

Numbers Tell, Words Sell

Michael Thaler*, Mattie Toma[†], and Victor Yaneng Wang[‡]

November 5, 2024

[Click here for most recent version](#)

Abstract

When communicating numeric estimates with policymakers, journalists, or the general public, individuals must choose between using precise numbers or less-precise natural language. This project uses two sender-receiver experiments and finds that senders strategically use language to communicate numeric estimates to persuade receivers. Study 1 asks senders to communicate probabilities of abstract events to receivers on Prolific, and Study 2 asks academic researchers to communicate the effect sizes in economics papers to government policymakers. When senders are given incentives to directionally persuade receivers instead of incentives to communicate accurately, they are 25-29 percentage points more likely to choose to send a language message instead of a numeric message. Senders with incentives to persuade tend to slant language messages more than numeric messages, and this effect is driven by senders who prefer to send language messages. Our evidence suggests that senders are strategically using the imprecision of language to slant more, leveraging plausible deniability. Message receivers are persuaded by senders overall, particularly when senders communicate using language.

*University College London. Email: michael.thaler@ucl.ac.uk.

[†]University of Warwick. Email: mattie.toma@warwick.ac.uk.

[‡]MIT. Email: vyw20@mit.edu.

We are grateful to Johannes Abeler, Despoina Alempaki, Benjamin Enke, Thomas Graeber, Matthew Rabin, and seminar audiences at Berlin Behavioral Economics, Middlebury, Stanford Institute for Theoretical Economics, and Vienna Center for Experimental Economics for helpful comments and generous feedback. Deelan Maru at the Behavioural Insights Team, Eleanor Williams at the Australian Centre for Evaluation, and all participating researchers and policymakers generously contributed their time and insights. We would also like to thank the British Academy/Leverhulme Small Research Grants and the University of Warwick Behaviour Spotlight Small Grant Scheme and Primary Data Funds for generous funding support. The studies were approved by the ethics board at the University of Warwick (HSSREC 151.23-24, HSSREC 154.23-24, and HSSREC 219.23-24). Both studies were pre-registered on the AEA Social Science Registry (Study 1: [AEARCTR-0013287](#); Study 2: [AEARCTR-0013947](#).)

1 Introduction

How people communicate is essential to how information spreads. In many environments, such as empirical research, company reports, or policy white papers, a well-informed expert chooses how to communicate information to a less-informed audience. This communication can take varied forms. It can involve precise statements using numbers, such as “There is a *66% chance* that the event will occur” or “The treatment effect is *6 percentage points*.” It can also involve less-precise statements using natural language, such as “The event will *probably* occur.” or “The treatment effect is *substantial*.” Whether communicators choose to use relatively precise numbers or relatively imprecise language is highly context-dependent, and we study one particular mediating factor: whether the communicator’s incentives are aligned or misaligned with their audience.

In this paper, we present evidence that experts are more likely to use numeric messages if they want to communicate the truth, and are more likely to use language if they want to persuade people to update their beliefs in a particular direction. Our findings suggest that this main effect is driven by the relative imprecision of natural language compared with numeric messages. Individuals who have incentives to directionally persuade others face a tradeoff when using numbers to communicate: If they choose a number that does not represent the truth, they may find this psychologically costly (e.g. Serra-Garcia, Van Damme, and Potters 2011; Deversi, Ispano, and Schwardmann 2021).¹ But if they choose a number that does represent the truth, they will be less likely to effectively persuade. Imprecise messages, such as those using natural language, offers communicators a way out of this dilemma. Experts can slant language an intermediate amount in a way that gives them plausible deniability that they are being truthful, while still persuading others. For instance, consider an expert who knows that the true probability of an event is 61% but wants to persuade their audience that it is higher. They expect their audience to infer the truth is 66% if they communicate a numeric message that says “66%”, and uniformly between 56-76% in response to a language message that says “probable.” Experts may prefer to say “probable,” because they can plausibly claim that the message could have referred to the probability being 61%. That is, if the expert says “probable” and their audience guesses 66%, the expert gets the benefit of persuading while avoiding the psychological cost of overt deception.

We run two pre-registered experiments, each with a simple sender-receiver design, to shed light on this effect. We vary the incentives that senders face, and study the causal effect of incentive misalignment on whether senders prefer to use natural language or numeric likeli-

¹Importantly, this relies quite heavily on the message *content* being important outside of the particular interaction. Without considering the meaning of messages outside the interaction, numbers are also imprecise in pooling equilibria, such as in models like Crawford and Sobel (1982).

hoods. Study 1 analyzes an abstract setting in which senders communicate information about the probability that a red ball is drawn from a box to receivers. Study 2 uses a particular population of interest, analyzing how academic researchers communicate information about the treatment effect identified in a research study to government policymakers.

In each study, the sender is given information about the true state (the probability of drawing a red ball, or the treatment effect in the research). Senders then choose what message to send to receivers. This occurs in three stages. In the first stage, senders choose a *numeric message* to send, such as “66 percent” for probabilities or “6 percentage points” for treatment effects. In the second stage, senders choose a *language message* to send from a dropdown menu, such as “it is probable” for probabilities or “the effect was substantial” for treatment effects. Then, the third stage gives us our primary outcome of interest: senders choose between their numeric message (“66 percent”) or their language message (“it is probable”). Finally, receivers observe the message and predict the true probability or treatment effect. Senders’ monetary incentives are based on receivers’ answers, and we vary whether senders face *aligned incentives* to have the receivers predict the answer accurately, or face *directional incentives* for the receivers to give high (or low) predictions.

Our main finding is that directional incentives significantly increase senders’ likelihood of using language. In our data, when senders’ incentives are aligned, they only choose to communicate using language 14% and 13% of the time in Studies 1 and 2, respectively. By varying senders’ incentives, we see clear evidence supporting our main hypothesis: senders are 25 and 29 percentage points (both $p < 0.001$), respectively, more likely to communicate using language when they have directional incentives.

Our second finding considers the slant of the messages. We first find that senders slant their number messages in the direction of their incentives, as expected. When using numeric messages in Study 1, senders report a likelihood of drawing the red ball that is 10.7 (11.9) percentage points higher (lower) when they have directional-high (directional-low) incentives than when they have aligned incentives (both $p < 0.001$). When using numeric messages in Study 2, researchers report a treatment effect estimate that is 12% larger when they have directional-high incentives than when they have aligned incentives ($p = 0.004$). We next compare language to numeric messages. Building on a related literature that surveys individuals and asks them to map words into their perceived numeric equivalents (Mosteller and Youtz 1990; Ho et al. 2015; Dhimi and Mandel 2022; Ott 2021), we use language-to-number mappings elicited at the end of our experiment as a tool to examine the slant of language relative to numeric messages. We see that senders slant their messages even more when choosing a language message relative to a numeric message. This effect is driven entirely by senders who *prefer* to send the language message. Senders in Study 1 who prefer

to communicate using language slant their messages by an additional 4.5 (5.9) percentage points when using language under directional-high (low) incentives ($p = 0.098$ and $p = 0.037$, respectively), while researchers who prefer to use language in Study 2 slant their treatment effect estimates by an additional 64% when using language ($p = 0.044$). By looking at the distribution of responses, we see evidence that the majority of respondents do not misreport numbers at all (presumably due to an aversion to outright lying), but are willing to slant their language messages, helping to explain the relative preference for language messages under directional incentives.

Our third finding relates to the state that the senders are communicating about. When senders are communicating about a “bad state” (for instance, they are communicating a number that is low when they face directional-high incentives), we see that they slant their message relatively more in the direction of their incentives. Moreover, senders are more likely to use language when they are communicating bad states, providing further evidence that a preference for language is driven by senders wishing to distort their messages while still maintaining plausible deniability.

