Additional Information: **Robustness study with belief distributions**
Author(s)
Submission
1) Have any data been collected for this study already?
No

2) What's the main question being asked or hypothesis being tested in this study?
We examine whether the process of transmitting information through speech distorts the information and how such distortions differ between different kinds of information content.
This experiment replicates our main design, but with a different approach to measuring beliefs based on distributions about state beliefs rather than point beliefs about the state of the world. This also means that we incentivize transmitters using this new belief elicitation.

3) Describe the key dependent variable(s) specifying how they will be measured.

Our experiments involve transmission of information about two unknown states: (i) the change in home price growth in a US city and (ii) the change in revenue growth of a US retailer, both for the upcoming year.
Each state has three key dependent variables associated with it: the respondent's belief movement about the unknown state as measured using the full stated belief distribution (difference between posterior mean, elicited after the respondent listens to a message about the state, and prior mean, elicited beforehand); the respondent's belief about the message originator's prediction about the state of the world; and the respondent's belief about the originator's reliability, as measured on a scale ranging from 0 (extremely unreliable) to 100 (extremely reliable).
For the state beliefs, we elicit the whole distribution of beliefs using a set of pre-specified bins (from -10% to 10% in 2% increments).

For all three outcomes, we will z-score within topic*manipulation type quadrants and then pool together the observations from both topics.

Transmitters will be incentivized to record messages that induce shifts in listeners' belief distributions that are as close as possible to the shifts induced by the original messages. Specifically, they will be told that their probability of receiving the $20 bonus will be paid out according to the formula below: 100-5*(average change in probability assigned to bucket 1 based on your message, in percent - average change in probability assigned to bucket 1 induced by the original message, in percent)^2+(average change in probability assigned to bucket 2 based on your message, in percent - change in probability assigned to bucket 2 induced by the original message, in percent)^2+..., and so on. While this formula might look complicated, what it means is simple: you should try to accurately convey the full information content of the messages you listen to.

All three beliefs are incentivized through random incentives. Respondents will be told at the start that 1 in 10 respondents will be randomly chosen to be eligible for a bonus payment and have one of the incentivized tasks be paid out according to the formulas below.
(a) Beliefs about the true state of the world are incentivized with the following formula:

(i)Probability of winning $20 [in %] = 100 minus 5 times the sum, over all ranges, of (Your % for that range − Correct %)², where the “correct %” is 100% for the range that actually contains the true value and 0% for all other ranges. While this formula might look complicated, what it means is simple: to maximize your chance of winning the bonus, set each bar to the percentage chance you truly think that range will occur.

(b) Beliefs about the message originator’s beliefs and reliability are incentivized as follows:
(i) For 50% of respondents, there will be no incentive and the question will be phrased as a direct question about the message originator’s beliefs.
(ii) For 50% of respondents, the question will be phrased as a second-order question (“your job is to predict what people would on average respond to the direct question”) and responses will be incentivized with the following formula:
(1) For beliefs about the originator’s prediction: Probability of winning $20 [in %] = 100 − 10*(Response [in %] - Average response to direct question [in %])^2
(2) For beliefs about the originator’s reliability: Probability of winning $20 [in %] = 100 − 2*(Response - average response to the direct question)^2
If a respondent is selected to be eligible for a bonus, one of the incentivized beliefs of that respondent will be randomly selected to be the one that counts for payment.

4) How many and which conditions will participants be assigned to?

The full experiment comprises two separate data collections that build on each other, a transmitter experiment and a listener experiment. The two experimental collections rely on different respondent samples.

Transmitter experiment:
Participants listen to two short recordings played consecutively and without a break, each one of an opinion piece providing a qualitative narrative about the future path of a different economic variable. Then, they record their own summary of these recordings, separately for the first and second variable.

A randomly chosen 50% of transmitters will be asked their prior belief about each variable before hearing the recordings, and all transmitters will be asked for the three beliefs described above after recording their transmitted message for each topic.

Recording treatment arms:
Within each topic, we randomize three key features of the original recordings:
Level of variable: We randomize whether the piece argues for an increase or a decrease in the level of the variable.
Reliability of message: Second, we randomize the reliability of the original message. We randomly assign respondents to one of two different types of reliability manipulations:
● Naturalistic (combination of explicit statements about confidence, source quality and speaker competence, as well as implicit markers of reliability): Respondents in the naturalistic condition are assigned to one of the following 2 conditions: (i) Strong reliability; (ii) Weak reliability.
● Modular (Insertion of explicit markers indicating high or low reliability (e.g., definitely

vs. possibly, will vs. might, etc.): Respondents in the modular condition are assigned to one of the following 3 conditions: (i) Strong reliability; (ii) No reliability markers; (iii) Weak reliability.
Sex of transmitter voice: We randomize whether the recording is a male voice or a female voice. This is not a focus of analysis and we randomize simply for symmetry.
Randomization is stratified: each transmitter hears two recordings, one with an “increase” and one with a “decrease,” one with “strong reliability” and one with “weak reliability,” and one with a male voice and one with a female voice. Then, if exactly one of the two topics is in the modular condition, that topic has a 33% chance of getting switched to “no reliability markers.” If both topics are in the modular condition, there is a 66% chance that one of the two topics is randomly switched to “no reliability markers.”

