



The Role of Inclusion, Control, and Ownership in Workplace AI-Mediated Communication

Kowe Kadoma
kk696@cornell.edu
Cornell University
Ithaca, New York, USA

Christin Munsch
christin.munsch@uconn.edu
University of Connecticut
Storrs, Connecticut, USA

Marianne Aubin Le Quéré
msa258@cornell.edu
Cornell University
Ithaca, New York, USA

Danaë Metaxa
metaxa@seas.upenn.edu
University of Pennsylvania
Philadelphia, Pennsylvania, USA

Xiyu Jenny Fu
xf89@cornell.edu
Cornell University
Ithaca, New York, USA

Mor Naaman
mor.naaman@cornell.edu
Cornell Tech
New York, New York, USA

ABSTRACT

Given large language models' (LLMs) increasing integration into workplace software, it is important to examine how biases in the models may impact workers. For example, stylistic biases in the language suggested by LLMs may cause feelings of alienation and result in increased labor for individuals or groups whose style does not match. We examine how such writer-style bias impacts inclusion, control, and ownership over the work when co-writing with LLMs. In an online experiment, participants wrote hypothetical job promotion requests using either hesitant or self-assured autocomplete suggestions from an LLM and reported their subsequent perceptions. We found that the style of the AI model did not impact perceived inclusion. However, individuals with higher perceived inclusion did perceive greater agency and ownership, an effect more strongly impacting participants of minoritized genders. Feelings of inclusion mitigated a loss of control and agency when accepting more AI suggestions.

CCS CONCEPTS

• Human-centered computing → Empirical studies in HCI.

KEYWORDS

Co-writing, large language models, autocomplete

ACM Reference Format:

Kowe Kadoma, Marianne Aubin Le Quéré, Xiyu Jenny Fu, Christin Munsch, Danaë Metaxa, and Mor Naaman. 2024. The Role of Inclusion, Control, and Ownership in Workplace AI-Mediated Communication. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3613904.3642650>

1 INTRODUCTION

Large language models (LLMs) have the potential to transform workplace communication but can also exacerbate existing societal biases. LLMs are already seeing widespread use in assisting in

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3642650>

everyday workplace tasks like writing emails and authoring documents. Recent scholarship has shown that workers can use LLMs to increase their productivity and produce higher-quality work [22]. Despite these potential benefits, LLMs also have drawbacks, including the potential to reproduce social biases [4, 31]. For instance, when LLMs generate text about particular identity groups, they may display social biases in the content and style of the resulting text [30], causing *representational harm* through stereotyping and demeaning language [26, 31]. Beyond representational harms, we focus in this work on another category of harm: *quality-of-service harm* [6, 26], in which systems unequally serve different groups, which may “result in feelings of alienation, increased labor, and service or benefit loss” for people of different identities [26]. For example, if an LLM produces text that is stylistically better aligned with one gender group over another, the result is a model that disproportionately benefits certain users over others.

Such stylistic biases, whether real or perceived by users, can impact how individuals and groups of end-users feel about and interact with AI tools. When an LLM's text style better suits one group over another, such mismatch may impact users' feelings of inclusion when using that system. Inclusion could affect other psychological factors like users' sense of the human experiences of control and ownership [25] during the co-writing process. For example, lower perceptions of inclusion may lead users to accept fewer AI suggestions; prior work has shown the relationship between accepting AI suggestions and feelings of agency and control [18]. Given the potential importance of feelings of inclusion on people's use of AI tools, we set out to examine the relationship between writer-style bias and perceptions of inclusion, control, and ownership while co-writing with AI.

In particular, we focus here on the potential impact of AI writing assistants that provide auto-complete suggestions as the user is typing text, for example, while writing an email. In this context, we ask the following questions:

- RQ1: Can the style of AI auto-complete suggestions impact people's sense of inclusion, control, and ownership over the text they write?
- RQ2: How do these perceptions differ by gender?

More specifically, in this work, we investigate the impact of a specific communication style—assertiveness—on writers and its impact on minoritized gender groups. We build on a rich history examining gender disparities, particularly in the workplace where

assertive language can be very consequential [1, 16] and associated with better work outcomes [1, 14]. While there are many cultural differences that guide expectations of gender presentation in the workplace, scholars in gender-based communication describe male talk in Western workplaces as being assertive, powerful, and authoritative [16, 20, 29]. If writing assistants making suggestions in a more assertive style are not as well-aligned with the communication style of women and members of other minoritized gender groups, the discrepancy in style alignment could cause an additional burden for women and gender minorities by increasing writing effort and task completion time.

This paper addresses the research questions above using an online experiment where participants completed a hypothetical but common and consequential workplace writing task: writing to a manager to ask for a promotion. Participants received auto-complete suggestions from one of two LLM-powered AI writing assistants, one assistant offering suggestions in an assertive style and the other in a hesitant style. We then asked participants about their feelings of inclusion, control, and ownership in the writing experience.

We found that the intervention—getting suggestions in an assertive or hesitant style from an AI assistant—did not impact participants' perceptions of inclusion. However, it did impact their perceptions of control and ownership. We did not find the assistant style treatments directly resulted in gender differences in our main measures. We do find evidence that gender plays some role: minoritized genders' perceptions of control and ownership are more greatly impacted by their (likely pre-existing) feelings of inclusion than men's. In light of these findings, we offer a new conceptual model for understanding the impact of task- and user-style alignment on people's feelings of overall agency in AI-mediated environments.

2 BACKGROUND

As LLMs become more robust with the ability to inspire ideas, revise text, and generate short stories, the writing assistants powered by large language models are increasingly seen as co-writers. There is a growing body of work studying the interaction between users and these co-writing systems [13, 19]. Recent scholarship in this area has focused on understanding how writers evaluate and integrate the suggestions provided by LLMs into their cognitive writing processes [5], and how writers proactively engage with system-generated content [27].