We use additional survey modules in Study 1 to identify two key predictors of a propensity to use language in response to incentives to persuade. First, we see that senders who self-report that they prefer to strategically use language are more likely to respond to incentives ($p = 0.001$), and this can explain much of the treatment effect we observe: we interpret this as indicating that the participants in our experiment indicate awareness of their preference to use language to persuade. Second, more numerate participants are more likely to strategically use language ($p = 0.011$), further suggesting that this behavior is driven by strategic sophisticates.

To further explore the mechanisms underlying senders’ behavior, in Study 1 we consider an additional treatment in which senders choose between numeric messages and interval messages, such as “between 15 and 25 percent.” Senders are 30 percentage points more likely to use intervals instead of numbers when they face directional incentives ($p < 0.001$). We also see similar results for the slant of interval messages as we do for language messages. Because intervals and numbers differ primarily in their relative precision, the similar findings point to imprecision as an important mechanism underlying strategic preferences for language. Intuitively, imprecision affords plausible deniability, allowing senders to “hide behind” and excuse their slanted communications under directional incentives.

Finally, we look at the messages’ effects on receivers. In Study 1, we see that when receivers are paired with a sender facing directional incentives their guess of the chance of drawing a red ball is slanted in the direction of the sender’s incentivized direction by 15 percentage points ($p < 0.001$). In Study 2, policymakers predict treatment effects that are 22% larger when paired with a directionally-incentivized researcher ($p = 0.042$). Receivers

are particularly persuaded by language messages. In other words, language messages make receivers give answers that are closer to the incentives of senders, and are further from the truth. In Study 1 we test whether an awareness of sender incentives mitigates this effect, and we see that receivers are still largely naive even when sender incentives are known. This is consistent with the more general finding in Cai and Wang (2006) that receivers in strategic communication games rely more on senders' messages than standard equilibrium models would predict.

Our work relates to a broader literature discussing language use and imprecise communication in other environments. For instance, Graeber, Noy, and Roth (2024b) examine how verbal communication distorts the transmission of economic information; and Graeber, Roth, and Zimmerman (2024) find that people's beliefs are more persistently affected by stories than statistics. Weiszsäcker (2023) discusses how communication can lead to imprecision and misunderstandings – this especially relates to our experimental findings of strong receiver naivety, where receivers often fail to account sufficiently for misaligned incentives when interpreting sender messages.

Importantly, there are many reasons unrelated to strategic incentives why senders may choose to communicate using language instead of numbers. For instance, receivers may misperceive numbers and probabilities in systematic ways (Tversky and Kahneman 1992; Prelec 1998). Language is also often seen as a more intuitive cognitive approach that provides important contextual information and avoids inducing math anxiety (Graeber, Noy, and Roth 2024a; Choe et al. 2019). In our environments, we find high levels of numeric communication when incentives are aligned; instead, we identify strategic incentives as an additional important mechanism that helps to explain preferences for language.

There have been recent calls for greater use of numeric communication in policy contexts, such as climate science and military intelligence (Mastrandrea et al. 2011; Ho et al. 2015; Dhimi and Mandel 2021; Hopster 2023). At the same time, relatively low rates of number use have been documented for academic research communications. For instance, Edlin and Love (2022) document the prevalence of numeric estimates in academic abstracts: 98% of abstracts in medicine report numbers, whereas only 37% in empirical economics report numbers.² This motivates our interrogation of evidence communications to policymakers in Study 2, which connects to an existing literature examining (biases in) how policymakers update their beliefs

²The high rates of reporting numbers in medicine can be attributed to the CONSORT guidelines, which provide a framework for the reporting of randomized controlled trials. The 2010 statement outlining these guidelines suggests awareness that precise numeric reporting can reduce the chance that the public is unduly influenced by researcher incentives: “To assess a trial accurately, readers of a published report need complete, clear, and transparent information on its methodology and findings....Explicit descriptions, not ambiguity or omission, best serve the interests of all readers” (Schulz, Altman, and Moher 2010).

and make policy decisions in response to research findings (DellaVigna, Kim, and Linos 2024; Hjort et al. 2021; Mehmood, Naseer, and Chen 2023; Toma and Bell 2024; Vivalt and Coville 2023). While there may be numerous reasons why researchers might use language to communicate — including to decrease the complexity of their message and to provide further context — our study provides evidence that in some cases researchers strategically leverage language to persuade.

Finally, our work relates to a large experimental literature on cheap-talk games where senders and receivers have misaligned incentives; see Crawford (1998) and Blume, Lai, and Lim (2020) for survey articles. Much of this literature finds more communication than Nash equilibria would predict (Dickhaut, McCabe, and Mukherji 1995; Cai and Wang 2006), consistent with our data. A common explanation for overcommunication in experiments is that senders face costs for lying (Gneezy 2005; Kartik 2009; Fischbacher and Föllmi-Heusi 2013; Abeler, Nosenzo, and Raymond 2019). However, there is also evidence that senders are more averse to direct lying than to vague communication or non-disclosure (Serra-Garcia, Van Damme, and Potters 2011; Deversi, Ispano, and Schwardmann 2021; Jin, Luca, and Martin 2022; Alempaki, Burdea, and Read 2023). We contribute to this literature by testing the hypothesis that many senders with misaligned incentives prefer imprecise language to precise numerical communication for this reason.³ This strategic use of verbal imprecision also relates to studies examining language communication as a device for “saving face” (Holtgraves and Perdeu 2016; Juanchich, Sirota, and Butler 2012). Relative to such studies we examine explicitly the sender’s choice to communicate in language vs numbers. In most cheap-talk experiments, senders are tasked with communicating numerical messages such as states of the world. However, a few studies vary the message space by allowing senders to communicate with language as well (Charness and Dufwenberg 2006; Wood 2022; Zhang and Bayer 2023). Relative to these existing cheap-talk studies with language, we add variation in incentives to our setup, allowing us to causally demonstrate the role of incentives on message format choices. We also extend this literature by showing, in Study 2, that these mechanisms extend outside of abstract lab environments.

The rest of the paper proceeds as follows: Section 2 describes the hypotheses, design, and results for Study 1, which explores our question in an abstract environment. Section 3 presents the same for Study 2, which extends our findings to researcher communications with policymakers. Finally, Section 4 discusses implications of these findings and highlights promising opportunities for future research.

³Lipman (2009) discusses how the meaning of words can be ill-defined and language use is hard to reconcile with game-theoretic predictions. While we can partially explain such behavior, we do not focus on the “vagueness” of language, instead exploring “imprecision,” i.e. its ability to refer to a range of values.

2 Study 1

2.1 Design Overview

Study 1 was run in April 2024 on Prolific.⁴ In the study, “senders” are paired with “receivers”. Senders are tasked with sending a message to communicate the chance, from 0% to 100%, that they will draw a red ball from a box of red and blue balls. Receivers don’t know the true probability of drawing a red ball and are incentivized to correctly guess this based on the message they receive from their sender.

The key source of variation in the design is the incentives senders face when communicating their message: one-third of senders are randomly assigned to face incentives to persuade their paired receiver that the chance of drawing a red ball is high (*directional-high*); one-third are randomly assigned to face incentives to persuade their paired receiver that the chance of drawing a red ball is low (*directional-low*); and one-third are randomly assigned to face incentives to accurately communicate the chance of drawing a red ball (*aligned*).

Half of senders are randomized into the main treatment arm, in which they face two choices about how to communicate their messages:

1. Message content: Senders choose the messages they would prefer to send using both numbers and language. For numbers, they complete the sentence: “The chance that you will draw a red ball is [X] percent.” For language, they complete the sentence: “It is [WORD] that you will draw a red ball.”
2. Message format: Senders indicate whether they would prefer that their numeric or language message is sent to their paired receiver. Their preferred message is the one that is more likely to be sent.

