Overcoming Barriers to AI Disclosure: The Role of Granularity and Sociocultural Enablers in Reducing Perceived Stigmas

JESSICA HE, IBM Research, USA
AVINASH AKELLA, IBM, USA
HYO JIN DO, IBM Research, USA
JUSTIN D. WEISZ, IBM Research, USA

Despite growing use of generative AI in the workplace, common stigmas depict AI users as lazy or incompetent and their work as low-quality or untrustworthy. Concerns about these stigmas deter people from disclosing their use of AI, resulting in widespread concealment and harms arising from this lack of transparency. Through a survey of workers who use AI (N=162), we explore whether disclosures that offer additional granularity beyond simple acknowledgments of AI usage can mitigate reservations around disclosure. We found that disclosures indicating compliance with AI policies, how AI contributed, and presence of human review are effective in increasing users' comfort with disclosure and reducing perceived stigmas. However, granular disclosures are insufficient to meaningfully change disclosure decisions on their own. We contextualize these findings in an ecosystem of sociocultural enablers and barriers identified from participants' disclosure experiences and share implications for supporting AI disclosure in the workplace.

CCS Concepts: • **Human-centered computing → Empirical studies in HCI**; • **Social and professional topics →** *Socio-technical systems.*

Additional Key Words and Phrases: Disclosure, Transparency, Stigma, Generative AI

ACM Reference Format:
Jessica He, Avinash Akella, Hyo Jin Do, and Justin D. Weisz. 2026. Overcoming Barriers to AI Disclosure: The Role of Granularity and Sociocultural Enablers in Reducing Perceived Stigmas. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26), June 25–28, 2026, Montreal, QC, Canada.* ACM, New York, NY, USA, 27 pages. https://doi.org/10.1145/3805689.3806737

1 Introduction

The adoption of generative AI in the workplace has been accompanied by the emergence of stigmas associated with its use. People who use or are suspected of using AI are perceived as lazier, less competent, and less diligent [36, 53]. These stigmas extend to hiring outcomes and content created with AI, which is perceived as lower quality and less trustworthy [30, 35, 36, 56, 58, 69]. As a result of these stigmas, users often feel deterred from disclosing their use of AI [53, 76]. Even so, transparency remains critical, especially as generative AI adoption in the workplace continues to grow, and organization leaders increasingly prioritize responsible AI use [12, 63]. Undisclosed uses of AI can prevent people from identifying AI-generated content, limiting their ability to appropriately calibrate their trust and increasing risks of deception [10, 17, 40, 73]. Transparency is also

Authors' Contact Information: Jessica He, jessicahe@ibm.com, IBM Research, Cambridge, MA, USA; Avinash Akella, Avinash.Akella@ibm.com, IBM, New York, NY, USA; Hyo Jin Do, dohyojingina@gmail.com, IBM Research, Cambridge, MA, USA; Justin D. Weisz, jweisz@gmail.com, IBM Research, Yorktown Heights, NY, USA.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
FAccT '26, Montreal, QC, Canada
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2596-8/2026/06
https://doi.org/10.1145/3805689.3806737

critical for informed decision-making, consumer protections, and maintaining stakeholder trust [31, 54]. As such, there is a pressing need to overcome the barriers that currently preclude AI disclosure in the workplace.

Organizations and regulatory sectors have sought to address these barriers by requiring users to label AI-generated content. However, current disclosure guidelines vary widely across domains and organizations, specifying different types of information that must be revealed. For example, ACM's guidelines encourage authors to disclose which part of the work was generated by AI [3]. Some journals, such as Science [59], ask authors to disclose full prompts, AI tools used, and version numbers. Despite these requirements, concealment of AI usage remains prevalent [29, 76], highlighting the need to better understand users' views on different forms of disclosure. Prior work has suggested that granular disclosure details may reduce users' concerns about negative judgment [27, 36], but little is known about which specific types of information can help them overcome barriers to disclosure.

Our work investigates whether different types of granular disclosures that describe *how* AI was used – specifically, those that convey human oversight, responsibility, and ownership – can mitigate users' reservations around disclosure. Giray [27] posits that "by being transparent about their AI use, academic writers and researchers can demonstrate that AI enhances rather than undermines their contributions." Kadoma et al. [36] similarly suggests "signal[ing] the extent to which AI was involved in the creative process" to mitigate AI suspicion. Motivated by these ideas, we conducted a scenario-based survey (N=162) with workers who use AI to investigate the impacts of including different details of their AI workflows on their comfort level in disclosing AI use, perceived stigmas in disclosing, and decisions to disclose. Our findings indicate that compared to a simple acknowledgment of AI use, disclosures that specify compliance with AI policies, how AI was used, and human review of AI outputs can increase comfort in disclosure and reduce perceived stigmas. However, the detailed disclosures we tested were not more effective in changing decisions to disclose compared to a simple disclosure, revealing the limits of transparency in overcoming barriers to disclosure. We find that additional sociocultural factors, such as communication from leadership and peer practices, may be needed in conjunction with granular disclosures to shift disclosure behaviors.

Our work provides empirical evidence of effective strategies for reducing perceived stigmas of AI use and disclosure. We make the following contributions:

(1) We identify types of disclosure information and sociocultural factors that can help users overcome perceived stigmas that act as barriers to AI disclosure in the workplace.
(2) We provide a more holistic view of the interplay between AI stigma and disclosure by elucidating the ecosystem of considerations, enablers, and barriers that inform users' decisions to reveal or conceal their use of AI.
(3) Based on these findings, we share implications for developing more effective AI disclosure approaches and mitigating harms that can result from disclosure in professional contexts.

By informing the development of future AI disclosure approaches, our findings will advance transparency practices that support the needs and preferences of people who create content with AI.

2 Related Work

Our work builds on prior research on AI stigmas and contributes to work in developing transparency practices and de-stigmatizing AI.

2.1 Prevalence of AI stigma

People face stigma for using AI at work across many domains, including technology [26], education [4], academic research [27], healthcare [23], and journalism [69]. Both AI users and their work can be targets of stigma [56]. For users, stigmas include negative judgments of their competence and diligence [26, 36, 53, 56, 76]. In professional

contexts, these stigmas can manifest as anticipated or actual impacts on users' jobs, such as unfavorable hiring outcomes, lower pay, and displacement [27, 34, 36, 42, 53]. Importantly, such consequences are results of the social devaluation of AI users rather than outcomes of poor performance. For content produced with AI, stigmas include negative judgments of quality, authenticity, and trustworthiness [27, 30, 35, 36, 56, 58, 69]. Stigmas can reflect actual judgments from others or judgments that users anticipate receiving (i.e., perceived stigmas), and they can arise from actual uses of AI or merely suspected use [36, 53]. We organize stigmas identified by prior work into four categories, which we investigate in our study: self-judgment, external judgment of the user, external judgment of the content, and job impacts. We define and identify the provenance of these stigmas in Table 1.

Stigma	**Definition**	**Provenance**
Self-judgment	Users' judgments of themselves for using AI, including feelings of shame and incompetence	Zhang et al. [76]
External judgment of user	Negative judgments of the person who used AI enacted by other people	Gai et al. [26], Kadoma et al. [36], Reif et al. [53], Sarkar [56], Zhang et al. [76]
External judgment of content	Negative judgments of content produced (either wholly or in-part) with AI enacted by other people	Giray [27], Grigsby et al. [30], Jakesch et al. [35], Kadoma et al. [36], Sarkar [56], Schilke and Reimann [58], Toff and Simon [69]
Job impacts	Concerns about job and career-related impacts, such as displacement and adverse hiring outcomes	Giray [27], Hwang [34], Kadoma et al. [36], Lin and Parker [42], Reif et al. [53]

Table 1. Types of stigma surrounding AI use and disclosure in the workplace, which informed our study design.

2.2 Policies and barriers of AI disclosure

Disclosure of AI use has become increasingly crucial as transparency policies take effect and AI-generated content becomes less distinguishable from human work. The EU AI Act requires providers of AI systems to ensure outputs are "detectable as artificially generated or manipulated" [24]. In the United States, four states have similarly signed AI governance bills with labeling requirements [60]. Similar policies have been enacted at the organization level, such as by publication platforms, journals, and workplaces [2, 6, 45, 47, 64, 72]. Among these policies, there is no singular standardized approach to disclosure. Some approaches have high-level requirements, such as the Microsoft Enterprise AI Services Code of Conduct, which requires customers to "disclose when the output, decisions, or actions are generated by AI" [47]. Others specify details that should be included in a disclosure, such as prompts, names of AI tools, and descriptions of how AI was used [2, 14, 49, 50, 59]. Beyond legal obligations, transparency on AI use is also important to help users calibrate their trust, make informed decisions, and avoid harms of deception [13, 54, 73].

Despite the prevalence of such requirements and their importance, adoption of AI disclosure remains inconsistent [29, 32]. In the Academ-AI Dataset, Glynn [29] found 500+ research publications that include undeclared AI-generated content, with one in five evading detection by publishers that ban such content. Prior literature has identified several barriers to AI disclosure, including inconsistencies in disclosure guidelines across disciplines, as well as individual differences in awareness and views of the need for disclosure [20–22, 25, 33, 37, 55, 74, 75]. Another key barrier is stigma. Regardless of whether stigmas are substantiated or perceived, several studies have found that they lead to hidden uses of AI [29, 53, 76]. In stigmatized domains, AI disclosure can also result in tangible consequences. In journalism, disclosure can decrease audience trust in articles, even when their accuracy or fairness are not compromised [43, 69], and freelance creators who declare AI usage experience reduced project

earnings [34]. Given the importance of disclosure for both users and consumers, our study seeks to understand how to mitigate users' stigma-related reservations in disclosing AI usage.

2.3 De-stigmatizing AI use

No empirical research currently exists on interventions that effectively reduce perceived AI stigmas. Researchers and practitioners have suggested possible means of de-stigmatization, ranging from sociocultural shifts to design interventions in AI systems.

