



Believing Anthropomorphism: Examining the Role of Anthropomorphic Cues on Trust in Large Language Models

Michelle Cohn
mdcohn@ucdavis.edu
University of California, Davis
Davis, California, USA
Google Research
San Francisco, California, USA

Mahima Pushkarna
mahimap@google.com
Google Research
Cambridge, Massachusetts, USA

Gbolahan O. Olanubi
femio@google.com
Google Research
New York, New York, USA

Joseph M. Moran
joemoran@google.com
Google
Cambridge, Massachusetts, USA

Daniel Padgett
padgettd@google.com
Google Research
San Francisco, California, USA

Zion Mengesha
zmengesh@stanford.edu
Stanford University
Stanford, California, USA
Google Research
San Francisco, California, USA

Courtney Heldreth
cheldreth@google.com
Google Research
Seattle, Washington, USA

ABSTRACT

People now regularly interface with Large Language Models (LLMs) via speech and text (e.g., Bard) interfaces. However, little is known about the relationship between how users anthropomorphize an LLM system (i.e., ascribe human-like characteristics to a system) and how they trust the information the system provides. Participants ($n=2,165$; ranging in age from 18-90 from the United States) completed an online experiment, where they interacted with a pseudo-LLM that varied in modality (text only, speech + text) and grammatical person (“I” vs. “the system”) in its responses. Results showed that the “speech + text” condition led to higher anthropomorphism of the system overall, as well as higher ratings of accuracy of the information the system provides. Additionally, the first-person pronoun (“I”) led to higher information accuracy and reduced risk ratings, but only in one context. We discuss these findings for their implications for the design of responsible, human-generative AI experiences.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Applied computing** → **Psychology**.

KEYWORDS

anthropomorphism, trust, large language models, text-to-speech (TTS), first-person pronoun “I”

ACM Reference Format:

Michelle Cohn, Mahima Pushkarna, Gbolahan O. Olanubi, Joseph M. Moran, Daniel Padgett, Zion Mengesha, and Courtney Heldreth. 2024. Believing Anthropomorphism: Examining the Role of Anthropomorphic Cues on Trust in Large Language Models. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3613905.3650818>

1 INTRODUCTION

Millions of individuals use voice-based interfaces to interact with technology, including voice assistants like Google Assistant, Amazon’s Alexa, and Apple’s Siri, to name a few [3]. More recently, people have started engaging with a new conversational agent: large language models (LLMs) such as OpenAI’s ChatGPT¹ and Google’s Bard² (for an overview, see [52]). LLMs are capable of generating responses that can remain contextually relevant to the user’s input, can be indistinguishable from human-generated responses [4, 25], and can exceed expert-level skills [43, 44](but see also [20]). These advances are likely to increase how people anthropomorphize LLMs, or the tendency to ascribe human-like attributes to a real or imagined system behavior (e.g., progress bars, as “thinking” behaviors) [16, 18]. As prior work has shown a positive relationship between increased anthropomorphism and trust with conversational agents, avatars, and robot systems [23, 27, 40, 41, 45, 50], it is possible that a similar positive effect will be at play with LLMs. The aim of the current paper is to test the relationship between

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 EA '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0331-7/24/05
<https://doi.org/10.1145/3613905.3650818>

¹<https://chat.openai.com/>

²<https://bard.google.com/>

anthropomorphism and trust in LLMs by varying the linguistic cues the system produces.

Why does trust in LLMs matter? A growing body of work has shown that LLMs can recreate harmful stereotypes (e.g., by ability in [19]; race/ethnicity in [13]; gender in [2]), replicating the real-world biases of their training data. Additionally, in cases where there is little training data coverage, LLMs can generate fictitious data ('hallucinate') [54]. Indeed, some have raised possible harms of trusting information provided by LLMs, such as misinformation and discrimination [7, 14, 15, 51]. Therefore, a deeper understanding of the factors that contribute to how users trust an LLM is needed for responsible human-AI experiences.

This paper addresses this gap and is, to our knowledge, the first study to explicitly manipulate two linguistic anthropomorphic cues in an LLM and measure the impact on multiple dimensions of trust. In our study, participants ($n=2,165$) from the United States engaged with a pseudo-LLM system to elicit information. We manipulated the degree of anthropomorphic cues along two dimensions, Modality (text only or text + speech) and Grammatical Person (first-person, "I" or third-person "the system"). Current LLMs, such as Bard and ChatGPT, use "I" (e.g., "As a large language model, I...") and most commonly people engage with them in a text-only interface, but text + speech options are now available as well [8].

We examine the impact of these manipulations on overall ratings of anthropomorphism and trust of the system, as well as trial-by-trial ratings of accuracy, risk, and follow-up validation across five areas: health/well-being, career, medications, travel, and cooking. In sum, this work makes the following contributions:

- We provide experimental evidence that adding a computer generated voice increases both anthropomorphism of LLMs and trust that the information they provide is accurate.
- We offer evidence that, in some contexts, users rate the information an LLM provides as more accurate and less risky when the system uses the first-person singular pronoun ("I").

2 RELATED WORK

2.1 Anthropomorphism

People of all ages have shown a propensity to anthropomorphize computers; that is, to ascribe human behaviors to the system [16, 49]. Social response theories of technology equivalence ([28, 34, 35] propose that application of social rules from human-human interaction to human-computer interaction are both subconscious and automatic, such as using politeness norms [33] or applying gender stereotypes [17].

2.1.1 Anthropomorphism and the Voice. Some work has explicitly manipulated the degree of anthropomorphism by varying the presence or absence of a voice [31, 37, 50]. For example, [50] found that participants who "drove" an autonomous vehicle that had a name ("Iris") and produced pre-recorded human audio recordings showed stronger anthropomorphism (a composite of 4 questions) and trust than when the system did not have its own voice. [37] found that avatars presented with a text-to-speech (TTS) voice increased ratings of trust, relative to conditions with no voice. In a similar vein, [26] found that interacting dyads better cooperated in computer-mediated communication when their speech was generated with

TTS, compared to a text-only chat. Furthermore, [31] found that external raters gave higher ratings of anthropomorphism and trust for a conversational interface when it used a TTS voice (relative to no voice).