Half of senders are randomized into a supplementary treatment arm, where they choose between communicating numbers and intervals rather than numbers and language. For intervals, senders complete the sentence: “The chance that you will draw a red ball is between [Y] and [Z] percent.” where we fix $Z = Y + 10$, and Y to be multiples of five. In the main treatment arm comparing numbers and language, there may be other reasons particular to language why a sender might prefer one format to another. The purpose of this arm is to isolate the role that imprecision plays, by comparing two formats that are both numeric but vary in their precision.⁵

⁴Eligible participants needed to be United States adults, have completed at least 100 prior submissions on Prolific, and have had an approval rating of 90% or greater. In total, 1086 participants completed our study. As specified in our pre-analysis plan, we do not include the 86 subjects (8 percent) who failed an attention check.

⁵Note that in both cases, we restrict to a smaller message space for language and intervals, and this

Finally, we include a benchmarking exercise at the end of the study to elicit beliefs about how each word used in the experiment maps onto a corresponding numeric estimate. This allows us to measure the degree to which messages are slanted in the direction of senders' incentives in the experiment.

Hypotheses

As pre-registered, our primary hypothesis is that when senders face directional incentives to persuade receivers, they will prefer more imprecise message formats. That is, they are less likely to prefer communicating with numeric messages, and more likely to prefer communicating with language or interval messages. By communicating imprecisely, senders can maintain some plausible deniability about the extent to which they are slanting the truth when communicating to receivers.

Hypothesis 1. Imprecise communication under directional incentives:

- When choosing between numeric and language messages, senders are more likely to choose to send language messages when they face directional incentives.
- When choosing between numeric and interval messages, senders are more likely to choose to send interval messages when they face directional incentives.

Our remaining hypotheses are all pre-registered as secondary or exploratory hypotheses. Our second hypothesis is that senders slant their messages in a manner consistent with their incentives to directionally persuade. Relatedly, we expect that senders using imprecise messages will slant the truth more:

Hypothesis 2. Message slant and imprecise communication:

- First, we consider all message formats. We hypothesize that senders slant messages in the direction of their directional incentives, relative to the case with aligned incentives.
- When senders face directional incentives to persuade, they slant their chosen language/interval messages more than the numeric messages.
- When senders face directional incentives to persuade, those who preferred the language/interval format slant their chosen imprecise messages more than those who preferred the number format, compared with the numeric messages they chose.

Our third hypothesis explores the idea that even within an incentives condition there are sometimes greater incentives to slant the truth. For instance, suppose the sender is

coarsening also necessarily makes messages more imprecise.

directionally incentivized to make the receiver believe the true probability is high. If the sender then is in the “bad state” such that the true probability is low, we hypothesize that they will slant their message more to persuade the receiver. We then hypothesize that such senders are more likely to choose imprecise message formats.

Hypothesis 3. Precise good news and imprecise bad news:

- Among senders with directional incentives, senders slant messages more in the direction of their incentives when the true state is misaligned with their incentives.
- Among senders with directional incentives, senders use language and intervals more when the true state is misaligned with their incentives.

Our fourth hypothesis examines whether directionally-incentivized senders succeed in persuading receivers; and furthermore, whether the use of language or interval messages increases the extent of receiver persuasion.

This is also a test of the extent of receiver naivety. In theory, sophisticated receivers might be able to adjust the way they interpret slanted or imprecise messages to account for sender incentives, and thereby recover the true probability (Kartik 2009). Such sophisticated receivers would not be easily persuaded by the use of language or interval messages. In practice, however, receivers may not display this level of sophistication due to limitations such as bounded rationality. If so, directionally-incentivized senders will be more successful at persuading receivers.

Hypothesis 4. Receiver persuasion:

- Receivers are persuaded by senders. That is, receivers’ predictions move in the direction of the sender’s incentives on average.
- Receivers are more persuaded by directionally-incentivized senders when language and interval messages are sent than when numeric messages are sent.

We also consider whether receivers exhibit greater naivety when they do not know the specific incentives of their paired sender; that is, does information about sender incentives help to guide receivers’ responses?

2.2 Design for the Senders Experiment

Senders begin by reading the instructions for the roles of both sender and receiver. There are two main parts in the senders experiment: “Communications” and “Benchmarking”. 500 participants completed both parts: 251 were assigned to the numbers versus language

condition, and 249 were assigned to the numbers versus intervals condition. In addition to a \$3.60 payment for completion, 10 participants were randomly selected to earn up to \$50 as a bonus payment based on their response in a randomly-selected question in one of the two parts. Appendix Section B.1 shows screenshots from the experiment.

2.2.1 Communications

In the Communications part, senders make decisions for eight separate communications. They first learn about the chance, from 0% to 100% (randomized within-person), that the receiver will draw a red ball from a box of red and blue balls.⁶ This information was conveyed via a slider bar that shows the share of red balls out of 100. While it is easy to see from the slider approximately the true share of red balls, the slider intentionally does not report the exact number for two reasons. First, communicating the message in a format that involves neither numbers nor language ensures that senders are not anchored to one format over the other when selecting their own message. Second, the slight ambiguity provides scope for self-deception.

At the same time, senders are informed of their randomly-assigned incentive condition: directional-high, directional-low, or aligned.⁷ On the decision page, senders are told either “If this question is selected for payment, you will be more likely to earn the bonus if your Receiver predicts that the chance of drawing a RED ball is HIGH (LOW)” for directional-high (directional-low) incentives or “If this question is selected for payment, you will be more likely to earn the bonus if your Receiver’s prediction of the chance of drawing a RED ball is MORE ACCURATE” for aligned incentives. Incentives are randomly assigned within-person across the eight communication decisions.

After learning about the chance of drawing a red ball, senders make their message content choices for numbers and language (or numbers and intervals, depending on the condition). For numbers, they complete the sentence “The chance that you will draw a red ball is [X] percent” by selecting any integer for X between 0 and 100. For language, they complete the sentence “It is [WORD] that you will draw a red ball” by selecting a word from a dropdown list containing 13 words such as “improbable” and “almost certain”.⁸ We introduced a dropdown

⁶Probabilities took the following possible values: 2%, 8%, 17%, 25%, 33%, 42%, 50%, 58%, 67%, 75%, 83%, 92%, 98%.

⁷All incentives follow the binarized scoring rule (Hossain and Okui 2013; Vespa and Wilson 2016). Senders receive more details about how their bonus payments in the instructions; for instance for directional-high they learn that “the probability you will earn a bonus will be equal to the receiver’s probability of earning a bonus.”

⁸Appendix Table A1 shows the complete list of words used in the dropdown list. These are randomly presented either smallest to largest or largest to smallest, based on pilot data indicating the mapping of the words to numbers.

list rather than allowing open text responses to allow for a systematic mapping from words to numbers, which we describe in Section 2.2.2. For intervals, they complete the sentence “The chance that you will draw a red ball is between [X] and [Y] percent” by selecting X and Y. X is presented in units of 5, from 0 to 90, and Y is always 10 units greater than X such that the interval is always a fixed width. We imposed these constraints on the interval messages to ensure they were treated differently than the numeric messages; participants could not always precisely match the midpoint of their interval to their numeric message, for instance.

Finally, after making their message content choices, senders select their preferred message format, numbers or language (or numbers and intervals). Senders are informed that their preferred format “is more likely (but not guaranteed) to be the one communicated” to their paired receiver; in practice, we select their preferred format 50% of the time and randomly select a format 50% of the time. This element of chance ensures that all content choices are incentive compatible. That is, if we implemented the format preference with certainty, then the sender’s content choice for the message format that isn’t preferred would be merely hypothetical.

2.2.2 Benchmarking

At the end of the experiment, senders are asked to complete a benchmarking exercise in which we elicit their mapping of words onto numbers, using the words from the language dropdown list in the communications part. The mapping is incentivized using the procedure developed in Krupka and Weber (2013): participants are told “At the end of the survey, we will compare each of your answers to the average answers given by all other participants. If a question is randomly selected for payment, you will earn the bonus if your answer is within is within 3 percentage points of the average response.” This benchmarking exercise allows us to compare the slant of the numeric and language messages, as described in Section 2.5.