Along with their potential benefits, however, these models carry potential risks—risks that may not be equally distributed across different user groups. Prior work has identified several categories of harms including social systems harms, allocative harms, representational harms, quality of service harms, and interpersonal harms [26, 31]. Prior work has pointed to the existence of such harms in the context of auto-complete suggestions. For example, societal biases in the suggestions can cause information harms (a subset of social systems harms) by influencing what content is written [10, 23], and even shifting writers' attitudes, potentially without their awareness [10, 11]. While we do not know of work showing representational harms in auto-complete environments, given the potential for content influence [10, 11] and the known

representational biases that exist in LLM text generation [26, 31], the issue is likely to persist in the auto-complete context as well.

Quality-of-service harms are another concern identified by researchers. These harms reflect developer's choices and result in unintended performance disparities based on identity [6, 26]. For example, prior work has shown that audio speech recognition (ASR) systems have significantly higher error rates for Black speakers than White speakers [12]. These error rates can have psychological and behavioral impacts on Black users such as feelings of frustration and disappointment; they can cause users to have to modify their linguistic patterns [15]. Such quality-of-service harms can lead to alienation and result in greater labor burdens and loss of benefits for marginalized communities [26]. In this short paper, we extend research on this topic to consider writer-style biases in auto-complete suggestions and their potential to cause differential impacts on inclusion, control, and ownership for different social groups, with a specific focus on gender.

3 METHODS

We addressed our research questions using an online experiment where participants completed a common workplace writing task. We asked participants to write a hypothetical email to their manager asking for a promotion. We chose the promotion task since it is a situation that requires high assertiveness in context in which minoritized genders face several challenges [1, 2, 14]. The experiment had three conditions. Participants in two of the treatments received auto-complete suggestions from a language model. In one of these treatments, the model made *assertive* auto-complete suggestions, and in the other, it made *hesitant* auto-complete suggestions. A control group of participants wrote without the assistance of any model.

After the writing task, participants answered a demographic questionnaire. Participants in the AI conditions answered an additional post-task questionnaire regarding their perceptions of inclusion, ownership, and agency while writing with the AI assistant. The questions and 5-point Likert scale answers were adapted from previous work [10, 18]. The research design was approved by Cornell's IRB. The design was preregistered and is available on OSF¹.

3.1 Writing App

We used a custom experimental platform combining a rich-text editor and a writing assistant, as seen in Figure 1. The platform code was adapted from prior experimental work on auto-complete systems [10]. The suggestions work such that when the participant pauses typing, the system displays suggestions for completion that are 20 words or less. The user can then choose to accept the suggestion by pressing *tab* or *right arrow key*, or they can ignore the suggestion by continuing to type. Pressing *tab* accepts the subsequent word in the suggestion, and a user can press multiple times to accept the full suggestion.

To generate the interactive text suggestions for the experiment, we used GPT-4 via an API call with prompts carefully designed to display assertive or hesitant text. To achieve that goal, we inserted the prompt “In a [self-assured/hesitant] manner ask for a raise” before the participants' written text. After the entered text, we

¹https://osf.io/kwbev?view_only=c5639dcc6292482b8bd79f2086b38d24

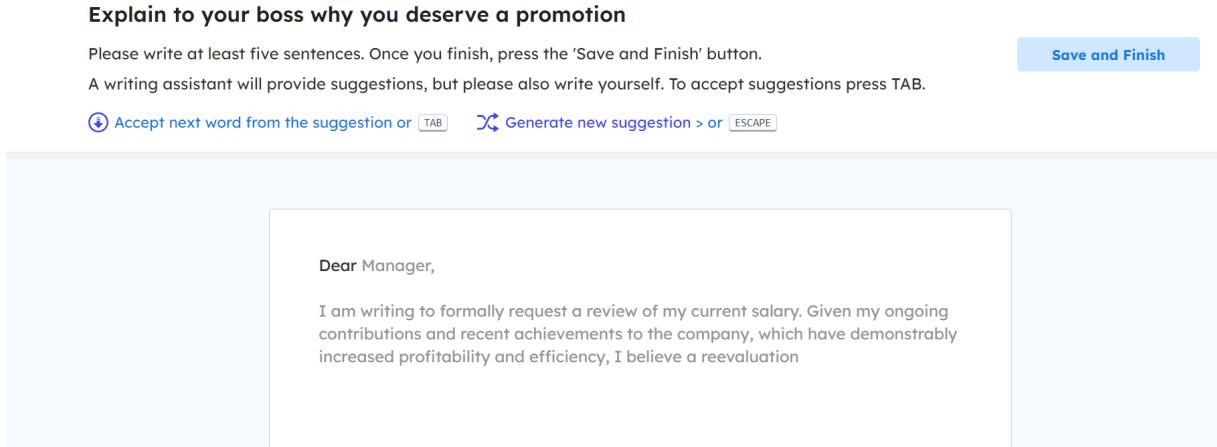


Figure 1: Screenshot of the writing task. Instructions are given in the panel at the top. Participants can hit ‘tab’ to accept the suggestion or ‘esc’ to generate a new suggestion. Written text is in black and suggestions appear in gray text.

added the affirming instruction “add one sentence continuing the email in a [self-assured/hesitant] manner asking for a raise”. This way, the suggestions maintained the desired style regardless of the text already entered by the participant.

3.2 Participant Recruitment

We recruited 750 participants from a gender-balanced sample of English-speaking US-based adults on the crowdsourcing service Prolific. The sample size was calculated using a two-way ANOVA with a power of 80% based on a pre-test survey. To ensure response quality, we required participants to have approval ratings greater than 95% on Prolific. We excluded participants who did not state their gender and manually verified that each writing sample was faithful to the writing task. In the end, we had usable data from 738 participants.

Our experiment considers two distinct gender categories: men, as category that is usually associated with greater societal influence and dominance, and minoritized genders who typically encounter societal disparities. While describing our results, we will refer to **Women**, **Non-Binary** or third gender, and “preferred to describe” (also known as **Self-Described**) as **WNBSD** participants. Of the 738 participants whose data is used, 48.9% identified as men and 51.1% as WNBSD: 47.7% identified as women, 3% identified as non-binary or third gender, and 0.4% preferred to describe their gender.