While often a human voice generates stronger feelings of anthropomorphism than a TTS voice [38], TTS voices still evoke human-like attributes, such as gender [17, 36], race/ethnicity [24], and age [55]. Additionally, advances in neural TTS mean that computer generated voices 'sound' increasingly human-like to listeners [11]. Taken together, these studies suggest that the presence of a voice – even if computer-generated – will evoke an anthropomorphism response in the current study.

2.1.2 Anthropomorphism and Nouns. While less studied than the voice or visual cues, there is some work showing that anthropomorphism can be reflected in the nouns used to describe a system. In a study of user reviews of Amazon Echo, [36] found that users who had more social interactions with the Echo (e.g., as a friend or family member) were more likely to also use the third-person singular feminine pronoun, "she", or the name "Alexa" to describe the system (more human-like), compared to "it" or "Echo" (less human-like).

Fewer studies have explicitly manipulated the nouns the system uses. [32] include, "It doesn't even refer to itself as 'I,'" as one of the reasons users know a computer is not human-like. While voice assistants have often been developed to avoid the first-person singular ("I") as a design principle (e.g., Apple's Siri "Avoid using personal pronouns"; [1]), current LLMs do use "I" and "me" (e.g., "Tell me what's on your mind"). While limited, some related work has explored introducing the first-person "I" pronouns. For example, [39] manipulated whether a conversational agent used more first-person singular ("I") or second-person singular ("you") pronouns; they found that users preferred statements oriented toward themselves (i.e., that used the second-person). The current study explores the effect of the system using different nouns – "I" vs. "the system" – in modulating degrees of anthropomorphism and trust, addressing a gap in our scientific knowledge.

2.2 Trust and Anthropomorphism

Trust is "a multifaceted concept that can refer to the belief that another will behave with benevolence, integrity, predictability, or competence" [50]. As expected, the accuracy (either apparent or real) [53] and credibility [29] of the system improve users' trust and reliance on it. At the same time, attributes unrelated to performance shape metrics of user trust with avatars, autonomous vehicles, robots, voice assistants, and other conversational agents [23, 27, 40, 41, 45, 50] (but see [42]). For example, [48] found that a robot that had static eye color (more human-like) was more trustworthy than a robot whose eye color changed according to its status (e.g., green when speaking; blue when it recognized a person). Other work has examined how presence or absence of a human-like face shapes trust [6, 22]. For example, [22] found users had greater trust for a computer when it was presented with a more human-like face.

As mentioned, the presence of a system's "voice" has also shown to increase both trust and anthropomorphism in related work, relative to no-voice or text-only conditions [31, 37, 50]. Relatedly, other work has shown that the properties of a system's voice shape

trust: [47] found that participants invested more money in a system that used a Southern Standard British English (SSBE) TTS voice compared to TTS voices that used regionally specific dialects (e.g., Liverpool, Birmingham), but only in the condition where the system gave participants more favorable returns, suggesting that anthropomorphic features interact to shape trust.

3 METHODS

The driving research question for this experiment was to understand the influence of implicit cues within language and in system feedback design on the extent to which users anthropomorphize the LLM system and trust its outcomes. In large part, our goal was to inform design principles around anthropomorphism for the responsible creation of LLM-powered conversational experiences.

3.1 Recruitment and Participants (n=2,170)

Participants were recruited from a pool of adults from the United States, using a Qualtrics Panel to ensure the participants in each condition were balanced by age, geographic region, and race/ethnicity (see Appendix A for demographic information). To be eligible for the study, all participants had to pass a technical qualifier to ensure that they could hear the sound files³, have prior experience with voice technology, have normal hearing (or be wearing a hearing aid), and have normal vision (or be wearing corrective lenses).

A total of 2,170 US English speaking adults completed the experiment. Three participants were removed as they did not enter a valid age and two participants were removed due to technical issues (n=1 background music, n=1 removed their headphones during the trials and did not hear the audio). Retained participants included data for n=2,165, all of whom indicated they either had normal vision (n=1,300) or that they were currently wearing corrective glasses or contacts (n=865). Participants indicated that they had normal hearing (n=2,100) or were currently wearing hearing aids (n=65).

Prior to data collection, the survey was reviewed by the Google Ethics Board. The following safety measures were taken to ensure the protection of participants: all participants were given the option to opt out of any sensitive questions (e.g., age, gender, race/ethnicity) and completed informed consent. No identifying information (e.g., name, address, etc.) was collected.

3.2 Procedure

An overview of the experiment is provided in Figure 1. Participants were randomly assigned to one of the four conditions, and were given instructions for the experimental trials (order randomized). On each trial, they asked the system a predefined question and rated the system's response (see Appendix B for full list of stimuli and system responses).

After the system provided a response, participants were given time to review the response in the modality of the system. Participants who experienced the chat (text) modality progressed through the following steps (see Appendix C for a schematic): (A) Participants were asked to click an "Ask" button to query the system. (B) A "submit" sound was used to confirm that the question had been asked after 100ms. (C) Following a 500ms delay, the question

was presented in a speech bubble indicating a user-asked question. The system's response was buffered with a progress cue, in which the user was shown "..." in the system's speech bubble for 1500ms. (D) After 2 seconds, the system's response was displayed in the system's speech bubble.

Participants who experienced the chat (text) and voice (TTS) modality progressed through steps 1 through 4, with the distinguishing factor being that the system's response was vocalized using a Google Studio AI voice⁴ (US-English Studio-O, female; 65 dB normalized in Praat) at step 4 in addition to the speech bubble.

In both cases, the system's responses were identical across all participants, while the initial frame varied according to the first-person noun ("I") or third-person noun ("the system") conditions applicable to them ("Here's what [I | the system] found"). After each trial or "interaction", participants were asked to rate the system's response for three dimensions related to trust:

- (1) **Perceived Accuracy:** "How accurate is the system's response?" Not at all accurate < Somewhat accurate < Mostly accurate < Completely accurate
- (2) **Perceived Risk:** "If you were relying on this information, how risky would it be if the system got this information wrong?" Not at all risky < Neither risky nor unrisky < Somewhat risky < Extremely risky
- (3) **Follow-up Validation:** "Would you validate the system's answer with other sources?" No < Maybe < Yes

Once participants completed all twenty trials, they were then asked to perform a set of post-trial assessments to rate the system as a whole.