Proposed strategies for sociocultural shifts include both top-down and bottom-up approaches. For example, researchers and practitioners have called for organization leaders to position AI's role as augmentative and encourage its use [18, 23, 48]. Some companies have taken this encouragement a step further by including AI usage as a criteria for good performance in employee evaluations [1]. Another top-down approach is to provide guidelines on accepted, ethical uses of AI, which may encourage disclosure of such uses. For example, in universities, educators have enacted course policies that explicitly outline approved and non-approved uses of AI [49, 65]. However, as Gai et al. [26] found, encouragement and even incentives from organization leaders are not always sufficient to overcome stigma and enable AI adoption. Longer-term, bottom-up strategies are also needed. In the context of AI chatbots, Dergaa et al. [18] suggest educating users to improve AI literacy and dispel misconceptions, promoting positive interactions with AI, and providing resources to address self-stigmas. Expressing respect and appreciation for honest AI disclosures may also help remedy the shame that authors experience [8].

Prior work has also suggested exploring design interventions as a means of de-stigmatization. For example, Nielsen [48] proposed eight UX design strategies, such as setting appropriate expectations of AI capabilities to calibrate user trust, providing controls that preserve user agency, and clearly defining the role of AI in a workflow. Nielsen [48] also suggests social-proofing (e.g., sharing team AI usage statistics) to signal acceptance, as well as integrating AI features into existing tools to address possible feelings of shame in using standalone AI assistants. Related work on certification labels indicate that communicating auditing outcomes, i.e., compliance with regulations and ethical standards, improve end user perceptions of AI [57]. Similarly, researchers have called for increased transparency as a means of de-stigmatization [27, 36, 48], presenting a Catch-22 — while normalizing disclosure can encourage usage, people may not feel comfortable disclosing until AI usage is normalized. Nevertheless, these works fall short in identifying the specific types of information that can help people feel more comfortable using AI and disclosing their usage.

3 Study

We conducted a scenario-based, within-subjects survey (N=162) to understand users' perceptions of detailed AI disclosures and factors that influenced their disclosure decisions. Our work explores the following research questions:

RQ1. How do different types of AI disclosure details affect users' comfort in disclosing, perceived stigmas, and decisions on whether to disclose AI use?
RQ2. Beyond perceived stigmas, what factors inform users' decisions to disclose or conceal their use of AI?

3.1 Disclosure variants

We devised eight disclosure variants to test in our survey, including a simple acknowledgment of AI use as a baseline. As this baseline disclosure was intended to reflect the high-level AI acknowledgments used today, we fashioned it after existing AI labels (such as Meta's [46], TikTok's [68] and Medium's [45]): "Parts of this work were created with AI."

We then developed seven granular variants that showcase human oversight (presence of human review and iteration), responsibility (compliance with AI policies, benefits added by AI), and ownership (efforts in prompting, human vs. AI contributions) – facets of human agency that we hypothesized would lend credibility to AI-assisted work by demonstrating rigor in the creation process. We iteratively developed descriptions and examples of each variant with feedback from colleagues to ensure clarity. The variants are summarized in Table 2, and we refer to them by the names shown in this table.

Variant	Description	Example
Simple (Baseline)	A simple acknowledgment that you used AI.	*"Parts of this work were created with AI."*
Policies	An acknowledgment that you followed all applicable AI usage policies.	*"The author's uses of AI were approved by company and platform policies."*
Human contributions	A description of your contributions to the work.	*"The author was responsible for content and idea decisions."* or *"The author was responsible for refining AI-generated content."*
AI contributions	A description of how you used AI (i.e., the AI's contributions to the work).	*"AI was used to conduct background research and make stylistic refinements."*
AI benefits	A description of the benefit that AI provided to the work.	*"AI improved the quality of this content."* or *"AI improved the efficiency of the creation process."*
Prompt	A summary of the AI prompt(s) that you used in the work.	*"The author prompted AI to act as a domain expert and brainstorm new ideas."*
Review	An acknowledgment that you reviewed and approved AI-generated content, modifications, or suggestions before using them in your work.	*"The author reviewed AI-generated content to ensure quality and accuracy."*
Iteration	An acknowledgment that you iterated on AI-generated content, modifications, or suggestions before using them in your work.	*"The author iterated on AI-generated content multiple times to ensure quality and accuracy."*

Table 2. Disclosure variants tested in this study with different types of granular information about how AI was used, highlighting human oversight, responsibility, and ownership. Variant names were used for internal purposes; descriptions and examples were shown to participants in the survey.

3.2 Measures

We assessed user perceptions of each disclosure variant with three measures: **comfort level** in using the variant to disclose AI use, **perceived stigmas** when using the variant, and **decisions** on whether they would disclose with the variant. These measures are supported by the Theory of Planned Behavior, which posits that human behavior is a result of attitudes, subjective norms, and perceived behavioral control [5]. These three facets inform behavior intention, which then drives behavior. In our survey, we evaluated behavior intention (disclosure decision), along with the attitudes (comfort level) and subjective norms (perceived stigmas) that mediate it. Perceived behavioral control is out-of-scope for our study since it is not directly tied to stigma; further research is needed to investigate attribution mechanisms that can improve perceived behavioral control to encourage disclosure.

Comfort level was measured on a forced-choice[1] 6-point scale: "1- Very uncomfortable," "2 - Uncomfortable," "3 - Somewhat uncomfortable," "4 - Somewhat comfortable," "5 - Comfortable," and "6 - Very comfortable." To assess perceived stigmas, we presented multi-select options for the four types of stigma identified from prior work (Table 1), along with the option, "I would not have any of these reservations." Finally, we asked participants whether they would disclose their use of AI with the information given in the variant with options for *yes, no,* and *unsure.*

3.3 Procedure

In the survey, participants were first asked to describe a scenario of AI use in which they created content for work and experienced at least one form of stigma, as well as whether they had chosen to disclose their use of AI and why. Each participant was then shown descriptions and examples of the disclosure variants in a randomized order. The (1) policy, (2) review, and (3) iteration variants were only shown if participants indicated that they had (1) complied with applicable AI usage policies, (2) reviewed AI-generated content, and (3) iterated on AI-generated content via direct changes or re-prompting. All other variants were shown to all participants.

For each variant, participants were asked about their perceptions of using it to disclose their use of AI in the scenario they shared by selecting their comfort level, perceived stigmas, and disclosure decision. We grounded questions in participants' actual uses of AI rather than standardized hypotheticals to assess attitudes, perceived norms, and behavioral intentions with greater ecological validity. Participants were then asked to indicate all the information they would want to include in an AI disclosure for their scenario, with options to select any of the variants previously presented to them, as well as the option to write in additional information. We refer to this condition as AGGREGATED disclosure. Participants were asked to rate their aggregated disclosure on the same three measures to help us understand whether the combination of multiple types of information would affect disclosure outcomes. The survey concluded with demographic questions about participants' geography, job role, and general perceptions of generative AI. The full survey instrument can be found in Appendix A.

3.4 Participants

Participants were recruited internally from a large technology company and externally via Prolific. Our screening criteria included experience using AI to create content for work that was shared with others, as well as experiencing at least one form of stigma stemming from their use of AI. An a priori power analysis (assuming a small-to-moderate effect size of 0.3 [70], $\alpha = 0.05$, and power = 0.95) indicated a required sample size of 145. In total, we collected 194 complete responses. Filtering out low-quality responses and those that failed attention checks resulted in a final sample size of 162 participants, with 60 from internal recruitment and 102 from Prolific [2]. All participants met the same screening criteria, and Kruskal-Wallis tests showed no statistical differences between internal and external participants; hence, we combine them for analysis. The survey took approximately 15 minutes to complete. Internal participants were incentivized via a raffle of 15 prizes with values equivalent to $25 USD. Prolific participants were compensated $5 USD. All survey responses were anonymous. In accordance with our institution's policies, we were unable to collect demographic information such as gender and age, and collection of geographic information and job roles was limited to pre-specified options.

Our sample spanned a global geography, with 46.3% based in EMEA (Europe, Middle East, Africa), 33.3% in the Americas, and 9.3% in APAC (Asia-Pacific)[3]. Internal participants worked in diverse job roles, with the most

[1]We chose to forgo a "neutral" option because in the context of this question, neutrality can be interpreted as lack of discomfort, making it redundant with the "comfortable" options.

[2]Responses that were nonsensical or did not answer the question were deemed low-quality. Attention checks included an instructional manipulation check and nonsensical item [52]. Since we required participants to share a scenario in which they had experienced stigma, we also filtered out those who selected "I did not have any of these reservations" when describing their scenario of AI use.

[3]Two participants chose not to disclose their geography.

common being software development & support (23.3%), marketing & communications (16.7%), consulting (13.3%), design (10.0%), and sales (6.7%). Among external participants, the most common job industries were information technology (23.5%), education & EdTech (14.7%), healthcare (9.8%), advertising & marketing (6.9%), and finance & financial services (6.9%). Participants also varied in approaches to AI use, with 51.9% taking an enthusiastic approach, 40.1% taking a pragmatic approach, and 8.0% trying to avoid using generative AI. These approaches were adapted from Gkinko and Elbanna [28]'s taxonomy of AI chatbot users.

3.5 Analysis

Our quantitative analysis focused on understanding the impact of granular AI disclosure variants on comfort levels, perceived stigmas, and decisions to disclose, compared to a simple acknowledgment. Hence, we use the SIMPLE disclosure as the baseline and compare its responses with each of the variants. To analyze comfort levels, we determined average comfort ratings given to each variant. We first tested the distribution of the difference of pair-wise comfort levels between the baseline and each of the variants for normality. A Shapiro-Wilk's test [62] showed that they were not normally distributed, as comfort levels tended to skew towards a preferred variant type. Hence, we used the non-parametric Wilcoxon signed-rank test [16] to make pairwise comparisons between the baseline and each variant. We report effect sizes using rank-biserial correlation, and we interpret effect sizes $r < 0.3$ as small, $0.3 \leq r < 0.5$ as moderate, and $r \geq 0.5$ as large [70]. To analyze perceived stigmas, we converted responses to a binary stigma-indicator variable, with the presence of stigma indicated by selection of any of the four stigma types, and no stigma indicated by selecting "I would not have any of these reservations." For both perceived stigmas and disclosure decisions, we used McNemar's test [44] to test for significant categorical shifts: from presence to absence of stigma (or vice versa) and from *no* or *unsure* disclosure decisions to *yes* (or vice versa). We report effect sizes using Cohen's g, and we interpret effect sizes $.05 \leq g < .15$ as small, $.15 \leq g < .25$ as moderate, and $g \geq .25$ as large [15].