The majority of the participants were between 25-34 years old (56%), identified as white (61.9%), received a bachelor’s degree (38.2%), reported working a full-time job (59.4%), and reported working entirely in person (40.6%).

3.3 Measures

We used the following measures to examine how participants interacted with the AI assistant and their perceptions of the writing task with the AI assistant.

AI reliance. Prior work on how users interact with AI writing assistants uses the measure of mutuality to describe the level of *interaction* the writer has with the AI assistant [13, 18]. However,

in this work, we are interested in the final text output, not the event-based interactions. We introduce a new measure, *AI reliance*, as we believe this measure captures the construct more intuitively. AI reliance is the fraction of AI-written characters in the entire text. The measure uses a score from 0 to 1, where 1 means the assistant wrote the entire text (reflecting complete reliance on the assistant) and 0 means the human wrote the entire text (no reliance).

Inclusion. We draw on the concept of self-extension, which explains how objects become integral to one’s self-concept [3], to develop the notion of inclusion in AI-mediated communication. We focus on identity self-extension where the technology is a reflection of the user [18, 24]. This literature suggests that if a user feels the technology reflects themselves, they will feel included. We therefore measure inclusion with the following statements: *the writing assistant was made for people like me*, and *the writing assistant sounded like I would write myself*. Response options ranged from *strongly disagree* (1) to *strongly agree* (5).

Control. Prior work has shown that feeling in control is not limited to one action or outcome. Feeling in control can encompass engaging in an action, control over the outcome, or both [18, 27]. For example, some users may feel a lack of control since they cannot choose which suggestions they are seeing (control during the writing process). However, the fact that they can select which suggestions are incorporated into the message can impact their sense of control over the process and the final version [5]. We therefore measure control by asking participants *how much control did you feel over the process of writing the message* and *how much control did you feel over the final version of the message* [18]. Response options ranged from *not at all* (1) to *a great deal* (5).

Ownership. One aspect of ownership involves “mineness”—the idea that a specific target belongs to an individual. The target can include physical objects but can also encompass intangible objects such as words and thoughts. In our experiment, we consider two targets—the message itself and the message style. Therefore, we also investigate ownership over the writing content and style. Our questions included *to what extent do you feel like the message you wrote is yours* and *thinking back on the message writing activity, how*

much did the message sound like you [18]. The response options were the same as those for the control questions.

4 RESULTS

Our results expose the impact of writing style on perceptions of inclusion, control, ownership, and agency and hint at potential differences between genders in that respect. Our primary, preregistered analysis finds some significant effects of AI writing style on perceptions of control and ownership but not on inclusion. A follow-up, exploratory analysis suggests that perceptions of inclusion mediate the loss of agency as AI reliance increases. Additionally, the analysis suggests that heightened perceptions of inclusion lead to increased agency, particularly among minoritized genders.

We first provide an overview of the effects of using the AI autocomplete suggestions on the writing task. We examined three components of the writing task: the length of text written by participants, their writing time, and their degree of AI reliance. We performed a two-way ANOVA to compare the means between the three treatment groups (self-assured AI writing style, hesitant AI writing style, and control) and two gender groups (men and WNBSD participants). Overall, participants who wrote with the assistant wrote longer emails ($M=160.18$, $SD=82.16$) and spent less time on the task ($M=303.82$, $SD=218.53$) than participants in the control group (email length $M=95.53$, $SD=35.64$; time on task $M=363.66$, $SD=292.68$). The differences are statistically significant (length: ANOVA $F=67.86$, $p<0.001$; time: ANOVA $F=6.35$, $p=0.002$). Post-hoc tests did not observe any differences between the two AI conditions in email length and time on task. In contrast, we *did* observe significant impacts of AI writing style on AI reliance: participants who wrote with the self-assured writing style model relied on the model more ($M=0.55$, $SD=0.34$) than participants who wrote with the hesitant writing style model ($M=0.45$, $SD=0.35$). These differences were also significant ($t=3.16$, $p=0.002$). We did not observe gender effects or interaction effects between gender and the AI treatment on any of these metrics.

4.1 Inclusion

We examine the impact of AI writing style on our key measures, beginning with perceptions of inclusion. Figure 2 shows the frequency of different survey responses for each gender condition (different columns) for the inclusion statements: *the writing assistant was made for people like me* (“Assistant Was Made for You” on the left), and *the writing assistant sounded like I would write myself* (“Assistant Sounds Like You” on the right). In the condition with the hesitant AI writing style (left side of Figure 2a), 65.1% of men (first column) and 63% of WNBSD participants (second column) either agreed or strongly agreed (blue and dark blue bars in the column) that the writing assistant was made for people like them. As the figure also shows, in the self-assured AI writing style group (right side of Figure 2a), 68.3% of the men and 68.6% of WNBSD participants either agreed or strongly agreed with this statement. In general, participants more readily agreed with the “made for you” statement (Figure 2a) compared to the “sounds like you” statement (Figure 2b).

In Figure 2b, of the participants who received the hesitant treatment (left), 42.8% of the men and 45.7% of the WNBSD participants

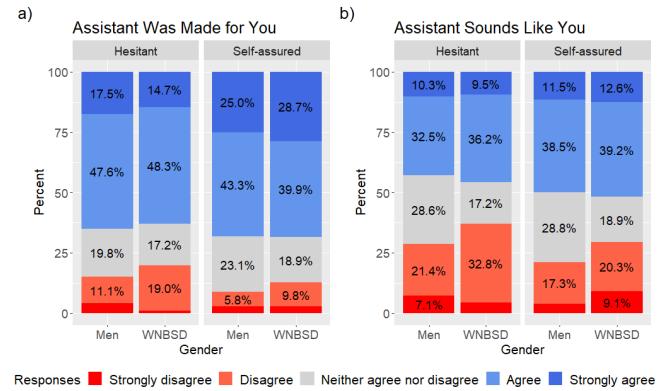


Figure 2: Participants' assessment of inclusion. Participants are more likely to say the assistant was made for them than sounds like them. Responses less than 5% are not labeled on the figure.

either agreed or strongly agreed that the assistant sounded like how they would write themselves. In the condition with the self-assured AI writing style, on the right side of Figure 2b, 50% of the men and 51.8% of the WNBSD participants either agreed or strongly agreed. A two-way ANOVA analysis revealed no significant effects for gender, AI treatment group, or the interaction between gender and treatment group across both inclusion questions shown in Figure 2.