- **Trustworthiness:** Rate the overall trustworthiness of the system (5 point scale) (Untrustworthy / Trustworthy).
- **Anthropomorphism:** Rate the system on five attributes adapted from the Godspeed Questionnaire [5] (all 5 point scales): Authenticity, Humanism, Awareness, Realism, and Competence, where each property is out of a possible score of 5 (maximum possible score = 25) (see Appendix D).

Finally, participants completed a listening comprehension question⁵ to ensure that participants in the chat + voice condition were listening to the audio throughout the experiment. All participants correctly answered the question.

4 ANALYSIS AND RESULTS

4.1 Overall Anthropomorphism by Condition

We modeled Anthropomorphism Score with a linear regression, with the effects of Modality (speech + text, text only), Grammatical Person (first-person ("I"), third-person ("the system")), and their interaction. Factors were sum coded. The model showed no effect of Grammatical Person (first vs. third-person) on anthropomorphism score. As seen in Figure 2, there was an effect of Modality, wherein participants' anthropomorphism scores were higher when the system responded with both speech and text [Coef = 0.36, $t = 3.85$, $p < 0.001$]. There was no interaction between Grammatical Person and Modality.

³Nancy had considered the sleeves.", generated with a TTS voice (US-English male; "Miles" from TTS Maker; 65 dB); options: "sleeves", "team", "sleds".

⁴For an audio illustration, see <https://cloud.google.com/text-to-speech/docs/voices>.

⁵They had a problem with the cliff" in a US-English male voice ("Miles" from TTS Maker; 65 dB) and identified the target word ("cliff", "slip", "trip").

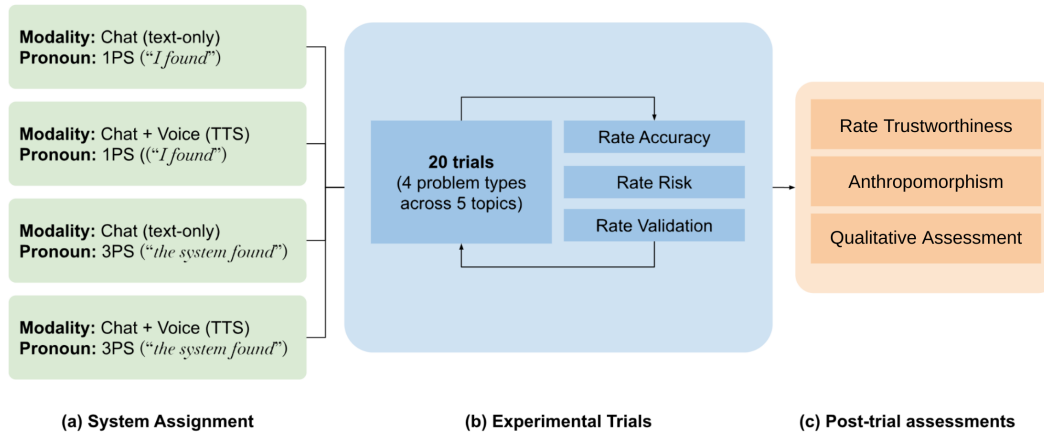


Figure 1: Procedure for the current study. All participants underwent a technical qualifier and comprehension assessment question. (a) System Assignment: First, all qualifying participants were assigned one of four conditional systems in our experiment, controlled for randomization and even distributions. (b) Experimental trials: Participants were asked to complete 20 trials. After each experimental trial, participants rated accuracy, risk, and validation. (c) Post-experimental trial assessments: After completing all experimental trials, participants rated their overall perceptions of the system, including trustworthiness rating, an Anthropomorphism questionnaire, and a qualitative assessment. In this study, we present the results of anthropomorphism and trust assessments.




Figure 2: Mean anthropomorphism score (out of a total of 25) across Modality (speech + text = orange, text only = dark blue) and Grammatical Person ("I found", "the system found"). Error bars indicate standard error of the mean.

4.2 Overall Trustworthiness Rating of System

4.2.1 Trustworthiness by Anthropomorphic Manipulations. We modeled participants' overall trustworthiness rating of the system with an ordinal regression with the ordinal R package [9], with the predictors of Modality (speech + text, text only; sum coded), Grammatical Person (first-person, third-person, sum coded), and

their interaction. The model showed no effects of Modality or Grammatical Person, or interaction.

4.2.2 Trustworthiness by Anthropomorphism Score. We modeled trustworthiness rating with a second ordinal regression to test the overall relationship with Anthropomorphism Score (centered). Results showed individuals with higher anthropomorphism scores rated the system as being more trustworthy [Coef = 0.26, $z = 23.35$, $p < 0.001$] (see Appendix E for a visualization).

4.3 Experimental Trials: Dimensions of Trust

We modeled participants' Likert responses for the three dimensions of trust in the information the system provided (1. Perceived accuracy, 2. Perceived risk, 3. Follow-up validation) with separate ordinal regressions. In each model, fixed effects included Modality (speech + text, text only), Grammatical Person (first-person, third-person), and their interaction. We additionally included predictors of Scenario (health, career, medications, travel, and cooking). Factors were sum coded.

4.3.1 Information Accuracy. The full model output for accuracy ("How accurate is the system's response?") is provided in Appendix F. The model showed an effect of Modality: participants who heard the TTS voice gave higher accuracy ratings ($p < 0.001$), as seen in Figure 3. There was no simple effect of Grammatical Person. There were effects of Scenario: perceived accuracy ratings were higher for the system's responses for health and career, but lower for medication and travel (all $p < 0.001$). Interactions between Modality and Scenario revealed reduced accuracy for medication and career responses (both $p < 0.001$), but higher accuracy for health in the "speech + text" condition ($p < 0.05$). Two interactions between Grammatical Person and Scenario were observed: when using "I", the

system’s responses were rated as more accurate for medication but less accurate for career (both $p < 0.001$). No other interactions were observed.