To analyze qualitative data, two authors conducted an inductive, team-based thematic analysis [11] to code open-ended responses in which participants explained how they reasoned about their disclosure decisions. The two researchers first independently coded the same subset of responses, discussed and agreed on a set of codes, then independently coded the remainder of the data, discussing and iterating on any disagreements and new codes until consensus was reached. The researchers who conducted this analysis were situated in technology-oriented work contexts where AI use is prevalent and encouraged but not unilaterally accepted and sometimes socially scrutinized. We acknowledge that our positionalities may have shaped our interpretations; we engaged in iterative discussion to surface assumptions and identify interpretive differences during analysis. Through this process, we constructed a set of themes that describe sociocultural enablers and barriers to disclosure.

4 Results

4.1 Scenarios of AI use for work

Participants shared scenarios of AI use across various contexts and modalities. Most participants (69.1%) used AI to create text-based content, though some created visual content (16.7%) or code-based content (11.7%). Participants shared the content with various audiences, including internal colleagues (77.1%), internal leadership (53.7%), external stakeholders (25.9%), the general public (16.6%), and a venue-specific audience (e.g., social media followers or conference attendees) (9.8%). The content was largely a mix of human and AI contributions, with 53.09% being a human-AI blend, 30.25% primarily AI-created, 9.26% entirely AI-created, and 7.41% primarily human-created. Kruskal-Wallis and post-hoc Dunn tests [19, 38] revealed no major differences between different contribution types or amounts; hence, we analyze scenarios in aggregate. Additional scenario details and minor differences are described in Appendix B.

Just over half of participants disclosed their use of AI (55.6%). 29.6% of participants experienced self-judgment, 58.0% perceived external judgments of themselves, 70.9% perceived external judgments of their content, and 25.9% were concerned about impacts on their job. P32 explained, *"my career and reputation were built based on how well I could write, but now I'm expected to use GenAI to do the first draft. It's humiliating to have learn to use such a tool."* For P16, perceived external judgments were validated when they disclosed their AI use: *"some audience members looked at the AI-generated pictures and seemed disappointed, saying that AI-generated images degrade the worth of your work."* Two participants were concerned that signaling AI use would cause misconceptions about how they used it, indicating a need for more granular disclosures:

> *"I take my job very seriously so I never use [AI] to actually make decisions for me, but it's handy when I just need to create a chart or a fast email, but sometimes I don't feel like telling other[s] because they might misunderstand and think I've used to create the data I'm talking about instead of just making a chart or something like that."* (P117)

> *"I had to put in the report the use of AI in the project recommendations and I felt that the client may think that I did not put enough effort in the task because I used AI and they may deem it negatively and not well researched which was not the case at all."* (P118)

4.2 Impact of disclosure variants (RQ1)

We find evidence that some types of granular disclosure information can increase users' comfort in disclosing AI use and reduce their perceived stigmas. Specifically, compared to the SIMPLE baseline, HUMAN CONTRIBUTIONS details and AGGREGATED[4] disclosures increased comfort levels, while POLICIES, AI CONTRIBUTIONS, and REVIEW variants were effective in both increasing comfort and reducing perceived stigmas. However, compared to the SIMPLE baseline, none of the variants tested significantly shifted participants' disclosure decisions from *no* or *unsure* to *yes*, and the PROMPT and ITERATION variants significantly swayed disclosure decisions towards *no*. We summarize these results in Table 3 and describe additional factors that impact barriers to disclosure through our qualitative analysis in Section 4.3.

4.2.1 Impact on comfort level. For the SIMPLE baseline, participants' average comfort level was 3.95 ($SD = 1.33$). Comfort levels showed statistically significant increases in the POLICIES variant with large effect size, as well as HUMAN CONTRIBUTIONS, AI CONTRIBUTIONS, REVIEW, and AGGREGATED. In contrast, including PROMPT details decreased comfort.

We further separated participants by those who had chosen to disclose in the scenario they provided and those who did not disclose (Figure 1). We compared their responses to identify the disclosure variants that were most effective in influencing people in scenarios where they had not disclosed. Average comfort level for the SIMPLE baseline was 4.21 ($SD = 1.25$) for those who disclosed and 3.58 ($SD = 1.35$) for those who did not disclose. The POLICIES variant showed a statistically significant increase in comfort for both the disclosed and undisclosed groups with large effect size. Within the group that did not disclose, AI CONTRIBUTIONS and REVIEW additionally resulted in higher comfort. The PROMPT variant resulted in a non-significant difference for participants who did not disclose, but significantly lower comfort for participants who did disclose, with large effect size. Hence, including details on POLICIES, REVIEW, and AI CONTRIBUTIONS were most effective in helping people who did not disclose feel more comfortable to do so. However, including PROMPT information adversely impacted people who were previously comfortable disclosing. All values, including p-values and effect sizes, can be found in Appendix D.

[4]We report details on selected compositions of the AGGREGATED disclosure in Appendix C.

Disclosure Variant	Comfort Level			Perceived Stigmas			Disclosure Decision		
	Mean (SD)	*W*	r	*No stigma*	χ^2	g	*Yes decisions*	χ^2	g
Simple (Baseline)	3.93 (1.32)	-	-	22.83%	-	-	67.90%	-	-
Policies	4.72 (1.21)	472.0***	**.63**	50.45%	16.45***	**.36**	81.08%	1.39	.17
Human contributions	4.14 (1.23)	2641.0*	.21	31.48%	3.01	.13	65.43%	.20	.04
AI contributions	4.22 (1.24)	2202.0*	.27	34.56%	6.35*	**.19**	70.37%	.25	.06
AI benefits	4.07 (1.39)	2394.0	.14	30.24%	2.52	.13	65.43%	.21	.05
Prompt	3.49 (1.43)	1810.0***	**.40**	21.60%	.02	.02	47.53%	19.32***	**.31**
Review	4.26 (1.30)	1582.5**	**.36**	39.24%	15.72***	**.31**	65.82%	.11	.04
Iteration	3.81 (1.32)	2378.5	.02	21.83%	0.00	.01	53.52%	6.92**	**.23**
Aggregated	4.40 (1.19)	1292.5**	**.37**	31.46%	1.56	.11	72.72%	.15	.06

Table 3. Summary of the impact of disclosure variants on comfort level, perceived stigmas, and disclosure decision. For each variant, we report average comfort level (rated on a scale of 1–6), percentage of participants who indicated no perceived stigmas, and percentage of participants who selected *yes* as their disclosure decision in the variant. We also report statistical significance of differences for each variant compared to the SIMPLE disclosure – for perceived stigmas and disclosure decisions, this change was measured by the *change in categorical shifts*. Blue highlights indicate values that significantly increased from the baseline; orange highlights indicate outcomes that significantly decreased. Asterisks indicate the level of statistical significance: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Bold values indicate significant differences with moderate ($.3 \leq r < 5$ or $.15 \leq g < .25$) or large ($r \geq .5$ or $g \geq .25$) effect sizes.

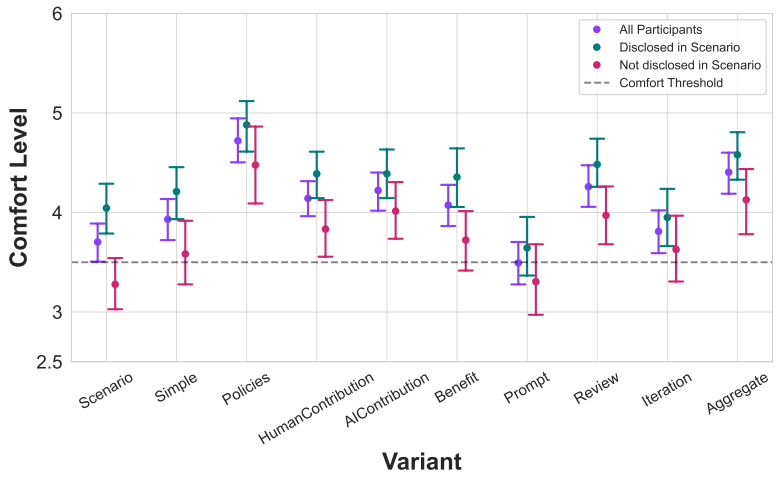


Fig. 1. Means and 95% confidence intervals of comfort level ratings for each disclosure variant. This chart compares all participants, participants who had disclosed in their scenario, and participants who had not disclosed.

4.2.2 Impact on perceived stigmas. Participants were primarily concerned about external judgments of the content, followed by perceived external judgments of themselves for having used AI. Fewer participants experienced self-stigmas and concerns about impacts on their job. We summarize results on perceived stigmas in Table 4.

Disclosure Variant	Self-judgment			External-user			External-content			Job impacts		
	%	χ^2	g	%	χ^2	g	%	χ^2	g	%	χ^2	g
Simple (Baseline)	15.43			45.06			52.47			16.05		
Policies	3.60	7.69**	**.42**	29.73	7.84**	**.30**	36.94	6.62*	**.24**	10.81	.31	.12
Human contributions	15.43	0.04	.00	36.42	3.13	.13	43.83	2.91	.12	16.05	.04	0.00
AI contributions	16.05	0.00	.02	33.95	5.35*	**.17**	43.83	3.52	.14	16.67	0.00	.02
AI benefits	14.81	0.00	.02	35.80	4.00*	**.15**	45.06	2.16	.11	15.43	0.00	.02
Prompt	20.99	1.83	.13	40.12	1.02	.08	54.32	.08	.03	22.84	3.23	.18
Review	14.56	0.00	.02	31.01	12.02***	**.27**	44.30	3.67	.15	15.19	.05	.05
Iteration	19.01	.96	.12	45.77	.33	.05	50.00	0.83	.08	16.90	.04	0.0
Aggregated	11.19	.94	.15	38.46	.41	.06	50.35	.18	.04	17.48	1.14	.14

Table 4. Proportion of participants who perceived each stigma in each variant. Statistically significant decreases in participants who perceived a stigma are highlighted in blue. Asterisks indicate the level of statistical significance: *p <0.05, **p <0.01, ***p <0.001. None of the variants caused a significant increase in perceived stigma.

With the SIMPLE baseline, 77.17% of participants indicated that they would feel stigma in disclosing. The proportion of participants who perceived stigma decreased significantly in the POLICIES and REVIEW variants with large effect size, as well as AI CONTRIBUTIONS (Table 3). We further examined reductions in specific stigmas across variants (Table 4). Compared to the SIMPLE baseline, there were significant decreases in the proportion of participants who perceived external judgment of themselves with POLICIES and REVIEW, with large effect sizes, as well as AI CONTRIBUTIONS and AI BENEFITS. The POLICIES variant additionally reduced self-judgments and perceived external judgments of the content. We did not observe meaningful differences in concerns on job impacts across variants.