Respondents receive incentives for transmitting all information contained in the original messages in a way that preserves the induced belief movements of listeners to those messages. Respondents are informed that 1 in 10 people will be selected for a bonus of up to $20. In particular, we tell our respondents that their task is to record a message that Induces average belief changes that are as close as possible to the average belief changes induced by the original message, measured over the full distribution of elicited beliefs; We explicitly explain to respondents that beliefs are measured using a distribution.

In order to do this, they should pass on anything from the original message that they think would be relevant for how people change their beliefs. They are told that if selected for the bonus, their voice message will be played to some other participants, and we will measure their belief changes after hearing the voice message. They are further told that the likelihood of receiving the bonus payment depends on how close the average belief change induced by their message within each bucket of the belief distribution elicitation is to the average belief change induced within the corresponding bucket by the original messages.

Listener experiment:

This involves a separate set of respondents. For each of the two topics, respondents first state their prior belief about the outcome variable of interest and then listen to a recording about the variable. As before, the order of the topics is randomized. For each topic, respondents are randomly matched to a transmitter and listen either to the same original recording as the one the transmitter heard, or that transmitter’s transmitted recording. There is a 30% chance of hearing the original and 70% chance of hearing the transmitted recording.

After listening to a recording, respondents are incentivized to forecast the future development of the variable as well as to guess the prediction of the message originator and the reliability of the original message. (The same 3 outcomes described above, incentivized in the same way).

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

Our three key outcome variables are belief movements about the state of the world (posterior minus prior), beliefs about the message originator’s prediction, and beliefs about the message originator’s reliability. Our main analyses compare the sensitivity of these

outcomes to the increase prediction/decrease prediction and strong reliability/weak reliability randomizations, between the original and transmitted recordings.

Specifically: to examine distortions of information about the level of a variable, we compare the following two quantities: (1) the difference between average beliefs among listeners who heard original recordings in the “increase” condition and average beliefs among listeners who heard original recordings in the “decrease” condition, and (2) the difference between average beliefs among listeners who heard transmitted recordings in the “increase” condition and average beliefs among listeners who heard transmitted recordings in the “decrease” condition. By “belief” here, we mean the belief about the message originator’s prediction, not the listener’s own belief about the state of the world (because the latter is a function of transmission of both level and reliability information).

Similarly, to examine distortions of information about the reliability of a message, we will conduct an identical comparison except using listeners’ beliefs about the reliability of the message originator and comparing between the strong reliability/weak reliability conditions. Finally, to examine distortions in downstream beliefs about the world, we will compare listeners’ average belief movements (posterior minus prior) between original and transmitted messages across our four key quadrants (decrease + strong reliability, decrease + weak reliability, increase + weak reliability, increase + strong reliability).

We will also examine the transcripts of transmitted messages. Using a large language model (such as GPT-5-2) we will classify each transcript according to whether it contains explicit statements about the original message’s level prediction and the original message’s reliability. We will repeat the first analysis above separately by whether the script contains an explicit level statement or not, and the second analysis separately by whether the script contains an explicit reliability statement or not.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

We will exclude responses below the 5th and above the 95th percentile of prior beliefs, taking these extreme responses as indications of either inattention or very outlandish and potentially distortionary beliefs about the subject matter. To deal with potential outliers in the beliefs about the state of the world and beliefs about the originator’s beliefs about the state of the world, we winsorize them at the 2nd and 98th percentile. Beliefs about the originator’s reliability are naturally bounded between 0 and 100, so we don’t need to winsorize them. We include an attention check and a set of comprehension questions on the study instructions. Subjects that fail the attention check are excluded from the study. If a subject fails to answer all comprehension questions correctly within the first two trials, we exclude that subject from the data collection. When moving from the transmitter experiment to the listener experiment, we exclude transmitted recordings that are below the 5th percentile of recording length or below the 5th percentile of transcript word count, as a proxy for recordings that are silent or otherwise completely lack content.

7) How many observations will be collected or what will determine sample size?

No need to justify decision, but be precise about exactly how the number will be determined.
In the transmitter experiment, we plan to collect 200 completes.
In the listener experiment, we plan to collect 800 respondents.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)