extremely”
Annoyed	2.82 (1.24)	31.1%
Frustrated	2.63 (1.26)	29.2%
Confused	2.31 (1.17)	19.1%
Angry	1.81 (1.00)	7.6%
Offended	1.74 (1.02)	6.3%
Content	1.79 (1.11)	11.2%
Interested	2.05 (1.16)	15.0%
Happy	1.57 (1.00)	8.4%
Impressed	1.52 (0.96)	7.1%

Table A11. Change in items of Mayer et al. (1995) organizational trust framework.

Item	Less true	No change	More true
They are very capable of performing their job.	47.8%	38.5%	13.7%

They have knowledge about the work that needs to be done.	39.5%	45.8%	14.7%
They are well-qualified.	39.1%	47.5%	13.4%
They really look out for what is important to me.	37.2%	50.0%	12.8%
Sound principles seem to guide their behavior.	35.5%	53.0%	11.5%
Their behavior is very consistent.	29.5%	51.9%	18.6%
They will go out of their way to help me.	29.5%	53.6%	16.9%
They have a strong sense of justice.	26.8%	62.3%	10.9%
They have good values.	25.1%	59.6%	15.3%

Figure A1. Change in Stereotype Content Model perceptions.



Appendix 12: Regression models predicting the consequences of workslop

We fit four separate linear regression models to identify significant predictors of four outcomes: hours spent per month dealing with workslop (Hours Spent), negative emotional response to receiving an instance of workslop (Negative Emotion), desire to collaborate with the workslop sender again in the future (Collaboration), and perceived cognitive effort of dealing with an instance of workslop (Effort). All models contained the same set of predictors with the exception of Hours Spent, which excluded the amount of time spent dealing with the one participant-provided instance of workslop. As is recommended when searching for meaningful predictors across related models, we adjust coefficient p-values to account for a false discovery rate (see main text). Predictors significant below $q < .05$ are bold in the following table.

Table A12. Predictors of four consequence outcomes for receiving workslop.

term	Hours Spent	Negative Emotion	Collaboration	Effort
Gender [Woman]	0.1 (0.11)	0.12 (0.11)	-0.15 (0.1)	0.13 (0.1)
Age	-0.01 (0.06)	0.03 (0.06)	0.03 (0.05)	-0.13 (0.05)
Education	0.01 (0.05)	-0.04 (0.05)	-0.03 (0.05)	-0.16 (0.05)*
Extraversion	-0.12 (0.06)	0 (0.06)	0.01 (0.05)	-0.11 (0.05)
Openness	0.04 (0.06)	0.05 (0.06)	-0.08 (0.06)	0.04 (0.06)
Neuroticism	0 (0.06)	0.03 (0.06)	0.05 (0.06)	-0.09 (0.06)
Conscientiousness	-0.01 (0.06)	0.08 (0.06)	-0.11 (0.05)	-0.06 (0.05)
Agreeableness	-0.03 (0.06)	0 (0.06)	0.02 (0.06)	-0.07 (0.06)
Need for Cognition	0 (0.06)	0.12 (0.06)	-0.05 (0.05)	-0.03 (0.05)
Work Pride	-0.01 (0.07)	0.01 (0.07)	0.1 (0.06)	0.23 (0.06)*
Work Arrangement	0.1 (0.06)	0.03 (0.06)	0.07 (0.05)	0.02 (0.05)
Workplace AI Encouragement	0.17 (0.06)	-0.02 (0.06)	0.11 (0.06)	0.08 (0.06)
Manager [True]	0.3 (0.13)	0.37 (0.13)	0.08 (0.12)	0.44 (0.12)*
Psychological Safety	0 (0.07)	-0.2 (0.07)	0.12 (0.06)	-0.08 (0.06)
Team Climate (Shared Vision)	-0.02 (0.07)	-0.1 (0.07)	0.17 (0.07)	0.03 (0.07)
Team Climate (Task Orientation)	0.08 (0.07)	0.12 (0.07)	-0.07 (0.07)	-0.01 (0.06)
Quantitative Workload Index	0.02 (0.06)	0.03 (0.05)	-0.11 (0.05)	0.01 (0.05)
AI Agency	-0.06 (0.07)	-0.06 (0.07)	-0.06 (0.07)	0.03 (0.07)
AI Valence	0.16 (0.1)	-0.09 (0.09)	-0.03 (0.09)	0.12 (0.09)
Trust in AI	-0.1 (0.09)	-0.04 (0.08)	0.3 (0.08)*	-0.13 (0.08)
Sender [Peer on another team]	-0.14 (0.15)	-0.24 (0.14)	-0.34 (0.14)	-0.31 (0.14)
Sender [Client]	0.06 (0.29)	0.01 (0.29)	0.1 (0.28)	-0.69 (0.28)
Sender [Direct report]	0.24 (0.16)	-0.1 (0.15)	-0.13 (0.15)	-0.14 (0.15)
Sender [Supervisor]	0.03 (0.2)	-0.17 (0.2)	0.2 (0.19)	-0.48 (0.19)
Sender [Higher-up]	-0.23 (0.23)	-0.17 (0.23)	0 (0.22)	-0.58 (0.22)
Sender [Lower level]	-0.04 (0.29)	0.16 (0.29)	-0.48 (0.27)	0.13 (0.27)
Time spent on workslop example		0.07 (0.05)	-0.11 (0.05)	0.16 (0.05)*

Signif. codes: '****' = $q < .001$, '***' = $q < .01$, '**' = $q < .05$

Note. Coefficient estimates and standard errors in parentheses, followed by FDR-adjusted p-value (q).