4.2.3 Impact on disclosure decisions. About half (54.9%) of participants chose to disclose their use of AI in the scenario they shared. When shown a SIMPLE disclosure, this proportion increased to 67.90%. Among the variants, the proportion of *yes* decisions were highest for POLICIES (81.08%), AGGREGATED (72.72%), and AI CONTRIBUTIONS (70.37%) variants. Distributions of disclosure decision are shown in Figure 2.

In analyzing the proportion of participants whose disclosure decisions changed between the SIMPLE baseline and each variant, we found that a significant proportion of participants' decisions changed from *yes* to *no* or *unsure* for PROMPT and ITERATION. For all other variants, differences were not significant, including conversions to *yes*. In contrast, a comparison of participants' disclosure decisions to those they actually made in their scenario yields significant differences. The proportion of participants who converted to *yes* significantly increased for the following conditions: POLICIES, AGGREGATED, AI CONTRIBUTIONS, HUMAN CONTRIBUTIONS, AI BENEFITS, SIMPLE, and REVIEW (Appendix E). Hence, in some cases, a simple disclosure suffices, while in other cases, more granular disclosures can encourage users to positively change their disclosure intentions. Given the lower comfort ratings and greater prevalence of perceived stigmas of the SIMPLE baseline (Sections 4.2.1 and 4.2.2), granular disclosures may be more effective in mitigating reservations in certain cases, and additional sociocultural considerations likely influenced disclosure decisions when a simple disclosure was insufficient. These findings indicate that while granular disclosures can play a role by increasing sense of comfort in disclosing and reducing perceived stigmas, there is a broader ecosystem of factors that interact to determine users' ultimate decisions. We provide further insight into this ecosystem through qualitative findings in Section 4.3 and our discussion in Section 5.1.

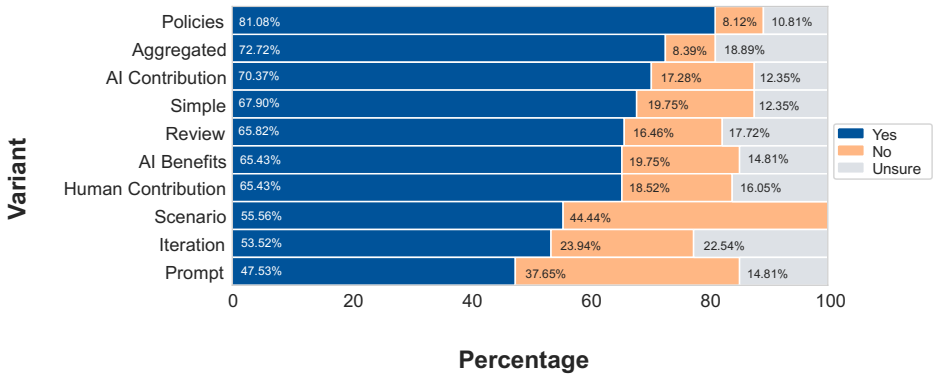


Fig. 2. Proportion of participants who selected *yes*, *no*, and *unsure* when asked whether they would disclose their use of AI using the information provided in each variant. We also include the breakdown of decisions in the scenarios of AI use that participants shared.

4.3 Factors influencing disclosure decisions (RQ2)

Participants were asked to share the reasons behind their disclosure decisions in the scenarios they shared. These responses provided further insight on the social enablers and barriers in disclosure. We summarize these findings in Table 5.

Reasons for disclosing	**Definition**
Transparency	General motivations to be transparent.
Calibrating trust and expectations	Helping the audience set appropriate trust and expectations in AI-assisted content.
Normalizing and promoting AI use	Encouraging others to use AI.
Education	Teaching others how AI can be used.
Ethical practice	Motives relating to fairness, morality, honesty, and openness.
Self-protection	Avoiding penalties for AI content issues or non-disclosure.
Encouragement from leadership	Demonstrating use of AI in workplaces where it is encouraged or required.
Compliance	Following set guidelines for disclosure.
Reasons for *not* disclosing	**Definition**
Fear of negative judgment	Concerns about how others would perceive them or their work for using AI.
Self-judgment	Feelings of shame or discomfort associated with AI use.
Take full credit	Wanting to take credit for the work without sharing with AI.
Unnecessary or irrelevant	Feeling that disclosure was not necessary or relevant for the context or audience.
Insufficient AI contribution	AI was not used enough to warrant disclosure, or the user made significant contributions and/or felt ownership over the content.
Policy gaps	Policies on disclosure are too vague or do not exist.

Table 5. Reasons that participants chose to disclose or not disclose their uses of AI. These reasons provide insight into the sociocultural enablers and barriers to disclosure.

4.3.1 Reasons for disclosing. Among participants who had disclosed, many did so for the sake of **transparency.** For example, P89 *"wanted to be transparent about the tools [they] used while also emphasizing that the final content reflected [their] own expertise and judgment."* Some went beyond general transparency motives and discussed a desire to help their audience **calibrate trust and expectations**, such as by alerting them of possible quality issues: *"I think it's important for the team to know that parts could be modified by AI so they can be...critic[s] about it."* (P13). Others recognized the value of transparency in **normalizing and promoting AI use** and sought to encourage this shift through proactive disclosure: *"I disclosed my use of AI and how so that I could help make it more acceptable and comfortable for others to do so"* (P38). P12 similarly hoped to influence their team's AI culture: *"I want to be able to trust my colleagues and therefore I have to create transparency first."* In addition to engendering AI acceptance, participants also used disclosure as a means of **education** on the capabilities and opportunities of AI. For example, P42 wanted *"to inform the team that it is so easy to get the scripts which can save time."*

Participants were also motivated by a desire to follow **ethical practices**, such as P130, who *"wanted to be open with [their] colleagues and not hide anything."* Others felt the need to disclose out of **self-protection** against consequences of AI errors or non-disclosure. For example, P67 wanted to shift accountability away from themselves: *"The AI has previously included erroneous information in summaries of my work and I don't want to be associated with it."* P162 disclosed to protect against social consequences: *"I was unsure of the quality of the code and because it clearly used coding conventions that I wouldn't use normally, so admitting that I used an LLM to generate it was less embarrassing than being 'accused' of having used it."* Beyond these internal motivators, participants also disclosed due to **encouragement from leadership** or for **compliance** with their workplaces' disclosure guidelines. For example, P102's team had been *"asked to proactively use AI,"* and P10's team *"wanted to benchmark how much GenAI [they] were using."*

4.3.2 Reasons for not disclosing. Among participants who did not disclose, reasons included social stigmas and other factors. Stigma-related reasons primarily included a **fear of negative judgment** from others as well as **self-judgments**. P91 wrote, *"I was worried what the client might think - it could seem lazy or careless to rely on AI for something I'm supposed to be doing myself."* Similarly, P85 described their own internal judgments: *"I still feel kind of ashamed that a person with my background, I have a bachelor's degree in communication and media, needs the help of AI in order to be more productive and to spread the message better."* Participants' comments also provided insight into the drivers behind these stigmas, such as beliefs that low AI literacy may worsen judgments: *"I sent out a company wide email about some changes. I felt a few people that aren't technologically hip would think it's cheating or lazy"* (P67).

Beyond stigma, some participants felt that their effort and ownership in the work may be threatened as a consequence of disclosure. Some participants wanted to **take full credit** for the content. As P80 explained, *"I did not want to take away from my effort"*, and P84 similarly *"was proud [of] the end-result and wanted all the credits."* Others felt that disclosure was **unnecessary or irrelevant** for the context or audience of their use case. For example, P32 described a lack of audience interest: *"When discussing with stakeholders, they were more interested in the final result than they were the process in which it came about."* Some participants also felt that disclosure was not needed for **insufficient AI contributions**, such as when *"the idea was mine, and AI only helped with the implementation"* (P152). Finally, **policy gaps** were a barrier to disclosure, particularly gaps in clarity on whether disclosure policies existed and what they entailed. This confusion was exemplified by P30, who wrote, *"I did not disclose because I was unsure whether occasional AI use needed to be reported. Basically, I wasn't sure whether such disclosure is required somewhere."*

4.3.3 Insight into disclosure variants. Participants' motives for disclosing can be supported by the variants tested in this study, which may provide insight on why they showed increased comfort and lessened perceived stigmas for certain variants. For example, details on HUMAN REVIEW and each party's CONTRIBUTIONS are useful for calibrating trust and expectations, and illustrating AI CONTRIBUTIONS specifically can be a means of education.

Ethics and compliance motives can in part explain the popularity of the POLICY variant, as users may be eager to demonstrate that their uses of AI abided by applicable policies.

Participants' comments also provided insight into why specific disclosure variants – namely, ITERATE, PROMPT, and in some cases, AGGREGATED – had less favorable perceptions per our quantitative analysis. For example, comments about participants' sense of ownership over AI-assisted content, especially after modifying it, indicate that disclosure may have been considered unnecessary or even counter-productive in the ITERATE variant. P7 described their co-creative workflow and why their iterations negated the need for disclosure: *"I had iterated on every bit of the content and reworked AI-generated output to ensure it didn't read AI-like, and didn't copy any part of the output verbatim. My first draft then went to a writer/editor to review and polish. I.e., there was one AI step followed by two human review/edit/polish steps, essentially neutralizing [AI-generated content]."* Additionally, as our survey did not specify where or how disclosures would appear, participants also may have had concerns about the length of the PROMPT and AGGREGATED disclosure statements, especially if they had used multiple prompts or selected multiple pieces of information to include. As P117 explained, *"I would want it to be short and simple, not in the way. I am ok with the disclosure but it should not be something someone would be drawn to on the page."*

5 Discussion

We found that compared to a simple acknowledgment of AI usage, granular disclosures that describe human oversight, responsibility, and ownership can be more effective in helping users feel comfortable disclosing and in reducing concerns about stigma. Our findings also indicate that overcoming barriers to disclosure involves an interplay of granularity and sociocultural factors. We first interpret these results (Section 5.1), then derive implications for supporting AI disclosure in the workplace and mitigating user harms (Section 5.2). We end with opportunities to explore interventions for disclosure barriers that were not addressed in this work (Section 5.3).