2.3 Design for the Receivers Experiment

Receivers similarly begin by reading the instructions for the roles of both sender and receiver. There are two main parts in the receivers experiment: “Predictions” and “Benchmarking”. The Benchmarking part is identical to the version used in the senders experiment. 500 participants completed both parts of the receivers experiment: 248 were assigned to the numbers versus language condition, and 252 were assigned to the numbers versus intervals condition. In addition to a \$3.70 payment for completion, 10 participants were randomly selected to earn up to \$50 as a bonus payment based on their response in a randomly-selected question in one of the two parts.

2.3.1 Predictions

In the Predictions part, receivers first receive the message from their paired sender about the chance of drawing a red ball.⁹ Then, receivers predict the chance of drawing a red ball. Receivers make eight predictions in total.¹⁰

Importantly, half of receivers are assigned to the *incentives-known* condition in which they are explicitly told their sender’s incentive assignment. The remaining half of receivers are assigned to the *incentives-unknown* condition, in which they are aware of the different incentive conditions that senders might face but do not know the condition to which their paired sender was assigned. This variation allows us to determine the degree to which receivers are able to anticipate and account for sender responses to incentives, which has important implications for policy.

After receivers make each prediction, they indicate on a Likert scale how informative they think the sender’s message is.

2.4 Empirical framework

Our primary specification tests whether senders are more likely to choose language messages under directional incentives, as per Hypothesis 1. (For senders who choose between interval and numeric messages, we use the equivalent specification for interval messages.) For sender i who learns that the true probability of a red ball is p , we estimate the following OLS equation:

$$Language_{ip} = \beta_0 + \beta_1 Directional_{ip} + \delta FE_i + \alpha FE_p + \epsilon_{ip} \quad (1)$$

where $Language_{ip}$ is an indicator equal to one when the sender chooses language (or intervals) over numbers, and $Directional_{ip}$ is an indicator equal to one when the sender faces directional rather than aligned incentives. We include individual fixed effects, FE_i , and fixed effects for each true probability of drawing a red ball, FE_p . We cluster standard errors at the individual level.

2.5 Senders’ Results

Figure 1 presents the raw data, where we see clear evidence that incentives affect senders’ choice of messages when they choose between natural language messages and numeric messages. In support of Hypothesis 1, we see that senders in the aligned condition choose language over


⁹Recall, in 50% of cases the sender’s message was communicated using their preferred and in 50% of cases we randomly selected a format. Receivers knew that the preferred format was not selected with certainty.

¹⁰We randomly pair the receiver with a new sender and message for each decision.

numbers 14 percent of the time (s.e. 2 pp). This share increases substantially in both the directional-high condition (to 35 percent, s.e. 3 pp) and the directional-low condition (to 40 percent, s.e. 3 pp). The differences between each directional condition and the aligned condition are highly statistically significant ($p < 0.001$ each).

We see similar patterns when senders choose between numbers and intervals. Relative to language, senders choose interval messages somewhat more often, and the comparisons across conditions are qualitatively similar. In the aligned condition, senders choose interval messages over numeric messages 22 percent of the time (s.e. 2 pp). This share increases substantially in both the directional-high condition (to 51 percent, s.e. 3 pp) and the directional-low condition (to 54 percent, s.e. 3 pp); and the differences between directional and aligned incentives are again highly statistically significant ($p < 0.001$ each).

Figure 1: Message Format Choices



(a) Numbers vs Language

(b) Numbers vs Intervals

Notes: This figure plots the raw data for the likelihood of choosing imprecise communications (language in Panel A and intervals in Panel B) over numbers, by incentive condition (aligned, directional-high, and directional-low). Bars reflect 95% confidence intervals. The sample includes 4,000 choices across 500 senders.

Table 1 shows the results from Equation 1. We see that overall directional incentives increase the likelihood of using language (intervals) to communicate by 25 (30) percentage points ($p < 0.001$).

Table 1: Format choice

	(1)	(2)	(3)	(4)
	Use Language	Use Language	Use Intervals	Use Intervals
Directional Incentives	0.250*** (0.025)		0.297*** (0.029)	
High Incentives		0.222*** (0.027)		0.288*** (0.030)
Low Incentives		0.279*** (0.028)		0.305*** (0.033)
Observations	2008	2008	1992	1992
Aligned Mean	.14	.14	.22	.22
Respondent FE	Yes	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of incentives on message format choice. In Columns 1 and 2, the dependent variable is an indicator equal to one when the sender chose to communicate using language rather than numbers. In Columns 3 and 4, the dependent variable is an indicator equal to one when the sender chose to communicate using intervals rather than numbers. Columns 1 and 2 are calculated over the sample of senders who had the option to communicate in language or numbers; Columns 3 and 4 are calculated over the sample of senders who had the option to communicate in intervals or numbers. In the *Directional Incentives* row, the independent variable is an indicator equal to one when the sender faced any kind of directional incentives, i.e. either directional-high or directional-low. In the *High Incentives* row, the independent variable is an indicator equal to one when the sender faced directional-high incentives. In the *Low Incentives* row, the independent variable is an indicator equal to one when the sender faced directional-low incentives. Fixed effects for the respondent and true probability are included as controls. Aligned mean calculates the likelihood of using language among aligned individuals in the corresponding sample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

2.5.1 Imprecise communication under directional incentives

There are many reasons why senders may choose different message formats when they have incentives to persuade. We now examine their within-format preferences to provide evidence that senders are strategically leveraging the imprecision of language to slant their messages.

First, we construct a measure of *slant* as follows: For numeric messages, slant is defined as ($number - true\ probability$). For interval messages, slant is defined as ($interval\ midpoint - true\ probability$). While numbers and numeric intervals have a natural meaning in relation to numeric probabilities, language messages require an extra conversion step. We use the

mapping from the benchmarking exercise in the senders experiment to create our measure of slant for language: (*language mapped to number* – *true probability*).¹¹

Consistent with the first part of Hypothesis 2, we see clear evidence that senders slant in the direction of their incentives on average. Appendix Table A2 shows that senders with directional-high incentives slant numeric messages upwards by 10.7 pp compared to senders with aligned incentives ($p < 0.001$), and senders with directional-low incentives slant numeric messages downwards by 11.9 pp compared to senders with aligned incentives ($p < 0.001$). By defining *directional slant* as being equal to *slant* for senders with directional-high incentives and equal to $-slant$ for senders with directional-low incentives, the average directional slant is 11.3 pp.

Testing the second part of Hypothesis 2, Appendix Table A3 looks at whether senders slant imprecise messages more than numeric messages when they face high versus low directional incentives. We see evidence supporting our hypothesis, but the effects are modest; the difference between imprecise and numeric messages is 2 percentage points larger for directional-high than directional-low incentives ($p = 0.003$).¹²

This average effect, however, belies a substantial heterogeneity. In Table 2, we run the same regressions as Appendix Table A3 but interact the gap in slant with the sender’s preferred format. Consistent with the last part of Hypothesis 2, when senders prefer to use language, they slant language messages significantly more relative to numeric messages: senders who prefer language slant their messages by an additional 4.5 percentage points under directional-high incentives ($p = 0.098$) and by an additional 5.9 percentage points under directional-low incentives ($p = 0.037$) when they use language. That is, senders appear to be opting to use language (or intervals, for which we see similar although somewhat muted effects) in exactly the cases that they are slanting their message more than they otherwise would with numbers—the imprecise communications seem to afford the degree of plausible deniability that might excuse a slanted message.

¹¹We use senders’ individual-level perceptions of the mapping from each word to number in our analysis. Results are little changed if we use average perceptions across all senders instead.

¹²The coefficient is similar when comparing language to numbers, and intervals to numbers. However, the noise generated by our language-mapping procedure leads to much larger standard errors, leading to different p-values ($p = 0.111$ for language; $p < 0.001$ for intervals; $p = 0.003$ for imprecise messages overall, pooling language and intervals).