4.2 Control

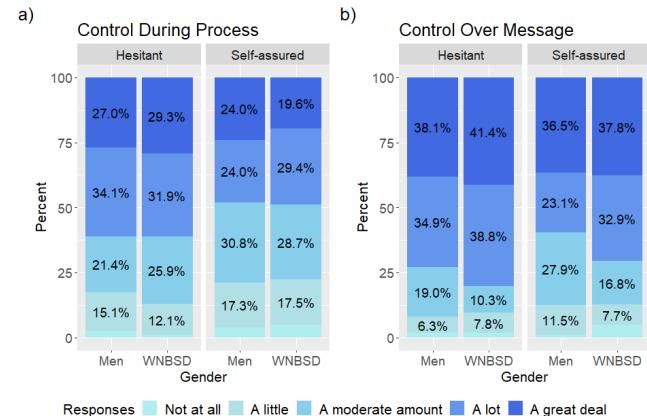


Figure 3: Participants' assessment of control. Participants assisted by the hesitant writing style model are more likely to feel greater control over the final version of the message and the writing process than participants assisted by the self-assured writing style model. Responses less than 5% are not labeled on the figure.

Figure 3 shows the frequency of different survey responses for each gender (different columns) for the following questions: *how much control did you feel over the process of writing the message* (“Control During Process” on the left), and *how much control did*

you feel over the final version of the message (“Control Over Message” on the right). In the condition with the hesitant AI writing style suggestions (left side of Figure 3a), 61.1% of the men (first column) and 61.2% of WNBSD participants (second column) reported feeling a lot or great deal of control (the darker blue and darkest blue bars in the columns) during the writing process. In the condition with the self-assured AI writing style (right side of Figure 3a), 48% of men and 49% of the WNBSD participants expressed a lot or a great deal of control during the writing process. A two-way ANOVA revealed that the difference between the AI treatment groups was significant ($p=0.004$), meaning that participants in the hesitant condition reported more control over the process. There were no significant differences for gender or the interaction between gender and the AI treatment groups for this question.

Looking at Figure 3b and control over the final version, in the condition with the hesitant AI writing style (left), 73% of men and 80.2% of WNBSD participants reported a lot or a great deal of control over the final version of the message. The right side of Figure 3b shows 59.6% of men and 70.7% of WNBSD participants who wrote with a self-assured writing style model reported a lot or a great deal of control over the final version of the message. For this question as well, a two-way ANOVA revealed that the difference between the AI treatment groups was significant ($p=0.046$), with participants in the hesitant AI writing style condition reporting greater control over the final version of the message. At the same time, there were no significant differences in relation to gender or interactions between gender and the AI treatment group for this question. In both treatment conditions, participants generally reported higher levels of control over the final version of the message (Figure 3b) than the during the writing process (Figure 3a).

4.3 Ownership

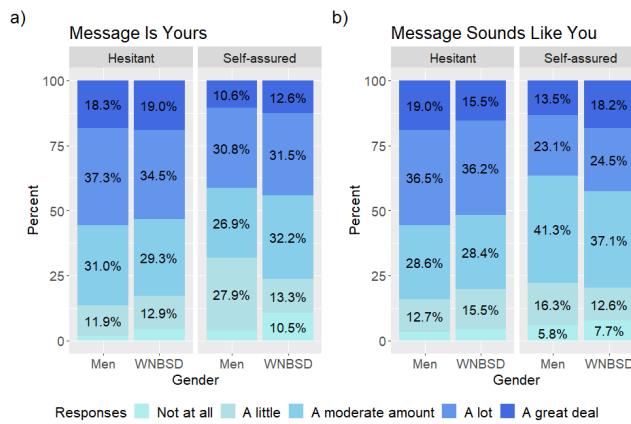


Figure 4: Participants assessment of ownership. Participants assisted by a hesitant model are more likely to say that they wrote the message and the message sounds like them. Responses less than 5% are not labeled on the figure.

Figure 4 illustrates the frequency of different survey responses for each gender condition (columns) for the following questions: *to what extent do you feel like the message you wrote is yours* (“Message

Is Yours”), and *thinking back on the message writing activity, how much did the message sound like you* (“Message Sounds Like You”). In the hesitant AI writing style condition (left side of Figure 4a), 55.6% of the men and 53.5% of WNBSD participants expressed a lot or a great deal of ownership over the message (the dark blue and darkest blue bars in the column). In the self-assured AI writing style condition (right side of Figure 4a), participants experienced a higher level of ownership, with 41.4% of men and 44.1% of the WNBSD participants expressing a similar degree of ownership over the message. A two-way ANOVA revealed that this difference between the AI treatment groups was significant ($p<0.001$). There was no significant difference for gender or the interaction between gender and the AI treatment groups for this question.

Reflecting on the style of the message, Figure 4b shows a similar trend: participants in the hesitant AI writing style condition (left) were more likely to say the final message reflected their style. Once again, a two-way ANOVA revealed that the difference between the AI treatment groups was significant ($p=0.032$). There was no significant difference for gender or the interaction between gender and the AI treatment groups and gender for this question.

Overall, our results show AI writing style impacts control and ownership. Participants who wrote with the hesitant model expressed greater control over the writing process and the final message and felt the message sounded more like them than participants who wrote with the self-assured model. However, our analysis did not reveal an impact of the AI writing style on perceived inclusion. Additionally, we did not see any effects of gender or interaction between gender and treatment group across inclusion, control, and ownership questions. We further explore these factors in the analysis below.