5.1 Ecosystem of AI disclosure factors

5.1.1 The role of granular disclosure. AI disclosure in the workplace is influenced by a complex ecosystem of enablers and barriers that encompass both the mechanisms of disclosure and the broader sociocultural dynamics surrounding it. On one hand, some forms of granular disclosure can increase users' comfort in disclosing and reduce perceived stigmas. **The efficacy of these granular disclosures depends on what they signal rather than the level of detail.** In particular, disclosures that specify compliance with AI usage policies, how AI was used, and human review were effective in both increasing comfort and reducing perceived stigmas compared to a simple, non-granular disclosure. These disclosure variants may have been effective because they clarify the extent of AI use. In doing so, they prevent misconceptions about how or why AI was used and address the concerns of losing ownership shared by our participants.

5.1.2 Beyond disclosure. **At the same time, granular disclosure cannot overcome stigma alone because stigma is socially constructed.** We observed that granular disclosures did not significantly shift participants' behavior intention [5] towards choosing to disclose compared to the simple disclosures common today [46, 51, 68]. In the terms of Ajzen [5]'s Theory of Planned Behavior, subjective norms, which are shaped by social dynamics and workplace culture, may not be sufficiently impacted by additional disclosure information. Accordingly, we found that sociocultural dynamics shaped by workplace leaders impact attitudes toward disclosure – disclosure was motivated by top-down encouragement of AI use, whereas a lack of clear policies acted as a barrier. Bottom-up interventions are also needed. In Gai et al. [26]'s study, adoption for an internal AI coding assistant remained low despite encouragement from company leaders, largely due to engineers' fears of how others would perceive their capabilities or work as a consequence of using AI. Their findings indicate that peer-level interventions are crucial in mitigating these stigmas. For example, as Kadoma et al. [36] and Dergaa et al. [18] suggested, user education can reduce stigma by dispelling misconceptions about AI-assisted work. In other fields where stigma is prevalent,

such as mental health, education has also proven to be an effective intervention for attitudinal change [67]. Our findings suggest that similar educational interventions should be investigated to de-stigmatize AI in the workplace, as some participants were especially concerned about stigmas stemming from people who were less familiar with generative AI (Section 4.3.2). Social proofs may be another effective bottom-up intervention. For example, Nielsen [48] suggested sharing team-wide AI usage statistics to signal acceptance, and Fang and Lee [25] proposed prevalence of peer disclosure (i.e., "how often others disclose their AI use") as a social factor that shapes writers' disclosure decisions. Given that some participants were motivated to disclose to encourage their peers do so, our findings support the viability of this approach.

In some cases, non-disclosure can be rational, not deceptive. In line with our findings on users' sense of ownership over co-created work (Sections 4.3.2 and 4.3.3), prior work has also identified factors that increase ownership and diminish the perceived need to disclose, such as the degree of effort invested and the nature of contributions [25, 33, 37, 75]. These perceptions are justified in some scenarios. For example, ACM's policy states that it is not necessary to disclose "if you are using generative AI software tools to edit and improve the quality of your existing text" [7], and recent papers have outlined uses of AI that should not require disclosure in scientific publishing [39, 55]. However, in other cases, there may be gaps in awareness of transparency requirements [32] and of the audience's transparency needs. To address these gaps, future work should explore the potential for nudges [66] to improve awareness. Akin to social nudges used in marketing to influence decision-making [71], design interventions could reveal labeling requirements or the audience's transparency needs just before users share or publish AI-assisted content.

Hence, rather than a singular approach, overcoming barriers to AI disclosure requires a socio-technical approach that addresses multiple components in parallel: disclosure guidelines and methods, communication from leadership on AI use, and signaling acceptance and transparency needs. Addressing each piece can help dismantle barriers faced by different people in different contexts and enable disclosure among more users, creating a domino effect in which increased transparency gradually reduces the societal stigmas against AI use that hinder disclosure.

5.2 Implications for AI disclosure policies

5.2.1 Towards user-centered disclosure approaches. Findings from our survey can inform improvements to one component of this socio-technical approach: disclosure guidelines and methods. In particular, clearer guidelines that consider AI users' disclosure preferences are needed to facilitate transparency practices. We found that a lack of clarity in AI policies are a barrier to disclosure (Section 4.3.2), supporting prior findings on the ambiguity of transparency requirements [21, 32, 74]. This ambiguity may also explain why even the SIMPLE baseline improved comfort and alleviated stigmas, as participants may have been uncertain about how to disclose in the scenario they shared until they were provided with a clear example.

Our findings can be operationalized to improve disclosure policies by leveraging the types of disclosure information that increased comfort and reduced perceived stigmas (Section 4.2). Organizations and practitioners can consider providing suggestions to disclose these types of information or integrating them into default disclosure statements when users share AI-assisted work. Based on our findings, one possible template is as follows: "AI was used to help with [brainstorming/drafting/editing]. All uses of AI complied with applicable policies, and all AI-generated content was reviewed by the author." Additionally, given the overall negative reactions to disclosing prompts – a practice used in journals [59, 77] and education [49, 50] – this form of disclosure may be counterproductive and motivate concealment.

5.2.2 The need for context-sensitivity. We recognize that disclosure can cause harm in workplaces where stigmas persist, especially to minority groups [26, 36]. For example, in software engineering professions, women suspected of using AI face heightened scrutiny compared to their male counterparts [1, 26]. While detailed disclosures

may help by clarifying the indispensability of human capabilities, they cannot entirely mitigate these harms, especially among audiences where *any* use of AI is perceived negatively. In these contexts, addressing stigmas via sociocultural interventions first is needed before granular disclosures can become effective. **Developing strategies to overcome barriers to disclosure is highly context-specific and must consider the broader culture of AI use and potential risks to individuals.**

5.3 Limitations and Future Work

Since the goal of this work was to support users in disclosing their use of AI, our study focused on the perspective of users rather than audiences of AI-assisted work. An important next step is to understand the audience point-of-view and whether their negative judgments of AI users, their work, and their job prospects are mitigated by granular disclosures. As prior work has found that the concealment of AI usage is mediated by external perceptions [76], dispelling the audience's stigmatized perceptions and conveying this de-stigmatization to users is critical towards promoting transparent practices. As AI-assisted content proliferates broadly, another important perspective to consider is non-users, especially given prior findings that this group tends to view AI usage less favorably compared to users [9, 61].

Additionally, although our survey was grounded in participants' real uses of AI, further research is needed to measure disclosure behaviors within users' creation and sharing workflows. For example, future work could implement different disclosure options into users' workflows and study the resulting disclosure outcomes. Finally, our participants were skewed towards information technology professionals and more openness to generative AI. These characteristics may not be representative of other groups, such as the general population of generative AI users, non-tech professionals, or people who are distrustful of AI. Since we could not collect certain demographic information per our institution's policies, further research is also needed to understand how to overcome disclosure barriers that are specific to demographic groups, especially given prior findings that minorities in workplaces are more susceptible to stigmatized penalties [26].

6 Conclusion

Our work sought to identify effective approaches for overcoming stigma-related barriers to AI disclosure in professional contexts. Through a survey of workers who use AI, we identified types of granular disclosure information and sociocultural enablers that mitigate users' reservations. Specifically, incorporating details on policy compliance, AI contributions, and human review can help – either in AI guidelines or default disclosure statements. Sociocultural interventions targeting workplaces and broader audiences, such as nudges from leaders or social proof of peer disclosure, are also necessary to overcome additional barriers and prevent harms, especially in environments that heavily stigmatize AI. Together, these approaches can be integrated to develop strategies tailored to the sociocultural dynamics of workplaces and audiences to encourage transparency. By helping more people feel comfortable in disclosing their use of AI, such strategies can foster a culture of transparency that gradually reduces unfounded negative judgments of AI use.

7 Generative AI Usage Statement

Microsoft Copilot (version number: 2.20260101.44.0) was used to improve the wording of text written by the authors, format tables, and format citations. All ideas, research, and writing originated from the authors; all AI-generated content was reviewed and refined by at least one author.

References

[1] Oguz A. Acar, Phyliss Jia Gai, Yanping Tu, and Jiayi Hou. 2025. Research: The Hidden Penalty of Using AI at Work. *Harvard Business Review* (1 August 2025). https://hbr.org/2025/08/research-the-hidden-penalty-of-using-ai-at-work

[2] ACM. 2023. ACM Policy on Authorship. https://www.acm.org/publications/policies/new-acm-policy-on-authorship.

[3] ACM SIGCHI. 2025. Papers. ACM CHI 2026 Conference Website. https://chi2026.acm.org/authors/papers/

[4] Dharel P Acut, Eliza V Gamusa, Johannes Pernaa, Chokchai Yuenyong, Anabelle T Pantaleon, Raymond C Espina, Mary Jane C Sim, and Manuel B Garcia. 2025. AI shaming among teacher education students: A reflection on acceptance and identity in the age of generative tools. In *Pitfalls of AI Integration in Education: Skill Obsolescence, Misuse, and Bias*. IGI Global Scientific Publishing, 97–122.

[5] Icek Ajzen. 1991. The theory of planned behavior. *Organizational behavior and human decision processes* 50, 2 (1991), 179–211.

[6] arXiv. Date unknown. arXiv moderation. https://info.arxiv.org/help/moderation/index.html#policy-for-authors-use-of-generative-ai-language-tools.

[7] Association for Computing Machinery. 2026. Frequently Asked Questions – Publications Policies. https://www.acm.org/publications/policies/frequently-asked-questions

[8] Aorigele Bao and Yi Zeng. 2025. AI disclosure, moral shame, and the punishment of honesty. *Accountability in Research* (2025), 1–14.

[9] Joanna Blahopoulou and Silvia Ortiz-Bonnin. 2025. Student perceptions of ChatGPT: benefits, costs, and attitudinal differences between users and non-users toward AI integration in higher education. *Education and Information Technologies* 30, 14 (2025), 19741–19764.

[10] Shannon Bond. 2024. How AI Deepfakes Polluted Elections in 2024. National Public Radio (NPR). https://www.npr.org/2024/12/21/nx-s1-5220301/deepfakes-memes-artificial-intelligence-elections

[11] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.

[12] Joseph Briggs and Sarah Dong. 2025. How Will AI Affect the Global Workforce? Goldman Sachs Research Insights. https://www.goldmansachs.com/insights/articles/how-will-ai-affect-the-global-workforce

[13] Olivia Burrus, Amanda Curtis, and Laura Herman. 2024. Unmasking AI: Informing authenticity decisions by labeling AI-generated content. *Interactions* 31, 4 (2024), 38–42.