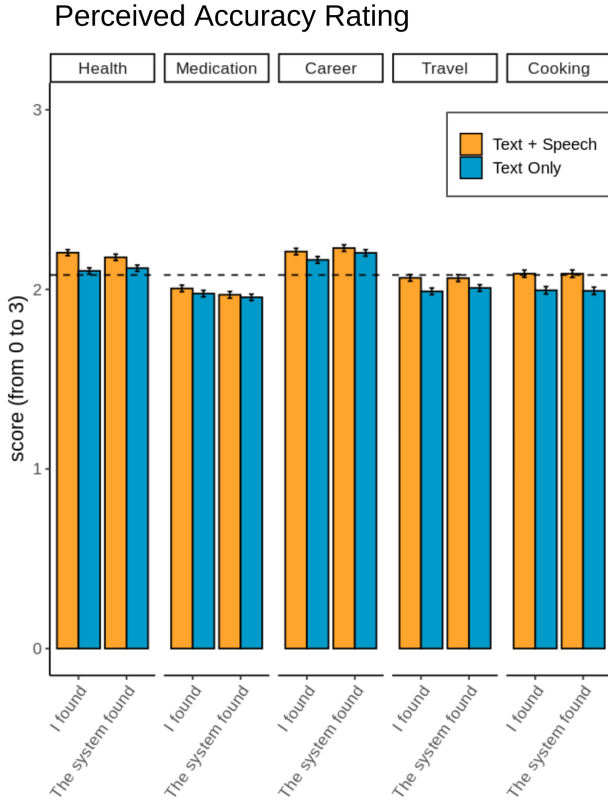


Figure 3: Mean ratings for “How accurate is the system’s response?” coded as numeric data (0= “Not at all accurate”, 1= “Somewhat accurate”, 2= “Mostly accurate”, 3= “Completely accurate”) across question contexts. Note that analyses were conducted on the Likert ordinal data and the numeric data is used for visualization purposes only. The experimental conditions included 1) Modality (speech + text = orange, text only = dark blue), and 2) Grammatical Person (“I found”, “the system found”). Error bars indicate standard error of the mean.

4.3.2 Perceived Risk. The full model output for perceived risk (“If you were relying on this information, how risky would it be if the system got this information wrong?”) is provided in Appendix G. The model showed no simple effects of Modality or Grammatical Person. Scenario was a significant predictor: participants indicated responses about health and medication were riskier to rely on if the system had incorrect information, on average, but scenarios in career and travel were less risky (all $p < 0.001$). Two interactions between Person and Scenario were observed: the systems that used “I” were rated as less risky when giving information about medications, but riskier for careers (both $p < 0.01$). No other effects or interactions were observed.

4.3.3 Follow-up Validation. The full model output for validation (“Would you validate the system’s answer with other sources?”) is provided in Appendix H. The model showed no simple effects of Modality or Grammatical Person. Yet, there were effects for Scenario. Participants indicated they would be more likely to validate the information given for questions about health and medication, but less likely to validate for responses about travel and careers (all $p < 0.001$). One three-way interaction was observed, where participants were more likely to validate responses for travel if the system used “I” and speech + text ($p < 0.05$). No other effects were observed.

5 DISCUSSION

This study investigated the role of two linguistic anthropomorphic cues, the presence of a voice and use of the first-person (“I”) by the system, on dimensions of trust in information provided by a pseudo-LLM system. We believe that this is an important step in understanding how people perceive and trust information from LLM models. In the following sections, we describe the five main contributions of this experiment.

5.1 Linguistic Cues of Anthropomorphism

First, we found that participants’ overall anthropomorphism scores were higher when the system responded with both text and speech (here a TTS voice), suggesting voice-based effects are not limited to naturally recorded human voices. These findings parallel those in the visual domain, where more human-like qualities tend to increase trust [23, 27, 41, 45, 50] (but see also [30]). However, contrary to our prediction, we did not observe effects of the grammatical person (“I” vs. “the system”) on users’ overall anthropomorphism scores.

5.2 A Computer Generated Voice Increases Perceived Accuracy of Information

Second, we found that the presence of a TTS voice shapes perceived accuracy of the system’s response. Relative to a text-only interface, participants believed the information a system gave was more accurate when they also heard the system talk. This effect is robust even in the absence of explicit human-like features, such as an image. This finding is consistent with related work showing that the presence of a pre-recorded human voice (relative to no voice) with an autonomous vehicle system increases both anthropomorphism and trust [50].

5.3 First-person (“I”) Increases Dimensions of Trust in Some Contexts

Third, we also found that, in limited contexts, the first-person “I” shaped dimensions of trust. When the system used “I”, its responses were rated as more accurate and less risky for medications; while speculative, these increases could reflect a perception of the system as a domain expert, particularly as the responses contained detailed medication information. Conversely, career information was rated as less accurate and riskier to rely on with “I”, raising the possibility that the first-person pronoun could increase perceived subjectivity. Unlike adding a voice, however, introducing “I” had more limited effects in general. One possibility is that participants might have

ignored the repeated preamble in each response (“Here’s what [I | the system] found”), whereas the voice was consistent throughout the system’s entire response.

5.4 Dimensions of Trust are not Equally Affected

While we observed an impact of the presence of a voice had an across-the-board effect on trial-level ratings of accuracy, we did not observe parallel effects on perceived risk and likelihood to validate. Furthermore, we did not observe that either the modality or person manipulations shaped overall ratings of the system’s trustworthiness, paralleling findings of equal trustworthiness ratings for content generated by humans compared to an LLM [25]. Additionally, we explored several question scenarios that could shape trust. While health and career information were rated as more accurate, medications and travel information were rated as less accurate overall. Participants indicated that systems’ responses about health and medications as riskier and ones they were more likely to externally validate. Participants’ reticence to rely on medical information given by an LLM reflects the gains needed to improve mutual trust [15].

5.5 Implications for UX Responsible Design

Our findings suggest that people believe information is more accurate and less risky when presented with anthropomorphic cues, which could lead to downstream harms if the system produces non-accurate data [54] or stereotypes [2, 13, 19]. We therefore make the following design recommendations:

- Using a voice with an LLM should be limited to cases where the model’s confidence in information accuracy is high. In cases where this cannot be avoided, we suggest introducing cues of speaker uncertainty in the auditory signal [46] and source attribution [15].
- Including a generated voice can improve trust and information uptake, which can be used for users’ benefit, such as adherence to a treatment plan in healthcare contexts.
- Consider alternatives to first-person pronouns (“I”) in LLMs. While we did not observe a strong overall effect of “I” as for the presence of a voice, it still increased trust in information given about medications.