4.4 Exploratory Analysis: Inclusion and Agency

To better understand the relationship between the different measures and the participants’ interaction with the AI model, we perform an exploratory analysis that brings together multiple factors using a linear regression model. To perform this analysis, we combined various measures into simplified high-level constructs. First, we use a compound measure of agency as our dependent variable. The concept of agency has been defined in several different ways but broadly captures the idea of “the capacity to alter a situation” [8, 21]. The variables that reflect that in our study were control and ownership. While distinct, the literature shows these concepts are closely related: whether a person has control greatly influences their perceived ownership [7, 18]. An analysis of our measures confirmed the control and ownership questions were indeed highly correlated in our data. We thus created a score for agency by performing a row-wise average across the control and ownership questions.

The independent variables in the model consisted of a simplified inclusion variable, the AI reliance score, the AI writing style, and demographic variables. To simplify the model and analysis, we combined the two highly correlated inclusion measures into one variable by performing a row-wise average across the two inclusion questions (Table 3 in the appendix shows the correlation of the inclusion, control, and ownership measures). The model also included the measure of AI reliance as defined above, namely how much the participants accepted the AI suggestions in the writing

Table 1: OLS linear regression analysis predicting perceived agency based on demographic features, AI writing style, AI reliance, and perceived inclusion. The constant corresponds to the baseline perceived agency.

Dependent variable:	
	Agency
Age	0.024 (0.020)
Race	0.013 (0.070)
Income	0.023 (0.034)
Gender	0.571* (0.261)
‘AI Reliance’	-3.451*** (0.367)
Inclusion	0.136* (0.066)
‘AI Writing Style’	-0.244** (0.090)
Gender:‘AI Reliance’	0.251 (0.191)
‘AI Reliance’:Inclusion	0.666*** (0.098)
Gender:Inclusion	-0.204** (0.071)
Gender:‘AI Writing Style’	0.050 (0.133)
Constant	3.661*** (0.248)
Observations	489
R ²	0.354
Adjusted R ²	0.339
Residual Std. Error	0.715 (df = 477)
F Statistic	23.747*** (df = 11; 477)

Note: *p<0.05; **p<0.01; ***p<0.001

process. Finally, we include in our model the AI writing style of the model that participants wrote with (the AI treatment), as well as more detailed demographic variables: age, gender, race, and income. The race and gender features were operationalized as binary variables where 1 reflected the privileged group and 0 reflected the minoritized group (e.g., for gender, 1=men, 0=WNBSD). We also mapped the AI writing style as a binary variable (0=hesitant model, 1=self-assured model). We operationalized the age and income demographic features as ordinal variables. Note again that we consider this analysis exploratory, as it was not part of the preregistered analysis presented above that focused on the main effects.

Table 1 shows the parameter estimates of a linear regression model predicting perceived agency. We tested several versions of the model that accounted for additional interactions (like AI writing style and AI reliance), and we present the model with the best fit. The model shows a main effect of AI writing style (Table 1, line 7) on agency. The negative value of the parameter signifies that writing with a self-assured model contributes to a loss of agency. We also see main effects of gender (line 4), showing men contributed to higher levels of agency. Perceptions of inclusion (line 6) also positively impacted agency, and AI reliance (line 5) had a negative impact. These effects were qualified by interaction effects between AI reliance and perceived inclusion (line 9) and gender and perceived inclusion (line 10), which we explore next.

We explore the interaction effects to better understand their impact on perceived agency. Figure 5a depicts the interaction between AI reliance and perceived inclusion. AI reliance is the x-axis, and perceived agency is the y-axis. The graph compares the groups that showed high (dark blue), medium (blue), and low (light blue) levels of inclusion. Participants with a high level of inclusion, as indicated by the dark blue line in the figure, maintain their sense of agency even with increasing reliance on the writing assistant. Participants with a low (and to a lesser degree, medium) sense of inclusion lose their sense of agency rapidly when they rely more on AI.

Figure 5b explores how the interaction between inclusion (x-axis) and gender impacts agency. The graph compares men (blue line) and minoritized genders (WNBSD, light blue line). While inclusion and agency have a direct relationship for all gender identities, the slope for WNBSD participants is greater than the slope for men, indicating that perceived inclusion is more strongly tied to perceived agency for minoritized genders.

5 DISCUSSION

Our findings enhance the current models of “writing with AI” by introducing new factors and furthering our understanding of their relationships. The findings also hint at the role that gender may play in such a context, though this line of investigation requires additional substantiation. Overall, the findings offer implications for research and design of co-writing interactions with AI.

By introducing the construct of *inclusion*, our work extends the previous research on AI-MC, which has examined perceptions of agency while interacting with an AI writing assistant [17, 18]. We measure the concept of inclusion after the co-writing process by asking participants if they could envision themselves in the model and whether the model’s writing style matched their own. We show that this concept contributes to the perception of agency, a combined metric reflecting ownership and control over the writing process and the final product.

Informed by these findings, Figure 6 presents a new conceptual model of the factors contributing to agency in AI co-writing environments. This model is tentative: not all of these concepts were measured, let alone tested, in our work here. The model suggests that there are two paths for feeling agency in the co-writing process and its outcome. On the top, the AI writing style affects AI-user alignment—which we define as the similarity between the AI writing style and the user’s own preferences that may be derived from their identity. High AI-user alignment leads to increased perceived inclusion, positively contributing to perceived agency. In our research, we did not measure AI-user alignment. Instead, we built on the literature suggesting that minoritized genders are less likely to be aligned with assertive AI suggestions [9, 16]—which was not the case in this study. At the bottom of the model, the AI writing style affects AI-task alignment, which we define as the perceived usefulness of the AI assistant for the task (e.g., writing to ask for a promotion with an assertive AI). If there is high AI-task alignment, the user will rely on the model to complete the task. However, such reliance, as we have shown, can *negatively* contribute to the user’s perceived agency. The conceptual model includes an interaction between AI reliance and inclusion, as suggested by our exploratory results.