Table 2: Relative slant by senders' preferred format

	(1)	(2)	(3)
	Language - Number	Interval - Number	Imprecise - Number
Use ImpreciseXHigh	4.530* (2.728)	2.935*** (1.053)	3.847*** (1.324)
Use ImpreciseXLow	-5.938** (2.833)	-0.784 (0.957)	-2.896** (1.330)
Observations	2008	1695	3398
Mean over Obs	-0.32	-0.13	-0.06
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Notes: This table reports the relative difference in slant between the chosen imprecise communication and numbers, by incentive condition and preferred message. In Column 1, the dependent variable is the language message selected for a communication, mapped to numbers using the benchmarking exercise, minus the numeric message. In Column 2, the dependent variable is the midpoint of the selected interval message minus the numeric message. Column 3 pools the differences across Columns 1 and 2. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and incentive type. *Use ImpreciseXHigh* interacts an indicator equal to one when the sender preferred the imprecise communication with an indicator for directional-high incentives. Similarly, *Use ImpreciseXLow* interacts an indicator equal to one when the sender preferred the imprecise communication with an indicator for directional-low incentives. Controls include the individual indicators as well as fixed effects for the respondent and true probability. Mean over obs reflects the average difference between the imprecise and numeric messages for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

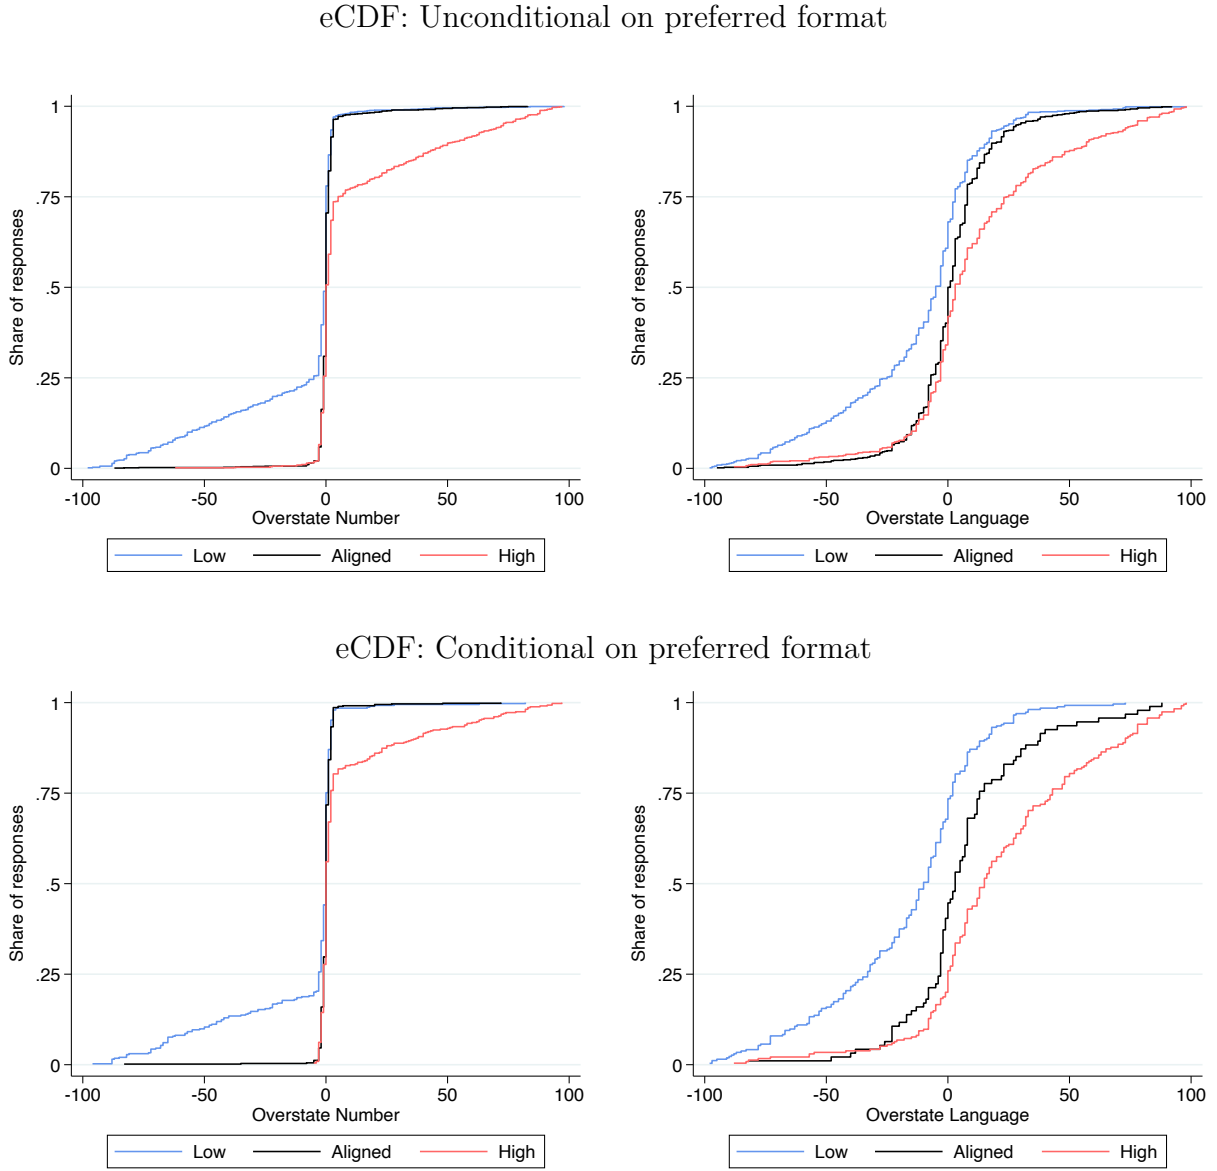
Distributions of messages

Next, we study how slant varies by message format at the individual level. In Figure 2, we visually compare slant by incentives and message format using empirical cumulative density function (eCDF) plots of the raw data.

The top-left panel shows the eCDF of senders' slant of numbers. We see that very few individuals with aligned incentives (shown with the black line) slant messages in either direction. In particular, Table 3 shows that the median slant is 0, the 30-70th percentile range is $[-1, 1]$, and the 10-90th percentile range is $[-2, 2]$. That is, senders in the aligned condition communicate the number corresponding to the true probability.

Next, the red line shows the distribution of slant for senders with directional-high incentives.

Figure 2: Cumulative Distribution Function of Slant by Incentives and Format



Notes: This figure plots the empirical cumulative distribution functions (eCDFs) for slant, by incentive conditions and message format. The x-axis reflects the degree to which senders overstate their message compared to the truth: “Overstate number” in the lefthand panels simply subtract the real probability from the numeric message. “Overstate language” in the righthand panels takes the language-to-number mapping from the benchmarking exercise and subtracts the true probability. The top panels show the distribution for all messages, while the bottom panels condition on the messages using the sender’s preferred format. The sample includes 2,008 choices across 251 senders. Similar eCDFs for interval messages are shown in Appendix Figure A1.

The median slant is again 0 and the 30-70th percentile range is $[-1, 3]$. In other words, most subjects do not slant numbers very much. However, there is a long tail: the 10-90th percentile

range is $[-2, 52]$, indicating that a positive average slant is primarily driven by the small fraction of senders who slant to extremes. We see similar patterns for the directional-low group (shown with the blue line). The median slant is 0, the 30-70th percentile range is $[-3, 0]$, and the 10-90th percentile range is $[-57, 2]$.

Table 3: Slant by Format and Preferred Format

Slanting Numbers					
Incentives	10th	30th	50th	70th	90th
Directional-low	-57	-3	0	0	2
Aligned	-2	-1	0	1	2
Directional-high	-2	-1	0	3	52

Slanting Language					
Incentives	10th	30th	50th	70th	90th
Directional-low	-57	-18	-5	2	17
Aligned	-15	-3	0	7	20
Directional-high	-15	-2	3	18	57

Slanting Language Conditional on Choosing Language					
Incentives	10th	30th	50th	70th	90th
Directional-low	-63	-28	-10	0	16
Aligned	-23	-3	3	12	38
Directional-high	-8	2	15	33	73

Notes: This table reports the raw data for the message reported by percentile and incentive condition. The top panel reports this in numbers, the middle panel reports this for language, and the bottom panel reports this for language conditional on preferring language. The sample includes 2,008 choices across 251 senders.