5.6 Limitations and Future Directions

While adding the TTS voice had a positive effect on anthropomorphism in the current study, TTS voices can lower ratings of human-likeness, particularly if they are compared to other TTS voices varying in roboticism [12, 55] or directly to a human voice [10, 21]. Future work directly comparing types of voices can further elucidate these factors for LLMs, as well as reveal any potential uncanniness effects [30]. Additionally, the current study used a female TTS voice, but future research with multiple TTS voices (e.g., varying in gender) can more fully explore the social factors shaping anthropomorphism of LLMs, including probing how their perception might vary across users (e.g., by age, gender, race/ethnicity). Moreover, this study controlled both the question contexts, as well as the exact wording participants “asked”. While this affords experimental control, additional research exploring the effects of a

voice and first-person (“I”) will be needed in more naturalistic types of tasks, with real user intents. Finally, our focus was on English in the United States; future work examining other languages and cultural contexts is needed to determine whether effects of anthropomorphic cues and trust for LLMs generalize cross-linguistically and cross-culturally.

REFERENCES

- [1] [n. d.]. <https://developer.apple.com/design/human-interface-guidelines/siri/overview/editorial-guidelines>
- [2] Alberto Acerbi and Joseph M Stubbersfield. 2023. Large language models show human-like content biases in transmission chain experiments. *Proceedings of the National Academy of Sciences* 120, 44 (2023), e2313790120.
- [3] Tawfiq Ammari, Jofish Kaye, Janice Y Tsai, and Frank Bentley. 2019. Music, search, and IoT: How people (really) use voice assistants. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 3 (2019), 1–28.
- [4] Sisith Ariyaratne, Karthikeyan P Iyengar, Neha Nischal, Napa Chitti Babu, and Rajesh Botchu. 2023. A comparison of ChatGPT-generated articles with human-written articles. *Skeletal Radiology* (2023), 1–4.
- [5] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* 1, 1 (Jan. 2009), 71–81. <https://doi.org/10.1007/s12369-008-0001-3>
- [6] Amy L Baylor and Jeeheon Ryu. 2003. The effects of image and animation in enhancing pedagogical agent persona. *Journal of Educational Computing Research* 28, 4 (2003), 373–394.
- [7] 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*. 610–623.
- [8] Cesar Cadenas. 2023. Google bard can now speak loud and clear as update introduces speech feature. <https://www.techradar.com/computing/artificial-intelligence/google-bard-can-now-speak-loud-and-clear-as-update-introduces-speech-feature>
- [9] Rune Haubo B. Christensen. 2015. Analysis of ordinal data with cumulative link models—estimation with the R-package ordinal. *R-package version* 28 (2015), 406. https://mran.revolutionanalytics.com/snapshot/2014-09-12/web/packages/ordinal/vignettes/clm_intro.pdf
- [10] Michelle Cohn, Bruno Ferenc Segedin, and Georgia Zellou. 2022. Acoustic-phonetic properties of Siri- and human-directed speech. *Journal of Phonetics* 90 (Jan. 2022), 101123. <https://doi.org/10.1016/j.wocn.2021.101123>
- [11] Michelle Cohn and Georgia Zellou. 2020. Perception of concatenative vs. neural text-to-speech (TTS): Differences in intelligibility in noise and language attitudes. In *Proceedings of Interspeech*.
- [12] Benjamin R. Cowan, Holly P. Branigan, Mateo Obregón, Enas Bugis, and Russell Beale. 2015. Voice anthropomorphism, interlocutor modelling and alignment effects on syntactic choices in human-computer dialogue. *International Journal of Human-Computer Studies* 83 (2015), 27–42. <https://doi.org/10.1016/j.ijhcs.2015.05.008>
- [13] Sunipa Dev, Tao Li, Jeff M Phillips, and Vivek Srikumar. 2020. On measuring and mitigating biased inferences of word embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 7659–7666.
- [14] Mateusz Dubiel, Sylvain Daronnat, and Luis A Leiva. 2022. Conversational Agents Trust Calibration: A User-Centred Perspective to Design. In *Proceedings of the 4th Conference on Conversational User Interfaces*. 1–6.
- [15] Adam G Dunn, Ivy Shih, Julie Ayre, and Heiko Spallek. 2023. What generative AI means for trust in health communications. *Journal of Communication in Healthcare* 16, 4 (2023), 385–388.
- [16] Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864.
- [17] Claus-Peter Ernst and Nils Herm-Stapelberg. 2020. Gender Stereotyping’s Influence on the Perceived Competence of Siri and Co.. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*. 4448–4453.
- [18] John Andrew Fisher. 1991. Disambiguating anthropomorphism: An interdisciplinary review. *Perspectives in ethology* 9, 9 (1991), 49–85.
- [19] Vinitha Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. 2023. “I wouldn’t say offensive but...”: Disability-Centered Perspectives on Large Language Models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 205–216.
- [20] Catherine A Gao, Frederick M Howard, Nikolay S Markov, Emma C Dyer, Siddhi Ramesh, Yuan Luo, and Alexander T Pearson. 2023. Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers. *NPJ Digital Medicine* 6, 1 (2023), 75.