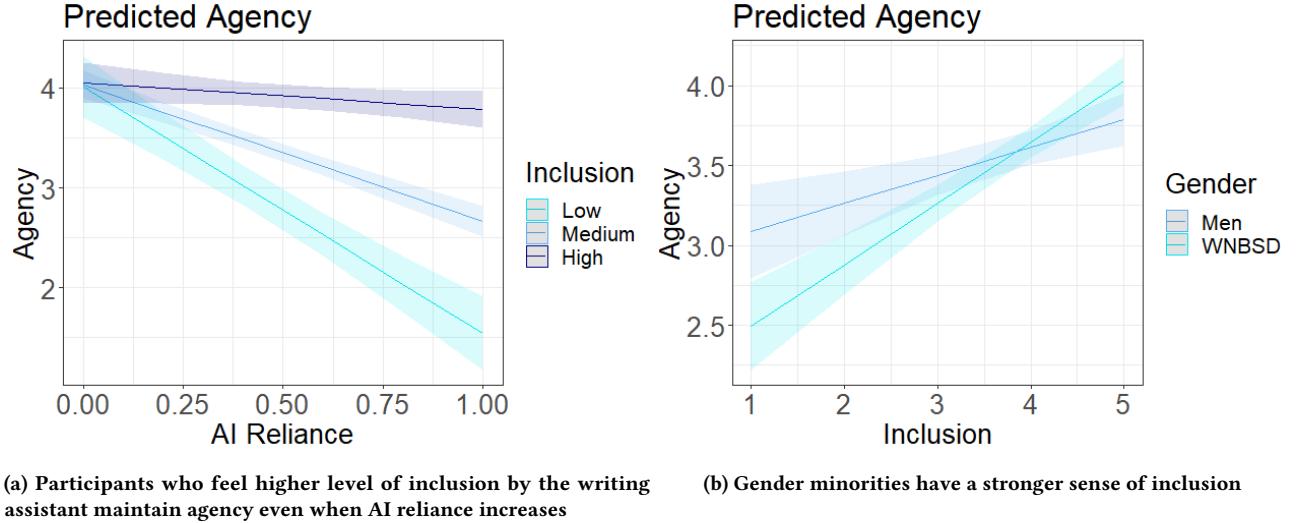


Figure 5: Interaction Effects in the Linear Regression. Figure 5a depicts the interaction between Inclusion and AI Reliance. Figure 5b depicts the interaction between Gender and Inclusion.

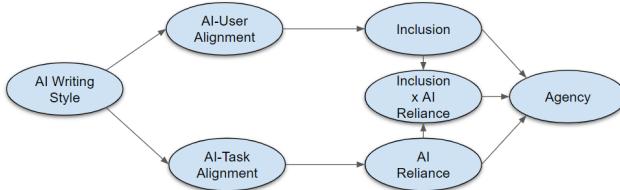


Figure 6: Conceptual Model of Agency: A proposed model of the constructs and measured variables that contribute to feelings of agency.

Our model raises several open questions that have implications for designers. The most salient question is the potential competition between AI-User alignment and AI-Task alignment. Or, how can developers strike a balance between models that are inclusive of various users but are also effective in completing the task? We note that even if the models align with both the user and task, there is remaining conflict: the AI-User alignment will ultimately contribute to agency, but the AI alignment with the task—which will increase reliance—will also *reduce* agency. This finding complicates the question of how developers can build models that help the user accomplish the task while also supporting the user’s agency, given the tension between those. Further work is needed to understand the trade-off between agency and task performance in real-world applications.

The idea of inclusion via AI-User alignment captures our hypothesis that gender identity will play a role in perceived inclusion and agency. We hypothesized that minoritized genders would feel more included in the hesitant AI writing style condition, at least compared to men. However, our main findings did not show such gender effects. Although existing literature highlights language use differs by gender in negotiation contexts, it often relies on experimental methods (such as laboratory studies) that encompass factors

beyond the scope of our experiment like job status, vocal pitch, or physical presence [1, 14]. Many of the laboratory studies show it is the small details, like the tone of voice or the wording, which makes gender differences salient [29]. The absence of physical cues, as shown in computer-mediated negotiations, reduces gender-based communication bias [28].

While gender differences did not account for our expected main effect, our exploratory analysis also suggests there are still, potentially, *some* differences between genders. A noteworthy finding in this analysis was the significant interaction between gender and inclusion in relation to perceived agency. The interaction suggests that for minoritized genders, perceptions of inclusion have a greater effect on the perception of agency compared to men. If they hold, these findings lend support to our initial hypothesis: that people of different genders may react differently to writing assistants of different styles. This finding suggests that it may be useful to devise, develop, and test language models that also optimize on generating perceptions of inclusion by people of different demographic backgrounds.

Our work explored the impact of AI auto-complete suggestions on feelings of inclusion, control, and ownership over the written text and considered the implications for users of various backgrounds. The work has a number of limitations. First, in asking people to write as if asking for a promotion, our study focused on an important and realistic workplace task. However, other workplace writing tasks should also be studied to generalize the results. Furthermore, we did not measure individual differences in assertiveness, which could have provided additional insight on the effect of the treatment. Pre-task measurement of such attitudes, though, could have introduced bias to the experimental task. We focus on English-speaking US residents. Additional work is needed to understand how users from various linguistic backgrounds and cultures interact with such AI agents. Despite these limitations, this short paper showed some of the first evidence that minoritized genders’

perceptions of control and ownership could vary in some ways from men's as a response to writer-style biases. We offer a new conceptual model for understanding the impact of these factors on people's feelings of overall agency in AI-mediated environments.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CHS 1901151/1901329.