Senders have a different pattern of behavior when choosing among language messages, as shown in the top-right panel of Figure 2. Unsurprisingly, given the noise in the mapping of language to numbers, and the coarser message space, we see that senders with aligned incentives sometimes slant upwards or downwards by a modest amount; the 30-70th percentile range is $[-3, 7]$.

With directional-high incentives, we now see a shift in the *median* slant. Here, the median slant is 3, the 30-70th percentile range is $[-2, 18]$, and the 10-90th percentile range is $[-15, 57]$.

In other words, while the right tail of the language-slant distribution looks similar to that of the number-slant distribution, there are more participants who slant a little bit. This suggests that there is a new cohort of senders who don't feel comfortable sending a numeric message other than the truth but do feel comfortable sending a language message that might overlap with the truth but errs in their preferred direction. We see similar patterns for participants with directional-low incentives. The median slant is -5, the 30-70th percentile range is [-18, 2], and the 10-90th percentile range is [-57, 17].

These differences are even starker when we condition on senders' preferred message format. The bottom-left panel of Figure 2 is similar to the top-left panel, showing that senders who prefer to send numeric messages slant numbers in a similar way to senders overall. However, the bottom-right panel shows a more extreme slant among the senders who prefer to send language messages. For senders who prefer language with directional-high incentives, the median language slant is 15, the 30-70th percentile range is [2, 33], and the 10-90th percentile range is [-8, 73]. And for senders who prefer language with directional-low incentives, the median language slant is -10, the 30-70th percentile range is [-28, 0], and the 10-90th percentile range is [-63, 16].


In other words, the majority of senders who choose to send language messages are sending slanted language messages, while the majority of senders who choose to send numeric messages are not. Slant for numbers is driven more by a small group of participants who slant to extremes, while slant for language occurs more uniformly across the distribution.

We show the equivalent plots for the distribution of slant for intervals in Appendix Figure A1. It visually looks in between the figures for numbers and language, suggesting that the imprecision of intervals leads to a related, but more muted, impact.

2.5.2 Precise good news and imprecise bad news

Next, we turn to heterogeneous effects by the true probability of drawing a red ball. In particular, in the directional conditions we classify the state as "good" when the true probability is higher (lower) for senders with high (low) incentives and "bad" when the true probability is lower (higher) for senders with high (low) incentives. Consistent with Hypothesis 3, Figure 3 shows that on average senders slant their messages more in the direction of their incentives when the state is bad: For states that are aligned with senders' incentives ("good states"), senders send numeric and language messages that are similar to the aligned condition. However, when the states are misaligned with senders' incentives ("bad states"), they substantially distort both numeric and language messages. In accordance with the second part of Hypothesis 3, Appendix Table A4 shows that senders send more imprecise messages when the state is bad.

Figure 3: Sender’s Message by True Probability



Notes: This figure plots the sender’s message on the y axis and the true probability on the x axis, by incentive condition. All message formats are pooled in this figure; language messages are included using the language-to-number mapping from the benchmarking exercise, and interval messages are included based on the midpoint of the interval. Bars reflect 95% confidence intervals. The sample includes 4,000 choices across 500 senders.

2.5.3 Predictors of the strategic use of language

We can glean further insights into the underlying mechanisms by exploring self-reports that we collected at the end of the experiment, along with heterogeneous treatment effects.

First, Appendix Figure A2 shows that senders self-report preferring to use language in the real world “to make it easier to withhold information” (87%), “to make it easier to lie to someone” (84%), and “to make it easier to persuade someone that the likelihood is higher than it actually is” (58%). Perhaps surprisingly, only 31% say they would use language to make their message easier to understand. We create an index that takes the average of the Likert responses for the former three reasons to use language, all of which relate to strategic persuasion. Interestingly, Column 1 of Table 4 shows that a self-reported preference to strategically use language almost entirely explains the response to incentives that we observe in our experiment ($p = 0.001$). This suggests at least some self-awareness. That is, when looking just at the message choices in our experiment we cannot determine whether senders

motivatedly self-deceive and believe their language and number communications are equally close to the truth. This additional data suggests, however, that self-deception is unlikely.

Second, at the end of the experiment we asked senders to work through a three-question module to assess their numeracy.¹³ We observe that the number of questions answered correctly on our three-question assessment is predictive of the strategic use of language: Column 2 shows that each additional correct answer increases the likelihood of using language when facing directional incentives by 8.5 percentage points ($p = 0.011$). In other words, more numerate respondents appear to be more likely to strategically leverage language to persuade.

Table 4: Heterogeneous treatment effects: Message format choice

	(1)	(2)	(3)	(4)
	X = Persuasion	X = Numeracy	X = Female	X = Resp Time
X*Directional	0.336*** (0.102)	0.085** (0.033)	0.049 (0.050)	0.001 (0.001)
Directional Incentives	0.019 (0.071)	0.057 (0.079)	0.216*** (0.038)	0.182*** (0.050)
Observations	2008	2008	1960	2008
Aligned Mean	.3	.3	.29	.3
Respondent FE	Yes	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes	Yes

Notes: This table reports heterogeneous treatment effects for message format choice. In all columns, the dependent variable is an indicator equal to one when the sender chose to communicate using language rather than numbers. The coefficient of interest is an indicator equal to one when the sender faces directional incentives, interacted with one of four predictor variables: the average response to three 1-5 scale Likert questions on the likelihood of using language to persuade in the real world (Column 1); the number of correct answers on a three-question numeracy module (Column 2); an indicator for whether the respondent is female (Column 3); and the percentile breakdown of the sum of response times in each main decision (Column 4). Controls include an indicator equal to one when the sender faces directional incentives as well as fixed effects for the respondent and true probability. Aligned mean reflects the likelihood of using language under aligned incentives for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Finally, we look at heterogeneous treatment effects by gender (Column 3) and response time (Column 4) and see that neither is predictive of strategic use of language.¹⁴ If we think of response time as a proxy for attention, this suggests that the strategic use of language is

¹³These questions were adapted from a longer numeracy assessment included in Kahan et al. (2012).


¹⁴Our measure of response time is percentile breakdown of the sum of response times in each main decision.

not merely an artifact of more or less attentive participants in the experiment—more subtly, an *understanding* of how to leverage language appears to be key.

Receivers’ Behavior

Next, we turn to receivers. We first show, consistent with the first part of Hypothesis 4, that receivers are persuaded overall. Figure 4 compares average receiver guesses when senders face aligned, directional-high, and directional-low incentives (pooling across message formats). Receivers’ guesses look remarkably similar to the messages senders sent (Figure 3). Appendix A5 shows that receivers’ guesses are 15 percentage points higher when their paired sender faced directional-high compared to directional-low incentives ($p < 0.001$).¹⁵

Figure 4: Receiver’s Guess by True Probability



Notes: This figure plots the receiver’s guess on the y axis and the true probability on the x axis, by incentive condition that the sender faced. Guesses based on all sender message formats (number, language, interval) are pooled in this figure. Bars reflect 95% confidence intervals. The sample includes 4,000 choices across 500 receivers.

Appendix Table A6 compares average receiver guesses by the format they receive. Consistent with the second part of Hypothesis 4, we see that receivers’ guesses are significantly more

¹⁵This reports the gap for receivers with paired senders deciding between numbers and language. Receivers’ guesses are 20 percentage points higher when their sender faced directional-high incentives and was deciding between numbers and *intervals*.

likely to be in the direction of senders’ incentives when they receive a language or interval message as compared to a numeric message.