[14] Coalition for Content Provenance and Authenticity (C2PA). 2025. *Content Credentials*. https://contentcredentials.org/ Official website.

[15] Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Routledge, New York, NY.

[16] William Jay Conover. 1999. *Practical nonparametric statistics*. John Wiley & Sons.

[17] Di Cooke, Abigail Edwards, Sophia Barkoff, and Kathryn Kelly. 2024. As Good As A Coin Toss: Human detection of AI-generated images, videos, audio, and audiovisual stimuli. *arXiv preprint arXiv:2403.16760* (2024).

[18] Ismail Dergaa, Feten Fekih-Romdhane, Jordan M Glenn, Mohamed Saifeddin Fessi, Karim Chamari, Wissem Dhahbi, Makram Zghibi, Nicola Luigi Bragazzi, Mohamed Ben Aissa, and Noomen Guelmami. 2023. Moving beyond the stigma: understanding and overcoming the resistance to the acceptance and adoption of artificial intelligence chatbots. *New Asian Journal of Medicine* 1, 2 (2023), 29–36.

[19] Alexis Dinno. 2015. Nonparametric pairwise multiple comparisons in independent groups using Dunn's test. *The Stata Journal* 15, 1 (2015), 292–300.

[20] Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2024. The AI ghostwriter effect: When users do not perceive ownership of AI-generated text but self-declare as authors. *ACM Transactions on Computer-Human Interaction* 31, 2 (2024), 1–40.

[21] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI disclosure obligations: Who, what, when, where, why, how. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–11.

[22] Ziv Epstein, Antonio Alonso Arechar, and David Rand. 2023. What label should be applied to content produced by generative AI? (2023).

[23] Pouyan Esmaeilzadeh. 2024. Challenges and strategies for wide-scale artificial intelligence (AI) deployment in healthcare practices: A perspective for healthcare organizations. *Artificial Intelligence in Medicine* 151 (2024), 102861.

[24] European Union. 2025. *Article 50: Transparency Obligations for Providers and Deployers of Certain AI Systems*. https://artificialintelligenceact.eu/article/50/ EU Artificial Intelligence Act website.

[25] Jingchao Fang and Mina Lee. 2025. What Shapes Writers' Decisions to Disclose AI Use? *arXiv preprint arXiv:2505.20727* (2025).

[26] Phyliss Jia Gai, Jiayi Hou, and Yanping Tu. 2025. Competence Penalty Is a Barrier to the Adoption of New Technology. *Available at SSRN 5255039* (2025).

[27] Louie Giray. 2024. AI shaming: the silent stigma among academic writers and researchers. *Annals of Biomedical Engineering* 52, 9 (2024), 2319–2324.

[28] Lorentsa Gkinko and Amany Elbanna. 2023. The appropriation of conversational AI in the workplace: A taxonomy of AI chatbot users. *International Journal of Information Management* 69 (2023), 102568.

[29] Alex Glynn. 2024. Suspected undeclared use of artificial intelligence in the academic literature: an analysis of the Academ-AI dataset. *arXiv preprint arXiv:2411.15218* (2024).

[30] Jamie L Grigsby, Meg Michelsen, and Cesar Zamudio. 2025. Service ads in the era of generative AI: Disclosures, trust, and intangibility. *Journal of Retailing and Consumer Services* 84 (2025), 104231.

[31] Aaron Hall. 2025. When Employees Use AI Without Legal Review. https://aaronhall.com/when-employees-use-ai-without-legal-review.

[32] Jessica He and Hyo Jin Do. 2025. Exploring Industry Practices and Perspectives on AI Attribution in Co-Creative Use Cases. In *ACM International Conference on Intelligent User Interfaces*.

[33] Jessica He, Stephanie Houde, and Justin D Weisz. 2025. Which contributions deserve credit? perceptions of attribution in human-ai co-creation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.

[34] Angel Hsing-Chi Hwang. 2025. The Dual Impact of AI Use Disclosure: Harm for Creators vs. Gains for Non-Creative Freelancers. *International Conference on Computational Social Science* (2025).

[35] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T Hancock, and Mor Naaman. 2019. 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*. 1–13.

[36] Kowe Kadoma, Danaë Metaxa, and Mor Naaman. 2025. Generative AI and Perceptual Harms: Who's Suspected of using LLMs?. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–17.

[37] Donal Khosrowi, Finola Finn, and Elinor Clark. 2023. Diffusing the creator: Attributing credit for generative AI outputs. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 890–900.

[38] William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621.

[39] Marcel LaFlamme and Natalie Meyers. 2025. *Enacting AI Disclosure in Scholarly Publishing*. https://openanthroresearch.org/index.php/oarr/preprint/view/435/703 Preprint available via OpenAnthroResearch.

[40] Satish Lalchand, Val Srinivas, Brendan Maggiore, and Joshua Henderson. 2024. Generative AI is expected to magnify the risk of deepfakes and other fraud in banking. Deloitte Center for Financial Services. https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html

[41] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A Alghamdi, et al. 2024. A design space for intelligent and interactive writing assistants. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–35.

[42] Luona Lin and Kim Parker. 2025. *U.S. Workers Are More Worried Than Hopeful About Future AI Use in the Workplace*. Research Report. Pew Research Center. https://www.pewresearch.org/wp-content/uploads/sites/20/2025/02/ST_2025.2.25_AI-Workers_REPORT.pdf

[43] Chiara Longoni, Andrey Fradkin, Luca Cian, and Gordon Pennycook. 2022. News from generative artificial intelligence is believed less. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 97–106.

[44] Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* 12, 2 (1947), 153–157.

[45] Medium Help Center. 2024. *Artificial Intelligence (AI) Content Policy*. https://help.medium.com/hc/en-us/articles/22576852947223-Artificial-Intelligence-AI-content-policy Medium Help Center – Policies & Safety.

[46] Meta. 2024. *Our Approach to Labeling AI-Generated Content and Manipulated Media*. https://about.fb.com/news/2024/04/metas-approach-to-labeling-ai-generated-content-and-manipulated-media/ Official blog post.

[47] Microsoft. 2025. *Microsoft Enterprise AI Services Code of Conduct*. https://learn.microsoft.com/en-us/legal/ai-code-of-conduct Microsoft Learn policy page.

[48] Jakob Nielsen. 2025. AI Stigma: Why Some Users Resist AI's Help. https://jakobnielsenphd.substack.com/p/ai-stigma Substack post.

[49] Northern Illinois University Center for Innovative Teaching and Learning. 2023. Class Policies for AI Tools. https://www.niu.edu/citl/resources/guides/class-policies-for-ai-tools.shtml

[50] Acacia L Overono and Annie S Ditta. 2025. The use of AI disclosure statements in teaching: developing skills for psychologists of the future. *Teaching of Psychology* 52, 3 (2025), 273–278.

[51] Princeton University Library. 2025. *Disclosing the Use of AI*. https://libguides.princeton.edu/generativeAI/disclosure Research guide page.

[52] Prolific Researcher Help Center. 2025. *Prolific's Attention and Comprehension Check Policy*. https://researcher-help.prolific.com/en/article/fb63bb

[53] Jessica A Reif, Richard P Larrick, and Jack B Soll. 2025. Evidence of a social evaluation penalty for using AI. *Proceedings of the National Academy of Sciences* 122, 19 (2025), e2426766122.

[54] Elizabeth M. Renieris, David Kiron, and Steven Mills. 2024. *Artificial Intelligence Disclosures Are Key to Customer Trust*. https://sloanreview.mit.edu/article/artificial-intelligence-disclosures-are-key-to-customer-trust/

[55] David B Resnik and Mohammad Hosseini. 2025. Disclosing artificial intelligence use in scientific research and publication: When should disclosure be mandatory, optional, or unnecessary? *Accountability in research* (2025), 1–13.

[56] Advait Sarkar. 2025. AI Could Have Written This: Birth of a Classist Slur in Knowledge Work. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–12.
[57] Nicolas Scharowski, Michaela Benk, Swen J Kühne, Léane Wettstein, and Florian Brühlmann. 2023. Certification labels for trustworthy AI: Insights from an empirical mixed-method study. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*. 248–260.
[58] Oliver Schilke and Martin Reimann. 2025. The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes* 188 (2025), 104405.
[59] Science. [n. d.]. Science Journals: Editorial Policies. Science Author Portal Website. https://www.science.org/content/page/science-journals-editorial-policies#image-text
[60] Sentinella, Richard and Zweifel-Keegan, Cobun. 2025. US State AI Governance Legislation Tracker. https://iapp.org/resources/article/us-state-ai-governance-legislation-tracker/. Last updated: 15 July 2025.
[61] Nevenka Popović Šević, Aleksandar Šević, Milica Slijepčević, and Jelena Krstić. 2025. AI adoption in higher education: Exploring attitudes and perceived benefits between users and non-users. *Online Journal of Communication and Media Technologies* 15, 4 (2025), e202528.
[62] Samuel Sanford Shapiro and Martin B Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3-4 (1965), 591–611.
[63] Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Bryce Hall. 2025. *The state of AI: How organizations are rewiring to capture value*. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
[64] St. John's University Human Resources. 2024. *Policy 1038 – Artificial Intelligence in the Workplace*. https://www.stjohns.edu/who-we-are/leadership-and-administration/administrative-offices/office-senior-vice-president-and-chief-operating-officer/human-resources/policy-1038-artificial-intelligence-workplace Human Resources policy page.
[65] Stanford Teaching and Learning Hub. 2025. Course Policies on Generative AI Use. https://tlhub.stanford.edu/docs/course-policies-on-generative-ai-use/
[66] Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven, CT.
[67] Graham Thornicroft, Nisha Mehta, Sarah Clement, Sara Evans-Lacko, Mary Doherty, Diana Rose, Mirja Koschorke, Rahul Shidhaye, Claire O'Reilly, and Claire Henderson. 2016. Evidence for effective interventions to reduce mental-health-related stigma and discrimination. *The Lancet* 387, 10023 (2016), 1123–1132.
[68] TikTok. 2025. *AI-generated Content*. https://www.tiktok.com/tns-inapp/pages/ai-generated-content TikTok in-app transparency page.
[69] Benjamin Toff and Felix M Simon. 2024. "Or they could just not use it?": The dilemma of AI disclosure for audience trust in news. *The International Journal of Press/Politics* (2024), 19401612241308697.
[70] Maciej Tomczak and Ewa Tomczak. 2014. The need to report effect size estimates revisited. An overview of some recommended measures of effect size. (2014).
[71] Tina A. G. Venema, Floor M. Kroese, Jeroen S. Benjamins, and Denise T. D. de Ridder. 2020. When in Doubt, Follow the Crowd? Responsiveness to Social Proof Nudges in the Absence of Clear Preferences. *Frontiers in Psychology* 11 (2020). doi:10.3389/fpsyg.2020.01385
[72] Wiley. 2023. *Generative Artificial Intelligence Guidelines for Authors*. Retrieved 15-Aug-2024 from https://onlinelibrary.wiley.com/pb-assets/assets/15405885/Generative%20AI%20Policy_September%202023-1695231878293.pdf
[73] Chloe Wittenberg, Ziv Epstein, Adam J Berinsky, and David G Rand. 2024. Labeling AI-generated content: promises, perils, and future directions. (2024).
[74] Alexander J Wulf and Ognyan Seizov. 2024. "Please understand we cannot provide further information": evaluating content and transparency of GDPR-mandated AI disclosures. *AI & SOCIETY* 39, 1 (2024), 235–256.
[75] Yuxin Xu, Mengqiu Cheng, and Anastasia Kuzminykh. 2024. What makes it mine? exploring psychological ownership over human-ai co-creations. In *Proceedings of the 50th Graphics Interface Conference*. 1–8.
[76] Zhiping Zhang, Chenxinran Shen, Bingsheng Yao, Dakuo Wang, and Tianshi Li. 2025. Secret use of large language model (LLM). *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–26.
[77] Chris Zielinski, Margaret A. Winker, Rakesh Aggarwal, Lorraine E. Ferris, Markus Heinemann, Jr Jose Florencio Lapeña, Sanjay A. Pai, Edsel Ing, Leslie Citrome, Murad Alam, Michael Voight, and Farrokh Habibzadeh on behalf of the WAME Board. 2023. *Chatbots, Generative AI, and Scholarly Manuscripts: WAME Recommendations on Chatbots and Generative Artificial Intelligence in Relation to Scholarly Publications*. https://wame.org/page3.php?id=106 Revised recommendations published by the World Association of Medical Editors.