REFERENCES

- [1] Emily Amanatullah and Michael Morris. 2010. Negotiating Gender Roles: Gender Differences in Assertive Negotiating Are Mediated by Women's Fear of Backlash and Attenuated When Negotiating on Behalf of Others. *Journal of personality and social psychology* 98 (2010), 256–267. <https://doi.org/10.1037/a0017094>
- [2] Emily T. Amanatullah and Catherine H. Tinsley. 2013. Ask and Ye Shall Receive? How Gender and Status Moderate Negotiation Success. *Negotiation and Conflict Management Research* 6 (2013), 253–272. <https://doi.org/10.1111/ncmr.12014>
- [3] Russell W. Belk. 2013. Extended Self in a Digital World. *Journal of Consumer Research* 40 (2013), 477–500. <https://doi.org/10.1086/671052>
- [4] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3442188.3445922>
- [5] Advait Bhat, Saaket Agashe, Parth Oberoi, Niharika Mohile, Ravi Jangir, and Anirudha Joshi. 2023. Interacting with Next-Phrase Suggestions: How Suggestion Systems Aid and Influence the Cognitive Processes of Writing. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3581641.3584060>
- [6] Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. 2020. *Fairlearn: A toolkit for assessing and improving fairness in AI*. Technical Report MSR-TR-2020-32. Microsoft. <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>
- [7] Niclas Braun, Stefan Debener, Nadine Psychala, Edith Bongartz, Peter Sörös, Helge H. O. Müller, and Alexandra Philipsen. 2018. The Senses of Agency and Ownership: A Review. *Frontiers in Psychology* 9 (2018). <https://doi.org/10.3389/fpsyg.2018.00535>
- [8] Laura Cesafsky, Erik Stayton, and Melissa Cefkin. 2019. Calibrating agency: Human-Autonomy Teaming and the future of work amid highly automated systems. In *Ethnographic Praxis in Industry Conference Proceedings*. Wiley Online Library. <https://doi.org/10.1111/1559-8918.2019.01265>
- [9] Susan Herring and Anna Martinson. 2004. Assessing Gender Authenticity in Computer-Mediated Language Use/Evidence From an Identity Game. *Journal of Language and Social Psychology - J LANG SOC PSYCHOL* 23 (2004), 424–446. <https://doi.org/10.1177/0261927X04269586>
- [10] Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-Writing with Opinionated Language Models Affects Users' Views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3544548.3581196>
- [11] Celeste Kidd and Abeba Birhane. 2023. How AI can distort human beliefs. *Science* 380 (2023), 1222–1223. <https://doi.org/10.1126/science.adi0248>
- [12] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R. Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences* 117 (2020), 7684–7689. <https://doi.org/10.1073/pnas.1915768117>
- [13] Mina Lee, Percy Liang, and Qian Yang. 2022. CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3491102.3502030>
- [14] Jens Mazei, Joachim Hüffmeier, Philipp Alexander Freund, Alice F Stuhlmacher, Lena Bilke, and Guido Hertel. 2015. A meta-analysis on gender differences in negotiation outcomes and their moderators. *Psychological bulletin* 141 (2015), 85–104. <https://doi.org/10.1037/a0038184>
- [15] Zion Mengesha, Courtney Heldreth, Michal Lahav, Juliana Sublewski, and Elyse Tuennerman. 2021. "I don't Think These Devices are Very Culturally Sensitive"—Impact of Automated Speech Recognition Errors on African Americans. *Frontiers in Artificial Intelligence* 4 (2021), 169. <https://doi.org/10.3389/frai.2021.725911>
- [16] Karima Merchant. 2012. *How Men And Women Differ: Gender Differences in Communication Styles, Influence Tactics, and Leadership Styles*. Claremont McKenna College.
- [17] Hannah Mieczkowski, Jeffrey T. Hancock, Mor Naaman, Malte Jung, and Jess Hohenstein. 2021. AI-Mediated Communication: Language Use and Interpersonal Effects in a Referential Communication Task. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 17 (apr 2021), 14 pages. <https://doi.org/10.1145/3449091>
- [18] Hannah N. Mieczkowski. 2022. *AI-Mediated Communication: Examining Agency, Ownership, Expertise, and Roles of AI Systems*. Ph.D. Dissertation. <https://www.proquest.com/dissertations-theses/ai-mediated-communication-examining-agency/docview/2734696274/se-2>
- [19] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals (CHI '23). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3544548.3581225>
- [20] Louise Mullany. 2007. *Gendered discourse in the professional workplace*. Springer. <https://doi.org/10.1057/9780230592902>
- [21] Gina Neff and Peter Nagy. 2018. Agency in the digital age: Using symbiotic agency to explain human-technology interaction. In *A Networked Self and Human Augmentics, Artificial Intelligence, Sentience*, Zizi Papacharissi (Ed.). Routledge, 11. <https://doi.org/10.4324/9781315202082>
- [22] Shakked Noy and Whitney Zhang. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381 (2023), 187–192. <https://doi.org/10.1126/science.adh2586>
- [23] Ritika Poddar, Rashmi Sinha, Mor Naaman, and Maurice Jakesch. 2023. AI Writing Assistants Influence Topic Choice in Self-Presentation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3544549.3585893>
- [24] Morgan Quinn Ross and Joseph B Bayer. 2021. Explicating self-phones: Dimensions and correlates of smartphone self-extension. *Mobile Media & Communication* 9 (2021), 488–512. <https://doi.org/10.1177/2050157920980508>
- [25] Richard M Ryan and Edward L Deci. 2017. *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford publications.
- [26] Renee Shelby, Shahaleh Rismani, Kathryn Henne, Ajung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3600211.3604673>
- [27] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L. Glassman. 2023. Where to Hide a Stolen Elephant: Leaps in Creative Writing with Multimodal Machine Intelligence. *ACM Trans. Comput.-Hum. Interact.* 30, Article 68 (sep 2023), 57 pages. <https://doi.org/10.1145/3511599>
- [28] Alice F. Stuhlmacher and Amy E. Walters. 1999. Gender differences in negotiation outcome: A meta-analysis. *Personnel Psychology* 52 (1999), 653–677. <https://www.proquest.com/scholarly-journals/gender-differences-negotiation-outcome-meta/docview/220142492/se-2>
- [29] Caja Thimm, Sabine C Koch, and Sabine Schey. 2003. Communicating gendered professional identity: Competence, cooperation, and conflict in the workplace. *The handbook of language and gender* (2003), 528–549. <https://doi.org/10.1002/9780470756942.ch23>
- [30] Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. "Kelly is a Warm Person, Joseph is a Role Model": Gender Biases in LLM-Generated Reference Letters. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.243>
- [31] Laura Weidinger, Jonathan Usato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Jason Gabriel. 2022. Taxonomy of Risks posed by Language Models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3531146.3533088>