- [21] Iona Gessinger, Bernd Möbius, Sébastien Le Maguer, Eran Raveh, and Ingmar Steiner. 2021. Phonetic accommodation in interaction with a virtual language learning tutor: A Wizard-of-Oz study. *Journal of Phonetics* 86 (2021), 101029.
- [22] Li Gong. 2008. How social is social responses to computers? The function of the degree of anthropomorphism in computer representations. *Computers in Human Behavior* 24, 4 (2008), 1494–1509.
- [23] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. 2011. A meta-analysis of factors affecting trust in human-robot interaction. *Human factors* 53, 5 (2011), 517–527.
- [24] Nicole Holliday. 2023. Siri, you've changed! Acoustic properties and racialized judgments of voice assistants. *Frontiers in Communication* 8 (2023), 1116955.
- [25] Martin Huschens, Martin Briesch, Dominik Sobania, and Franz Rothlauf. 2023. Do You Trust ChatGPT?—Perceived Credibility of Human and AI-Generated Content. *arXiv preprint arXiv:2309.02524* (2023).
- [26] Carlos Jensen, Shelly D Farnham, Steven M Drucker, and Peter Kollock. 2000. The effect of communication modality on cooperation in online environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 470–477.
- [27] Youjeong Kim and S Shyam Sundar. 2012. Anthropomorphism of computers: Is it mindful or mindless? *Computers in Human Behavior* 28, 1 (2012), 241–250.
- [28] Kwan Min Lee. 2008. *Media Equation Theory*. Vol. 1. John Wiley and Sons, Ltd, Malden, MA, USA, 1–4. <https://doi.org/10.1002/9781405186407.wbiecm035>
- [29] Miriam J Metzger and Andrew J Flanagin. 2013. Credibility and trust of information in online environments: The use of cognitive heuristics. *Journal of pragmatics* 59 (2013), 210–220.
- [30] Masahiro Mori, Karl F MacDorman, and Norri Kageki. 2012. The uncanny valley [from the field]. *IEEE Robotics & automation magazine* 19, 2 (2012), 98–100.
- [31] Sara Moussawi, Marios Koufaris, and Raquel Benbunan-Fich. 2021. How perceptions of intelligence and anthropomorphism affect adoption of personal intelligent agents. *Electronic Markets* 31 (2021), 343–364.
- [32] Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues* 56, 1 (2000), 81–103.
- [33] Clifford Nass, Youngme Moon, and Paul Carney. 1999. Are people polite to computers? Responses to computer-based interviewing systems 1. *Journal of applied social psychology* 29, 5 (1999), 1093–1109.
- [34] Clifford Nass, Youngme Moon, John Morkes, Eun-Young Kim, and B. J. Fogg. 1997. Computers are social actors: A review of current research. *Human values and the design of computer technology* 72 (1997), 137–162.
- [35] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, Boston, Massachusetts, USA, 72–78. <https://doi.org/10.1145/259963.260288>
- [36] Amanda Purington, Jessie G. Taft, Shruti Sannon, Natalya N. Bazarova, and Samuel Hardman Taylor. 2017. “Alexa is My New BFF”: Social Roles, User Satisfaction, and Personification of the Amazon Echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, Denver, Colorado, USA, 2853–2859. <https://doi.org/10.1145/3027063.3053246>
- [37] Lingyun Qiu and Izak Benbasat. 2005. Online consumer trust and live help interfaces: The effects of text-to-speech voice and three-dimensional avatars. *International journal of human-computer interaction* 19, 1 (2005), 75–94.
- [38] Lingyun Qiu and Izak Benbasat. 2009. Evaluating anthropomorphic product recommendation agents: A social relationship perspective to designing information systems. *Journal of management information systems* 25, 4 (2009), 145–182.
- [39] Jianhong Qu, Ronggang Zhou, and Zhe Chen. 2022. The effect of personal pronouns on users and the social role of conversational agents. *Behaviour & Information Technology* 41, 16 (2022), 3470–3486.
- [40] Minjin Rheu, Ji Youn Shin, Wei Peng, and Jina Huh-Yoo. 2021. Systematic review: Trust-building factors and implications for conversational agent design. *International Journal of Human-Computer Interaction* 37, 1 (2021), 81–96.
- [41] Eileen Roesler, Dietrich Manzey, and Linda Onnasch. 2021. A meta-analysis on the effectiveness of anthropomorphism in human-robot interaction. *Science Robotics* 6, 58 (2021), eabj5425.
- [42] Eileen Roesler, Meret Vollmann, Dietrich Manzey, and Linda Onnasch. 2024. The dynamics of human-robot trust attitude and behavior—Exploring the effects of anthropomorphism and type of failure. *Computers in Human Behavior* 150 (2024), 108008.
- [43] Mustafa Safdari, Greg Serapio-Garcia, Clément Crepy, Stephen Fitz, Peter Romero, Luning Sun, Marwa Abdulhai, Aleksandra Faust, and Maja Mataric. 2023. Personality traits in large language models. *arXiv preprint arXiv:2307.00184* (2023).
- [44] Sivan Schwartz, Avi Yaeli, and Segev Shlomov. 2023. Enhancing Trust in LLM-Based AI Automation Agents: New Considerations and Future Challenges. *arXiv preprint arXiv:2308.05391* (2023).
- [45] William Seymour and Max Van Kleek. 2021. Exploring interactions between trust, anthropomorphism, and relationship development in voice assistants. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–16.
- [46] Eva Székely, Joseph Mendelson, and Joakim Gustafson. 2017. Synthesising Uncertainty: The Interplay of Vocal Effort and Hesitation Disfluencies. In *INTER-SPEECH*. 804–808.
- [47] Ilaria Torre, Jeremy Goslin, Laurence White, and Debora Zanatto. 2018. Trust in artificial voices: A “congruency effect” of first impressions and behavioural experience. In *Proceedings of the Technology, Mind, and Society*. 1–6.
- [48] Michelle ME Van Pinxteren, Ruud WH Wetzels, Jessica Rüger, Mark Pluymaekers, and Martin Wetzels. 2019. Trust in humanoid robots: implications for services marketing. *Journal of Services Marketing* 33, 4 (2019), 507–518.
- [49] Adam Waytz, John Cacioppo, and Nicholas Epley. 2010. Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science* 5, 3 (2010), 219–232.
- [50] Adam Waytz, Joy Heafner, and Nicholas Epley. 2014. The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of experimental social psychology* 52 (2014), 113–117.
- [51] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).
- [52] Tianyu Wu, Shizhu He, Jingping Liu, Siqi Sun, Kang Liu, Qing-Long Han, and Yang Tang. 2023. A brief overview of ChatGPT: The history, status quo and potential future development. *IEEE/CAA Journal of Automatica Sinica* 10, 5 (2023), 1122–1136.
- [53] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–12.
- [54] Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. *Advances in neural information processing systems* 32 (2019).
- [55] Georgia Zellou, Michelle Cohn, and Bruno Ferenc Segedin. 2021. Age- and Gender-Related Differences in Speech Alignment Toward Humans and Voice-AI. *Frontiers in Communication* 5 (2021), 1–11. <https://doi.org/10.3389/fcomm.2020.600361>

A PARTICIPANT DEMOGRAPHICS

Table 1: Participant Distribution Across the Four Conditions

		First-Person Text only	First-Person Voice + Text	Third-Person Text only	Third-Person Voice + Text
Age	Mean (sd)	46.9 years (18.3)	46.7 years (18.4)	47.1 years (18.0)	47.1 years (18.5)
	Range	18-90	18-85	18-86	18-89
Gender	Women	280	275	285	281
	Men	261	261	257	238
	Another gender	3	3	0	1
Race/ethnicity	white	390	396	396	386
	Hispanic or Latino	96	92	97	99
	African American or Black	69	69	69	71
	American Indian or Alaska Native	12	10	13	15
	Asian American	18	24	24	25
	Hawaiian or Pacific Islander	6	2	2	3
	multiracial	49	38	38	40
	Total n	544	539	542	540