A Survey instrument

Screener

(1) What is your employment status? [External only, single-select]
- Full-time
- Part-time
- Due to start a new job within the next month*
- Unemployed (and job seeking)
- Not in paid work (e.g. homemaker, retired, or disabled)*
- Other*

(2) Have you ever used generative AI to create content for work? [Single-select]
Work refers to any professional uses that are not for hobby or personal purposes. "Generative AI" refers to any features or applications (e.g., ChatGPT, Copilot, Gemini) that use AI to generate content (e.g., text, images, video, audio).
- Yes
- No*

(3) Was this work shared with other people and tied to you in some way (e.g., not shared anonymously or under your organization's name)? [Single-select]
- Yes
- No*

(4) Have you experienced any of the following reservation(s) about disclosing your use of AI, regardless of whether you disclosed or not? [Single-select]
- I felt ashamed about my use of AI.
- Other people may perceive me negatively.
- Other people may perceive my work negatively.
- There may be negative impacts on my job or career.
- Yes
- No*

**If participants selected any of these answers, they were screened out.*

Background

(1) When using AI to create content for work (e.g., text, code, visuals, audio), which of these describes your approach to disclosing your use of AI? [Single-select]
- I never disclose my use of AI, even if it's required
- I only disclose my use of AI when it's required
- I sometimes disclose my use of AI depending on how I used it, the audience, or other considerations
- I always disclose my use of AI, even if it's not required
- Other - please describe

Scenario

Think about a recent or memorable scenario when you used generative AI to create content for work (e.g., text, code, visuals, audio). Please identify a scenario where:

- The work was shared with other people
- The work was tied to you in some way (e.g., not shared anonymously or under your organization's name)
- You experienced at least one of these reservations about disclosing your use of AI, regardless of whether you disclosed or not:

  - You felt ashamed about your use of AI.
  - Other people may perceive you negatively.
  - Other people may perceive your work negatively.
  - There may be negative impacts on your job or career.

(1) Please describe the content you created with AI, including what type of content it was and its purpose. [Open response]
(2) Who did you share this content with? *Select all that apply.* [Multi-select]
    - Internal colleagues
    - Internal leadership (.e.g, managers, executives)
    - External stakeholders (e.g., collaborators, clients)
    - Venue-specific audience (e.g., followers of a private social media account, conference attendees)
    - The general public
    - Other - please describe
(3) Please indicate whether the following contributions were made by a person, AI, or both. [Matrix, single-select]
    Questions
    - Coming up with the initial idea
    - Conducting background research
    - Creating the structure, layout, or organization
    - Creating the initial draft
    - Making stylistic refinements (e.g., changes to wording, aesthetic, tone)
    - Making content refinements (e.g., changes to ideas, narrative, or message)

    Answer options
    - Person
    - AI
    - Both
    - Not applicable
(4) How much of the final work was created by AI? [Single-select]
    - Entirely AI-created
    - Primarily AI-created
    - Human-AI blend
    - Primarily human-created
(5) Did you review AI-generated content or suggestions before using them in your work? [Single-select]
    - I reviewed all of it.
    - I reviewed some of it.
    - I did not review it.
(6) Did you iterate on AI-generated content or suggestions before using them in your work? [Single-select]
    Iterating can refer to making direct changes to the content produced by AI or prompting the AI multiple times to produce the content you wanted.
    - Yes
    - No
(7) Did your use of AI comply with policies (e.g., of your organization, publication platform, or government) that govern or establish rules about your AI usage? [Single-select]
    - Yes
    - No

- I'm not aware of any applicable AI policies
- Unsure

(8) Did you disclose your use of AI to the people you shared this work with? [Single-select]
- Yes
- No

If disclosed:

(9) How did you disclose your use of AI? [Open response]
If you wrote a statement to disclose your use of AI, please write it here. Otherwise, please share details such as where you disclosed and what information you included when disclosing.

(10) For what reason(s) did you disclose your use of AI? [Open response]

(11) How comfortable did you feel disclosing your use of AI? [Single-select]
- 1 - Very uncomfortable
- 2 - Uncomfortable
- 3 - Somewhat uncomfortable
- 4 - Somewhat comfortable
- 5 - Comfortable
- 6 - Very comfortable

If not disclosed:

(9) For what reason(s) did you not disclose your use of AI? [Open response]

(10) How comfortable would you feel if asked to disclose your use of AI to the people you shared this work with?
- 1 - Very uncomfortable
- 2 - Uncomfortable
- 3 - Somewhat uncomfortable
- 4 - Somewhat comfortable
- 5 - Comfortable
- 6 - Very comfortable

All participants:

(11/12) Which of the following reservation(s) did you have about disclosing your use of AI in this scenario? [Multi-select]
- I felt ashamed about my use of AI.
- Other people may perceive me negatively.
- Other people may perceive my work negatively.
- There may be negative impacts on my job or career.
- I did not have any of these reservations.

In the following sections, you'll see different types of information that could be included in a disclosure of AI use. For each one, you'll be asked about your thoughts on including that information in a disclosure **for the scenario you just described. This disclosure would be visible to everyone you shared the work with.**

Disclosure Variants

Consider the following information that could be included in an AI disclosure for the scenario you shared. Note that this disclosure would be visible to everyone you shared the work with.

Information: [Variant description]

Example: [Variant example]

(1) How comfortable would you feel disclosing your use of AI with this information? [Single-select]
   - 1 - Very uncomfortable
   - 2 - Uncomfortable
   - 3 - Somewhat uncomfortable
   - 4 - Somewhat comfortable
   - 5 - Comfortable
   - 6 - Very comfortable

(2) Would you disclose your use of AI with this information? [Single-select]
   - Yes
   - No
   - Unsure

(3) Which of the following reservation(s) would you have about disclosing your use of AI with this information, if any? *Select all that apply.* [Multi-select]
   - I would feel ashamed about my use of AI.
   - Other people may perceive me negatively.
   - Other people may perceive my work negatively.
   - There may be negative impacts on my job or career.
   - Other - please describe
   - I would not have any of these reservations.

Aggregated Disclosure

(1) Which of the following information would you include in an ideal AI disclosure for your work? [Multi-select]
   This disclosure would be visible to everyone you shared your work with.
   - A simple acknowledgment that you used AI
   - An acknowledgment that you followed all applicable AI usage policies**
   - A description of your contributions to the work
   - A description of how you used AI (i.e., the AI's contributions to the work)
   - A description of the benefit that AI provided to the work
   - A summary of the AI prompt(s) that you used
   - An acknowledgment that you reviewed and approved AI-generated content or suggestions**
   - An acknowledgment that you iterated on AI-generated content or suggestions**
   - I would not feel comfortable disclosing my use of AI

(2) Is there any other information that you would include in an ideal AI disclosure for your work? [Optional, open response]

***Participants only saw these options if they met the criteria for them to be shown in the **Disclosure Variants** section (see Section 3.3).*

If (1) is NOT "I would not feel comfortable disclosing my use of AI":

(3) How comfortable would you feel disclosing your use of AI with a statement that summarized all of the information that you selected? [Single-select]
   - 1 - Very uncomfortable
   - 2 - Uncomfortable
   - 3 - Somewhat uncomfortable
   - 4 - Somewhat comfortable
   - 5 - Comfortable
   - 6 - Very comfortable

(4) Would you disclose your use of AI with a statement that summarized all of the information that you selected? [Single-select]
- Yes
- No

(5) Which of the following reservation(s) would you have about disclosing your use of AI with this information, if any? *Select all that apply.* [Multi-select]
- I would feel ashamed about my use of AI.
- Other people may perceive me negatively.
- Other people may perceive my work negatively.
- There may be negative impacts on my job or career.
- Other - please describe
- I would not have any of these reservations.