Receivers make significantly larger mistakes when they receive language versus numeric messages. One simple measure of this is average error: $|guess - probability|$.¹⁶ When receivers see numbers in the aligned condition, their average error is 2.5 pp (s.e. 0.2 pp). When receivers see language in the aligned condition, their average error is 8.7 pp (s.e. 0.5 pp). That is, numbers are interpreted more accurately than language.

In the conditions where senders face directional incentives, numbers lead receivers to an average error of 12.9 pp (0.8 pp), while language leads receivers to an average error of 26.1 pp (1.1 pp). That is, when senders are incentivized to persuade receivers, language leads to receiver beliefs that are even more inaccurate (all differences, including the interaction, have $p < 0.001$).¹⁷

Recall that half of receivers learn the incentives of their paired sender, while the other half does not. Thus far, we have combined analyses for receivers who do or do not know what senders’ incentives are. Next, we show that receivers do not fully adjust to the knowledge of senders’ incentives. In Appendix Table A7, we see that when receivers learn the incentives of senders, they are persuaded statistically-significantly less. However, 79% of the main persuasion effect remains, indicating that knowledge of senders’ incentives is not sufficient to avoid being persuaded.

Finally, we look at receivers’ ratings of how informative the message is. When receivers see a numeric message, Appendix Figure A3 shows that 72% rate the message as “very informative” and only 3% rate it as either “very” or “somewhat” uninformative. However, when receivers see a language message, 17% rate the message as “very informative” and 21% rate it as uninformative. This suggests that receivers might put less weight on language messages when updating their beliefs. In Study 1, where receivers’ priors are uninformed, it’s difficult to unpack the implications of this result, but Study 2 sheds further light on this.

3 Study 2

3.1 Design Overview

The benefit of the abstract “balls-and-urns” design of Study 1 is that the language-number mapping and the directional versus aligned incentives are clean and stripped of context.

¹⁶Specifically, we take the average after winsorizing at the 5- and 95-percent levels, with standard errors clustered at the receiver level.

¹⁷In both cases, receivers’ errors from interval messages lie in between their errors from numeric and language messages.

Study 2 considers communication behavior in a more real-world environment, in a setting where we expect the strategic use of language to have important implications. Specifically, in July 2024 we recruited a sample of academic social science researchers, most of whom are tenure-track or tenured professors, and asked them to communicate treatment effects identified in published social science papers to a sample of government policymakers we recruited from July to September 2024.

The design is similar in spirit to Study 1, with some key differences in addition to the population and type of information being communicated:

1. Researchers face either aligned or directional-high incentives. The incentive scheme is structured the same way as in Study 1. In addition, researchers see vignettes tailored to their incentive condition: researchers in the aligned condition are encouraged to “give the policymaker the best understanding of the data” while researchers in the directional-high condition are told to persuade their policymaker that the treatment effect is large for the sake of a government grant or policymaker attention.
2. Because we are less interested in identifying the underlying mechanism of imprecision and more interested in testing whether there is strategic use of language in this more real-world setting, we only compare the choice of language versus numbers and omit the comparison of intervals versus numbers.
3. Rather than asking people to communicate a probability, as in Study 1, we ask researchers to communicate the treatment effect of a policy intervention tested in a real research study, always reported in percentage points.

It is important to note that the aim of Study 2 is to identify whether researchers are more likely to opt for language when they face strategic incentives to persuade. Our data do not shed light on *how often* researchers face strategic incentives (strong enough to at least partially offset incentives to communicate accurately) in practice.

Our pre-registered hypotheses for Study 2 are broadly similar to those in Study 1, accounting for the differences described above.

3.2 Sample

3.2.1 Academic Researchers

We personally invited 241 social science researchers to participate in our study. 145 social science completed our survey, for a response rate of 60%.¹⁸¹⁹ 75% of participating researchers

¹⁸Appendix section B.3.1 shows an example recruitment email. Thank you to all who participated!

¹⁹This excludes eight individuals who attrited out of the sample and two individuals who were dropped because they indicated that they knew the overall hypothesis of our study.

are in tenure-track positions (Assistant Professor, Associate Professor, Professor, or equivalent) and 23% are post-docs or PhD students. 61% self-identify as men. 81% work with empirical or experimental data for most of their research projects, and 61% report working in a policy domain (most commonly economic policy, health, or education).

3.2.2 Policymakers

We recruited a sample of 66 policymakers through the University of Warwick’s *Policymakers Lab*. Policymakers in the sample work in 5 countries (United States, United Kingdom, Australia, India, and Belgium) and are primarily civil servants working in central government (for instance, U.S. policymakers in our sample work in the Department of State; Department of Health and Human Services; U.S. Agency for International Development; Department of Homeland Security; Department of Justice; and General Services Administration). 76% of policymakers in the sample report having been involved in policy adoption decisions. 52% have a Master’s or Professional degree, 26% have a PhD, and 21% have a Bachelor’s degree or equivalent. 54% of the sample identify as women. Policymakers are recruited with the motivation of “lending their expertise to contribute to and inform academic research”. Given that this sample is likely to be more experienced with and favorable towards research evidence than the typical policymaker, we expect that any observed naivete with respect to treatment effects is likely to represent a lower bound.

3.3 Design for the Researchers Experiment

In Study 2, 145 researchers complete just one main part, the Communications part.²⁰ A separate group of 22 researchers complete a Benchmarking part, described below. In addition to a \$10 payment for completion, researchers could earn \$10 as a bonus payment based on their message communication decisions.²¹ Appendix Section B.3 shows screenshots from the experiment.

3.3.1 Communications

In the Communications part, researchers first learn about two different, real research studies. The studies are randomly assigned from a broader set of six studies, shown in Table 5. All six studies identify the impact of a policy-relevant intervention, although the particular policy

²⁰Unlike in Study 1, to keep survey lengths shorter in Study 2 researchers and policymakers only read the full instructions for their own study versions.

²¹Payments were distributed in the form of Amazon gift cards, and we matched the currency based on the respondents’ preferences. Participants had the option to opt out of receiving a bonus payment; 23% took this option. Results are qualitatively similar if we exclude participants who opt out.

Table 5: Research studies included in Study 2

Paper	Policy Domain	Outcome	Effect
Domurat et al (AER, 2021)	Health insurance	Letter → takeup healthcare coverage	1.2pp
Adams et al (JFE, 2021)	Consumer finance, savings	Alt financial product info → switch accounts	1.9pp
Liebman & Luttmer (AEJ EP, 2015)	Labor market and aging	Social Security info → labor force participation	4.2pp
Fishbane et al (Science, 2020)	Criminal justice	New court summons → defendants appear	6.2pp
Burland et al (AER: I, 2023)	Education policy	Info about free tuition → applying to college	9.0pp
Bergman et al (AER, 2024)	Economic mobility, housing	Info, some support → moving to opportunity	13.8pp

domain varies across the studies. Importantly, each study reports the outcome of interest, the intervention’s effect, in percentage points. We present the effect in a bar plot which displays the control group (or baseline) mean alongside the treatment effect, as shown in Appendix B.3.1. The researcher’s task is to decide how to communicate the study results to their paired policymaker.

Before making their communication decisions, each researcher is assigned to one of two conditions: *directional-high* or *aligned*.²²

In the *directional-high* condition, researchers are told to “imagine you are the author of the study testing the policy intervention and you are trying to persuade a policymaker that the research is promising to increase your chance of getting a government grant or policymaker attention”. To ensure this framing is incentive-compatible, they are additionally told that they are more likely to earn the bonus payment if their paired policymaker predicts the intervention’s effect size is larger.²³

In the *aligned* condition, researchers are instead told to “imagine you are the author of the study testing the policy intervention and you are trying to give the policymaker the best understanding of the data to inform their policy decisions”. We also include incentives for this condition: researchers are more likely to earn the bonus payment if their paired policymaker correctly predicts the intervention’s effect size.²⁴

After learning about the results of the research studies and their incentive assignment, researchers make their message content choices for numbers and language. For numbers, they complete the sentence, using the example of the first paper on healthcare insurance takeup: “In the control group, 8.1% of individuals take up healthcare coverage. The letter led to an increase in healthcare coverage of [X] percentage points. That is, 8.1 + [Your

²²Unlike in Study 1, in Study 2 assignments are across-subjects, both to keep the survey as short as possible for bandwidth-constrained researchers and also to limit the possibility of guessing the research hypothesis in a sample that is likely to be more attuned to what we might be testing.