B USER QUESTIONS AND SYSTEM RESPONSES

Table 2: User Questions and System Responses across Question Contexts and Types

Question Context		Question Type	Question	System Response
Medication	You want to learn information about a new medication the doctor mentioned	Fact	What's the dosage for the medication, Levaquin?	Here's what [I the system] found: depending on the type of infection, the usual dose of Levaquin is 250–750 mg given once daily for 5–60 days.
		Risk	Does the medication Corgard have any dangerous side effects?	Here's what [I the system] found: side effects for Corgard include: numbness in your hands or feet, dizziness, feeling tired, and upset stomach.
		Advice	Should I take the medication Topamax if I'm already taking Tylenol?	Here's what [I the system] found: there are no interactions reported between Topamax and Tylenol. However, this does not necessarily mean no interactions exist. Always consult your healthcare provider.
		Estimation	What's the cost for the medication, Relenza?	Here's what [I the system] found: depending on the pharmacy you visit, the cost for Relenza inhalation powder could be around \$72 for a supply of 20 mg.
Health	You're interested in learning more about common health disorders	Fact	What are symptoms of stroke?	Here's what [I the system] found: common symptoms of stroke include sudden weakness on one side of the body, vision problems, and sudden difficulty speaking.
		Risk	Is cancer hereditary?	Here's what [I the system] found: up to 10% of all cancers may be caused by inherited genetic changes. It means that your risk of getting cancer is increased.
		Advice	What should I do to reduce my risk of diabetes?	Here's what [I the system] found: to reduce your risk of diabetes, you should eat a balanced diet and exercise regularly.
		Estimation	If my dad has heart disease, what are the chances I'll get it too?	Here's what [I the system] found: If one of your immediate family members has had a heart attack, a stroke, or was diagnosed with heart disease before the age of 60, this may indicate a family history of premature heart disease.
Career	You're interested in a career change.	Fact	How commonly do people switch jobs into healthcare?	Here's what [I the system] found: it is not as common for people to switch jobs into medicine as it is for other fields, such as computer science.
		Risk	What are the job-related health risks in finance?	Here's what [I the system] found: some risks in finance include chronic stress due to demanding workloads, tight deadlines, and making important financial decisions.
		Advice	Would I earn enough as an educational administrator?	Here's what [I the system] found: the average pay range for a Higher Education Administrator varies greatly (as much as \$55,816), which suggests there may be many opportunities for advancement.
		Estimation	How many years of training would I need to do to get into computer science?	Here's what [I the system] found: a bachelor's program in Computer Science generally takes four years to complete, while a master's degree may take an additional year or two years of coursework.
Travel	You're interested in learning more about the world.	Fact	Where is the world's tallest tree?	Here's what [I the system] found: the world's tallest tree is Hyperion, which is 344 feet tall and located in Redwood National Park in California in the United States.
		Risk	What are the risks in going in a cave?	Here's what [I the system] found: the most common causes of caving accidents include: falling, being struck by falling objects and hypothermia.
		Advice	Should I visit the world's tallest volcano on a vacation?	Here's what [I the system] found: Mauna Loa is a massive shield volcano. At 13,680 ft above sea level, hiking to its summit requires a good level of physical fitness and endurance.
		Estimation	How long would it take to reach the top of the world's tallest mountain?	Here's what [I the system] found: from base camp, it takes about 40 days to reach the summit of Mount Everest at 29,029 ft.
Cooking	You're interested in trying a new recipe.	Fact	What temperature do I need to cook chicken to?	Here's what [I the system] found: all poultry should reach a safe minimum internal temperature of 165 °F.
		Risk	Can I drink milk that's been expired for 2 days?	Here's what [I the system] found: as long as it's been stored properly, unopened milk generally stays good for 5–7 days past its listed date, while opened milk lasts at least 2–3 days past this date.
		Advice	If I have an allergy to peanuts, should I eat peanut oil?	Here's what [I the system] found: a peanut allergy is caused by an allergic reaction to the peanut protein. Peanut oil is typically safe because it's highly refined and has almost no detectable peanut protein.
		Estimation	How much is a dash of salt?	Here's what [I the system] found: a dash is somewhere between 1/16 and 1/8 teaspoon.

C TRIAL SCHEMATIC



Figure 4: First, participants were asked to click an “Ask” button to query the system. A “submit” sound was used to confirm that the question had been asked after 100ms. Following a 500ms delay, the question was presented in a speech bubble indicating a user-asked question. The system’s response was buffered with a progress cue, in which the user was shown “...” in the system’s speech bubble for 1500ms. After 2 seconds, the system’s response was in the system’s speech bubble. In the third-person condition, the response started with “Here’s what the system found”. In the first-person condition, the response started with “Here’s what I found”. In the text-only condition, participants saw the system’s typed response only. In the text + speech condition, they saw the typed response and heard a TTS voice reading the same response aloud. Note that there was a delay after displaying the system’s response before showing the three response options, where participants rated 1) perceived accuracy of the information, 2) perceived risk, and 3) follow-up validation.

D ANTHROPOMORPHISM QUESTIONNAIRE

(1) Please rate your impression of the system: (Fake – Natural)

- Completely fake
- Somewhat fake
- Neither fake nor natural
- Somewhat natural
- Completely natural

(2) Please rate your impression of the system: (Machine-like – Human-like)

- Completely machine-like
- Somewhat machine-like
- Neither machine-like nor human-like
- Somewhat human-like
- Completely human-like

(3) Please rate your impression of the system: (Unconscious – Unconscious)

- Completely unconscious
- Somewhat unconscious

- Neither unconscious nor conscious

- Somewhat conscious

- Completely conscious

(4) Please rate your impression of the system: (Artificial – Life-like)

- Completely artificial
- Somewhat artificial
- Neither artificial nor lifelike
- Somewhat lifelike
- Completely lifelike

(5) Please rate your impression of the system: (Incompetent – Competent)

- Completely incompetent
- Somewhat incompetent
- Neither incompetent nor competent
- Somewhat competent
- Completely competent

E TRUST AND ANTHROPOMORPHISM SCORE




Figure 5: Relationship between Trustworthiness Rating and Anthropomorphism score. Individual points represent group means at each score. (-2 = Completely untrustworthy, -1 = Somewhat untrustworthy, 0 = Neither untrustworthy or trustworthy, 1 = Somewhat trustworthy, 2 = Completely trustworthy). Error bars indicate standard error of the mean.