Demographics

(1) What is your primary job category? [Internal only, optional, single-select]
- Options: Architect, Consultant, Data Science, Design, Enterprise Operations, Finance, General Management, Hardware Development & Support, Human Resources, Information Technology & Services, Legal, Manufacturing, Marketing & Communications, Offering Management, Product Services, Project Executive, Project Management, Research, Sales, Services Solutions Management, Site Reliability Engineer, Software Development & Support, Supply Chain, Technical Services, Technical Specialist, Other

(1) What industry do you work in? [External only, optional, single-select]
- Options: Advertising and Marketing, Aerospace and Aviation, Agriculture and Farming, Art and Design, Automotive, Banking, Construction and Real Estate, Consulting and Professional Services, Education and EdTech, Energy and Utilities, Entertainment and Media, Environmental and Sustainability, Finance and Financial Services, Food and Beverages, Government and Public Services, Healthcare, Information Technology (IT), Insurance, Legal and Law Services, Manufacturing and Production, Retail and Consumer Goods, Social Services and Non-Profit, Telecommunications, Travel and Tourism, Transportation & Logistics, Other

(2) Do you research, design, or develop generative AI systems in your job role? [Single-select]
- Yes
- No

(3) Which of these statements best describes your approach to using generative AI for work? [Single-select]
These approaches do not include researching, designing, or developing generative AI as part of your job role.
- I try to avoid using generative AI and use it minimally.
- I take a pragmatic approach and use it after considering the risks and/or values it could bring to my work.
- I take an enthusiastic approach and actively explore new ways it can add value to my work.
- Other - please describe

(4) For the following statements, please select the point which best describes your feeling or impression of generative AI. [Matrix, single-select]
Statements:
- I am confident in generative AI
- Generative AI is reliable
- I can trust generative AI

Answer options:
- 1 - Not at all, 2, 3, 4, 5, 6, 7 - Extremely

(5) What are three words to describe how **you** perceive people who use generative AI? [Open response]

(6) What are three words to describe how **people in your workplace** perceive people who use generative AI? [Open response]
(7) What is the AI usage policy in your workplace? [External only, single-select]
- AI usage is not allowed
- AI usage is allowed for specific or approved uses
- AI usage is allowed
- I'm not aware of any AI usage policies

(7/8) What geography are you located in? [Optional, single-select]
- Americas
- APAC (Asia-Pacific)
- EMEA (Europe, Middle East, Africa)

B AI Usage Scenario Details

We conducted a deductive thematic analysis on participants' scenarios of AI use, applying the taxonomy of context types developed by Lee et al. [41] to describe uses of AI writing assistants. As our study focused on AI use in the workplace, most scenarios were in a professional context (61.7%); others were technical (24.1%), academic (4.9%), personal (4.9%), creative (4.3%), and journalistic (1.9%). Note that some participants described either one or multiple scenarios with more than one content modality or context; additionally, some scenario descriptions were too vague to capture content modality or context.

We also asked participants whether specific types of contributions within their scenarios of AI use came from a person, AI, or both. Table 6 reports these distributions.

Contribution Type	Person (%)	AI (%)	Both (%)
Initial Idea	75.16	12.42	12.42
Background Research	44.87	25.64	29.49
Structure	18.52	51.85	29.63
Initial Draft	27.78	53.09	19.14
Stylistic Refinements	28.57	40.99	30.43
Content Refinements	30.00	29.38	40.63

Table 6. Distribution of provenance across different types of contributions.

Within contribution types, there were no significant differences in the three measures between scenarios in which a user, AI, or both contributed the initial idea, background, research, structure, or initial draft of the content. Kruskal-Wallis and post-hoc Dunn tests revealed small differences in comfort levels between different contribution provenance for content and stylistic refinements. When AI was involved in stylistic refinements, participants were more comfortable disclosing using the HUMAN CONTRIBUTION ($Z = .023$, $p < .05$) disclosure. When AI was involved in content refinements, comfort levels were higher for the following disclosure variants: POLICY ($Z = .013$, $p < .05$), HUMAN CONTRIBUTIONS ($Z = .0064$, $p < .01$), PROMPT ($Z = .00016$, $p < .001$), REVIEW ($Z = .037$, $p < .05$), AGGREGATED ($Z = .019$, $p < .05$). Comfort in disclosing may have been generally higher when participants used AI for content or stylistic refinements because these are considered smaller contributions to content that they had produced themselves. However, perceived stigmas and disclosure decisions did not change.

Within contribution amounts, there were again no significant differences in perceived stigmas or disclosure decisions. Kruskal-Wallis and post-hoc Dunn tests revealed some differences in comfort levels. Participants felt more comfortable disclosing the PROMPT used when the content is entirely AI-generated as opposed to a human-AI blend ($Z = .00088$, $p < .001$), primarily human-created ($Z = .053$, $p < .01$), or primarily AI-created

($Z = .0053$, $p < .01$). Participants were also more comfortable using the AGGREGATED disclosure statement when the content is primarily AI-created as opposed to primarily human-created ($Z = .0068$, $p < .05$). Participants were also more comfortable using the AGGREGATED statement when the content is entirely AI-created as opposed to primarily human ($Z = .0039$, $p < .05$).

C Aggregated disclosure results

When asked which types of information they would include to form the AGGREGATED variant, the most commonly selected options and combinations were POLICIES (68%), SIMPLE (58%), REVIEW (53%), SIMPLE plus POLICIES (43%), and POLICIES plus REVIEW (40%). REVIEW was commonly selected in conjunction with other variants, including AI BENEFITS (26%), ITERATION (26%), and HUMAN CONTRIBUTIONS (21%).

Selected Choice	Support (%)
Policies	67.6%
Simple	58.8%
Review	52.9%
Simple & Policies	43.1%
Policies & Review	40.2%

Table 7. Frequently selected choices in the AGGREGATED question, ordered by Support which is defined as the proportion of responses where the choice is seen.

Antecedent	Consequent	Antecedent Support (%)	Consequent Support (%)	Support (%)	Confidence	Lift
Benefit	Review	36.27	52.94	26.47	0.73	1.38
Iteration	Review	34.31	52.94	25.49	0.74	1.40
Human Contributions	Review	29.41	52.94	20.59	0.70	1.32
Review & Policies	Iteration	40.20	34.31	19.61	0.49	1.42
Benefit	Review & Policies	36.27	40.20	19.61	0.54	1.34

Table 8. Results of an Apriori association analysis performed on the choices selected by users in the AGGREGATED question. We chose a minimum Support threshold=10%, and minimum Lift=1.3, and ordered the above results by Support (%).

Antecedent	Consequent	Antecedent Support (%)	Consequent Support (%)	Support (%)	Confidence	Lift
Benefit & Iteration & Policies	Review	11.76	52.94	11.76	1.00	1.89
Benefit & Iteration & Policies & Human Contributions	Review	11.76	52.94	11.76	1.00	1.89
Simple & Review & Iteration	Policies	12.75	67.65	11.76	0.92	1.36
Prompt & Policies	Review	13.73	52.94	11.76	0.86	1.62
Prompt & Policies	Simple	13.73	58.82	11.76	0.86	1.46

Table 9. Results of an Apriori association analysis performed on the choices selected by users in the AGGREGATED question. We chose a minimum Support threshold=10%, and minimum Lift=1.3, and ordered the above results by Confidence.

Antecedent	Consequent	Antecedent Support (%)	Consequent Support (%)	Support (%)	Confidence	Lift
Benefit & Iteration & Policies	Human Contributions & Review	11.76	20.59	11.76	1.00	4.86
Prompt	Human Contributions & Simple	18.63	18.63	10.78	0.58	3.11

Table 10. Results of an Apriori association analysis performed on the choices selected by users in the AGGREGATED question. We chose a minimum Support threshold=10%, and minimum Lift=1.3, and ordered the above results by Lift.

D Comfort level results for users who disclosed vs. did not disclose in their scenario

Disclosure Variant	Comfort Level (Scenario Disclosed)			Comfort Level (Scenario Undisclosed)		
	Mean (SD)	*W*	r	*Mean (SD)*	*W*	r
Simple (Baseline)	4.21 (1.25)	-	-	3.58 (1.35)	-	-
Policies	4.88 (1.09)	140.0***	**.62**	4.47 (1.35)	101.0***	**.64**
Human contributions	4.39 (1.20)	722.0	.18	3.83 (1.22)	624.0	.22
AI contributions	4.39 (1.21)	745.5	.16	4.01 (1.27)	394.0*	**.38**
AI benefits	4.36 (1.24)	667.5	.16	3.71 (1.28)	530.0	.13
Prompt	3.64 (1.38)	317.0***	**56**	3.31 (1.50)	608.0	.24
Review	4.48 (1.24)	488.0	.29	3.97 (1.33)	316.5**	**.44**
Iteration	3.95 (1.33)	558.0	.19	3.63 (1.30)	446.5	.17
Aggregated	4.58 (1.11)	331.5**	**.44**	4.13 (1.28)	320.5*	.29

Table 11. Summary of the impact of disclosure variants on comfort level, contrasted by their decision to disclose their use of AI in Scenario. For each variant, we report average comfort level (rated on a scale of 1–6). We also report statistical significance of differences for each variant compared to the SIMPLE disclosure. Blue highlights indicate values that significantly increased from the baseline; orange highlights indicate outcomes that significantly decreased. Asterisks indicate the level of statistical significance: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Bold values indicate significant differences with moderate ($.3 \leq r < 5$) or large ($r \geq .5$) effect sizes.

E Disclosure decisions compared to scenario decisions

Disclosure Variant	**Disclosure Decision**		
	Yes decisions	χ^2	g
Scenario (Baseline)	55.55%	-	-
Simple	67.90%	7.52^{**}	**.21**
Policies	81.08%	13.83^{***}	**.33**
Human contributions	65.43%	4.33^{*}	**.15**
AI contributions	70.37%	9.12^{**}	**.21**
AI benefits	65.43%	4.33^{*}	**.15**
Prompt	47.53%	2.23	.10
Review	65.82%	4.0^{*}	.10
Iteration	53.52%	0.17	.04
Aggregated	72.72%	4.5^{*}	**.16**

Table 12. Summary of the impact of disclosure variants on disclosure decision using Scenario case as the baseline. For each variant, we report percentage of participants who selected *yes* as their disclosure decision in the variant. We also report statistical significance of differences for each variant compared to the Scenario disclosure. This change was measured by the *change in categorical shifts*. Blue highlights indicate values that significantly increased from the baseline; orange highlights indicate outcomes that significantly decreased. Asterisks indicate the level of statistical significance: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Bold values indicate significant differences with moderate ($.15 \leq g < .25$) or large ($g \geq .25$) effect sizes.