A APPENDIX

A.1 Validating the Treatment

We used human annotators and computational methods to validate the effect of the auto-complete suggestions. We are able to analyze the auto-complete suggestions as well as the final message, as our writing app logs all auto-complete suggestions seen by the participants. For our analysis, we randomly sampled 100

participants from each AI treatment condition. We selected suggestions that appeared at the beginning, middle, and end of the writing task for a total of five suggestions per participant. We used LIWC-22 on the auto-complete suggestions to analyze the following features: positive tone, negative tone, pro-social behavior, politeness, conflict, clout, authenticity, and social behavior as a whole. Clout refers to the social status or confidence people display in writing. Authenticity refers to how much a person self-regulates, with low authenticity indicating the person is socially cautious and high authenticity indicating a person is spontaneous. Based on prior literature [16, 20, 29], self-assured language and hesitant language should differ along these metrics. With the exception of the positive tone metric, t-tests revealed significant differences between the means of the self-assured model suggestions and hesitant model suggestions across all the metrics (negative tone: $t=-6.73$, $p<0.001$; pro-social behavior: $t=3.29$, $p=0.0012$; politeness: $t= 3.77$, $p=0.0024$, conflict: $t=-2.27$, $p=0.025$; clout: $t=3.15$, $p=0.002$; authenticity: $t=-4.00$, $p<0.001$; social behavior: $t=3.24$, $p=0.001$). Table 2 shows several of the model suggestions with some selected self-assured suggestions highlighted in dark blue and selected hesitant suggestions highlighted in light blue.

Additional analysis showed the suggestions had an effect on the final message. We had three human annotators evaluate all of the final messages on a scale from 1-10, where 1 indicated no assertiveness and 10 indicated high assertiveness. T-tests revealed the means

between conditions to be significant (self-assured/control: $t=14.56$, $p<0.001$; self-assured/hesitant: $t= 14.56$, $p<0.001$; control/hesitant: $t= 10.09$, $p<0.001$). Using the same LIWC-22 analysis, t-tests showed significant differences between the text produced by these groups, with the exception of the positive tone metric and politeness (negative tone: $t=-7.09$, $p<0.001$; pro-social behavior: $t=4.83$, $p<0.001$; conflict: $t=-2.24$, $p=0.026$; clout: $t= 3.00$, $p=0.003$; authenticity: $t=-2.24$, $p=0.03$; social behavior: $t= 5.36$, $p<0.001$).

A.2 Exploratory Analysis Details

In the exploratory analysis (section 4.4), we performed a row-wise average across the inclusion questions and the control and ownership questions to create our agency measure. Table 3 shows the correlation between the post-task questions. The two ownership questions (*message sounds like me*, *message is mine*), two control questions (*control during process*, *control over message*), and two inclusion questions (*assistant made for me*, *assistant sounds like me*) are each highly correlated, as indicated by the bold text.

Table 4 shows the correlation analysis for the predictors in our model (Table 1 in section 4.4). The demographic features *race* and *gender* in the model are binary with 1 corresponding to the privileged group (white, male) and 0 corresponds to minoritized groups (non-white, WNBSD).

Table 2: We prompted GPT-4 to generate suggestions in self-assured style and hesitant style. We sampled a few suggestions participants saw while they were writing. Noticeable differences in the suggestions are highlighted.

Prompt	Sampled Suggestions
Come up with a continuation to the input text in a self-assured manner asking for a raise	<p>“Be justifiable to review my current compensation package. I believe an increase commensurate with my increased...”</p> <p>“Over the past year, I have consistently met and exceeded my performance objectives, taking on additional responsibilities without...”</p> <p>“My compensation should reflect this. I would appreciate the opportunity to discuss this with you.”</p> <p>“It is fair to align my compensation with these mounting responsibilities. I have done extensive research on industry.”</p>
Come up with a continuation to the input text in a hesitant manner asking for a raise	<p>“I certainly don’t want to come across as ungrateful or discontented.”</p> <p>“...my performance and contributions, if needed. I am not asking for an immediate response, but rather a thoughtful...”</p> <p>“In comparison to the industry standard, my current compensation seems a bit... well, it’s not entirely...”</p> <p>“I thought it might be time to discuss the possibility of... well, perhaps a slight adjustment in my...”</p>

Table 3: Correlation table for the post-task questions. The highest correlations are between *message is mine* and *message sounds like me* (the two ownership questions), *control during process* and *control over message* (the two control questions), and *assistant made for me* and *assistant sounds like me* (the two inclusion questions).

	message is mine	control over message	control during process	message sounds like me	assistant made for me	(assistant sounds like me)	ai_reliance
message is mine	1.00						
control over message	0.53	1.00					
control during process	0.45	0.65	1.00				
message sounds like me	0.67	0.51	0.45	1.00			
assistant made for me	0.23	0.21	0.21	0.37	1.00		
assistant sounds like me	0.13	0.17	0.22	0.21	0.56	1.00	
ai_reliance	-0.37	-0.27	-0.22	-0.28	0.13	0.16	1.00

Table 4: Correlation table for the predictors in the OLS linear regression

	Age	Race	Income	Gender	AI Reliance	Inclusion	AI Style
Age	1.00						
Race	0.16	1.00					
Income	0.22	0.06	1.00				
Gender	-0.02	-0.05	0.08	1.00			
AI Reliance	0.04	-0.01	0.09	0.07	1.00		
Inclusion	-0.08	-0.12	-0.08	0.02	0.17	1.00	
AI Style	-0.00	0.03	0.04	-0.10	0.14	0.10	1.00