F PERCEIVED INFORMATION ACCURACY

Table 3: Ordinal Mixed Effects Model Output for Perceived Information Accuracy

	Coef	Std. Error	z value	p value
Modality(speech+text)	0.1	4.20e-03	22.95	<0.001
Person("I")	5.70e-04	4.3e-03	0.14	0.89
Scenario(health)	0.17	0.02	10.47	<0.001
Scenario(medication)	-0.32	4.5e-03	-71.46	<0.001
Scenario(career)	0.36	4.5e-03	79.73	<0.001
Scenario(travel)	-0.16	4.5e-03	-35.71	<0.001
Modality(speech):Person("I")	0.01	4.3e-03	1.2	0.23
Modality:Scenario(health)	0.04	0.02	2.25	0.02
Modality:Scenario(medication)	-0.05	4.5e-03	-12.1	<0.001
Modality:Scenario(career)	-0.03	4.5e-03	-6.14	<0.001
Modality:Scenario(travel)	4.20e-03	4.5e-03	0.95	0.34
Person:Scenario(health)	0.01	0.02	0.42	0.68
Person:Scenario(medication)	0.05	4.5e-03	10.68	<0.001
Person:Scenario(career)	-0.04	4.5e-03	-8.14	<0.001
Person:Scenario(travel)	-0.01	4.5e-03	-1.79	0.07
Modality:Person:Scenario(health)	0.01	0.02	0.85	0.4
Modality:Person:Scenario(medication)	1.90e-03	4.5e-03	0.43	0.67
Modality:Person:Scenario(career)	-1.40e-03	4.5e-03	-0.32	0.75
Modality:Person:Scenario(travel)	0.01	4.5e-03	1.29	0.2
<i>Threshold coefficients</i>				
Not at all accurate Somewhat accurate	-3.99	0.03	-156.26	
Somewhat accurate Mostly accurate	-1.56	2.50e-02	-365.98	
Mostly accurate Completely accurate	0.6	0.02	38.33	
<i>Random effects</i>				
Participant	Variance			
	2.44			

G PERCEIVED RISK MODEL

Table 4: Ordinal Mixed Effects Model Output for Perceived Risk

	<i>Coef</i>	<i>Std. Error</i>	<i>z value</i>	<i>p value</i>
Modality(speech+text)	-0.03	0.03	-1.12	0.26
Person("I")	-3.8e-03	0.03	-0.12	0.90
Scenario(health)	0.44	0.02	24.35	<0.001
Scenario(medication)	0.53	0.02	29.31	<0.001
Scenario(career)	-0.20	0.02	-10.46	<0.001
Scenario(travel)	-0.99	0.02	-52.95	<0.001
Modality(speech):Person("I")	-0.02	0.03	-0.60	0.55
Modality:Scenario(health)	-2.6e-03	0.02	-0.14	0.89
Modality:Scenario(medication)	1.1e-03	0.02	0.06	0.95
Modality:Scenario(career)	-1.0e-02	0.02	-0.53	0.59
Modality:Scenario(travel)	-0.01	0.02	-0.63	0.53
Person:Scenario(health)	0.02	0.02	1.07	0.29
Person:Scenario(medication)	-0.05	0.02	-2.84	<0.01
Person:Scenario(career)	0.05	0.02	2.92	<0.01
Person:Scenario(travel)	-0.02	0.02	-1.32	0.19
Modality:Person:Scenario(health)	0.01	0.02	0.46	0.64
Modality:Person:Scenario(medication)	0.01	0.02	0.58	0.56
Modality:Person:Scenario(career)	-3.6e-03	0.02	-0.19	0.85
Modality:Person:Scenario(travel)	-0.01	0.02	-0.53	0.60
<i>Threshold coefficients</i>				
Not at all risky Neither risky nor unrisky	-1.52	0.03	-46.82	
Neither risky nor unrisky Somewhat risky	-4.2e-03	0.03	-0.13	
Somewhat risky Extremely risky	1.70	0.03		
<i>Random effects</i>				
Participant	Variance			
	1.91			

H FOLLOW-UP VALIDATION MODEL

Table 5: Ordinal Mixed Effects Model Output for Follow-up Validation

	<i>Coef</i>	<i>Std. Error</i>	<i>z value</i>	<i>p value</i>
Modality(speech+text)	-0.03	0.04	-0.74	0.46
Person("I")	-0.03	0.04	-0.67	0.49
Scenario(health)	0.47	0.22	21.68	<0.001
Scenario(medication)	0.70	0.02	31.04	<0.001
Scenario(career)	-0.34	0.02	-16.15	<0.001
Scenario(travel)	-0.48	0.02	-23.33	<0.001
Modality(speech):Person("I")	6.7e-03	0.04	0.17	0.87
Modality:Scenario(health)	-9.3e-03	0.02	-0.43	0.67
Modality:Scenario(medication)	0.01	0.02	0.49	0.62
Modality:Scenario(career)	0.02	0.02	0.65	0.52
Modality:Scenario(travel)	-0.01	0.02	-0.56	0.57
Person:Scenario(health)	0.04	0.02	1.89	0.06
Person:Scenario(medication)	9.1e-03	0.02	0.41	0.68
Person:Scenario(career)	-0.03	0.02	-1.66	0.10
Person:Scenario(travel)	-0.02	0.02	-1.15	0.25
Modality:Person:Scenario(health)	2.03-03	0.02	0.09	0.93
Modality:Person:Scenario(medication)	-0.01	0.02	-0.48	0.63
Modality:Person:Scenario(career)	-0.02	0.02	-1.03	0.30
Modality:Person:Scenario(travel)	0.04	0.02	2.01	0.04
<i>Threshold coefficients</i>				
No Maybe	-2.12	0.04	-49.88	
Maybe Yes	-0.33	0.04	-7.95	
<i>Random effects</i>				
Participant	Variance			
	3.36			