²³Specifically, if the paired policymaker predicts the effect size is X percentage points, the probability the researcher wins the bonus payment is equal to X%, with a minimum probability of 0%.

²⁴Specifically, if the true effect size is X percentage points, the researcher will win the bonus payment if their paired policymaker predicts the effect size is between X-1 and X+1 percentage points.

Response] % of those receiving the letter take up healthcare coverage.”²⁵ We were careful to design the communications such that they are clear even to those who might be less familiar with working in percentage points. For instance, $8.1 + [Your\ Response]$ dynamically updates in response to what the researcher inputs for X.

For language, researchers complete the sentence, again using the example of the paper on healthcare coverage takeup: “In the control group, 8.1% of individuals take up healthcare coverage. The letter led to a(n) [WORD] increase in healthcare coverage.” As in Study 1, researchers use a dropdown list to select a word or phrase—in this case including, for example, “tiny”, “intermediate”, and “very large”—to complete the sentence.

As in Study 1, after making their message content choices, researchers select their preferred message format, numeric or language. Again, researchers are told their preferred choice is more likely but not guaranteed to be the one communicated. The experiment ends with a short demographics questionnaire.

3.3.2 Benchmarking

To keep the Researchers survey as short as possible, in Study 2 we run the Benchmarking exercise among a separate group of 22 researchers. As in our main Researchers Experiment, we personally invited mostly tenure-track academic researchers to take part. In this separate study, for each word in the dropdown list in the Researchers Experiment, researchers make incentivized predictions about the “numeric effect size (as a percentage point increase) [they] think others would expect each of the following words or phrases would correspond to.” Researchers make separate predictions for each policy paper, to account for the possibility that language is perceived differently according to the policy context. As in Study 1, this benchmarking exercise allows us to measure the slant of language messages.

3.4 Design for the Policymakers Experiment

The policymakers experiment is focused entirely on policymakers’ predictions of the messages sent by their paired researchers. As in the researchers experiment, in addition to a \$10 payment for completion, policymakers could earn \$10 as a bonus payment based on their predictions. Appendix Section B.4 shows screenshots from the experiment.

In the experiment, policymakers first select the two policy domains that best reflect their area of expertise, from the set of six indicated in Table 5. They learn about the relevant research study in each of the domains (but not the results). Then, they receive a message from a randomly paired researcher communicating the treatment effect. After receiving the

²⁵Responses are constrained to be between 0 and 100 minus the baseline mean.

message, policymakers predict the actual treatment effect reported in the study. Policymakers are informed of the two incentive conditions faced by researchers but were not told which incentive scheme their paired researcher faced, akin to the *incentives-unknown* condition in Study 1.

After policymakers submit their predictions they report how likely they are to share information about the study with their colleagues as well as whether they would like to receive an infographic summarizing the study results. These outcomes point to the policy-relevant implications of their predictions—are policymakers likely to engage with the evidence in a meaningful way outside of the experiment? At the end of the experiment we include a short demographic survey and debrief policymakers on the real results of the research studies.

3.5 Results

Researchers' Behavior

Table 6: Format choice, overall and by good state and bad state

	(1)	(2)	(3)
	Use Language	Use Language	Use Language
Directional Incentives	0.291*** (0.054)		
Effect		-0.022** (0.010)	0.006 (0.006)
Observations	290	138	152
Mean over obs	0.26	0.41	0.13
Subject controls	Yes	Yes	Yes
Study FE	Yes	No	No
Sample	All	Directional-High	Aligned

Notes: This table reports the likelihood of researchers using language as the dependent variable. Column 1 shows the impact of directional incentives via *Directional Incentives*, an indicator equal to one when the researcher faces directional-high incentives. Columns 2 and 3 condition on responses when the researcher faces directional-high and aligned incentives, respectively. The key independent variable is *Effect*, the treatment effect from the study in question in percentage points. Subject controls include indicators for whether the researcher is tenure-track, male, and does empirical research. Study fixed effects are only included in Column 1. Mean over obs reflects the likelihood of using language for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

In the raw data, we see clear evidence that incentives affect senders’ choice of messages, in support of Hypothesis 1. Senders in the aligned condition choose language messages over numeric messages 13 percent of the time (s.e. 3 pp). This share increases substantially when they have directional incentives (to 41 percent, s.e. 4 pp). Column 1 of Table 6 reports the main effect: researchers are 29 percentage points more likely to use language when they face directional incentives ($p < 0.001$). Somewhat surprisingly given the differences across studies (e.g., one might imagine there are more reasons to use language to communicate context in Study 2), these levels are both qualitatively and quantitatively similar to those in Study 1.

Next, we exploit variation in the actual effect sizes of the research studies and find evidence that is consistent with Hypothesis 3 adapted to this study (here, small effect sizes are “bad states” for senders with directional incentives). First, incentives induce more language use for studies that have smaller effect sizes. Column 2 of Table 6 shows that for every percentage point increase in study effect size, researchers are 2.2 percentage points less likely to use language when in the directional-high condition ($p = 0.028$).²⁶ Column 3 shows a precisely-estimated null effect when incentives are aligned. That is, researchers’ decisions to use language are unaffected by the effect size when they are incentivized to accurately communicate the information to their paired policymaker.

As in Hypothesis 2 in Study 1, Table 7 demonstrates that researchers also behave differently in how they slant numbers and language. There is a modest difference in how researchers slant numbers across incentive groups. On average, researchers with aligned incentives send numeric messages that are slightly smaller than the true effect size. Column 1 of Table 7 shows that, compared to researchers with aligned incentives, researchers with directional incentives communicate using numbers that are 13% larger ($p = 0.004$).²⁷ This effect is statistically significant, but modest in size, suggesting that senders face a significant cost of misreporting the numerical estimates they observe.²⁸

Senders slant language significantly more than they do numbers. Column 2 of Table 7 shows that, compared to researchers with aligned incentives, researchers with directional incentives slant language that maps to numbers that are 49% larger. The frequency with which researchers use each word across the two conditions reflects this overall pattern: in the aligned condition, only 48% of researchers use a word to indicate a large effect size (“fairly large” or larger) while in the directional-high condition 74% of researchers use such language. Column 3 of Table 7 shows that the effect of incentives on slant is significantly larger for

²⁶This coefficient represents the interaction between sender incentives and study effect size in a regression that controls for individual-level characteristics and study-level fixed effects.

²⁷To get this percentage, we exponentiate the estimated coefficient in the table.

²⁸Much of the effect is driven by a negative overall level of slant in the aligned group; the directionally-incentivized senders slant upwards on average, but the effect is statistically-indistinguishable from 0.

language than for numbers ($p = 0.001$). Related to the last part of Hypothesis 2, Appendix Table A8 shows that the gap between language and numeric slant increases by 64% when looking at the interaction of incentives with an indicator for the choice to communicate using language ($p = 0.044$).²⁹ In other words, essentially the entire gap in slant is driven by senders who prefer to send language messages.

Table 7: Message slant by format

	(1)	(2)	(3)
	log(Number)	log(Language)	log(Lang)-log(Num)
Directional Incentives	0.122*** (0.044)	0.397*** (0.077)	0.275*** (0.085)
Observations	289	289	289
Mean over obs	1.49	1.28	-0.22
Subject controls	Yes	Yes	Yes
Study FE	Yes	Yes	Yes

Notes: This table reports researchers’ slant by message format and incentives. Column 1 regresses $\ln(\text{Number})$ on an indicator for *Directional Incentives*, which equals one when the researcher faces directional-high (versus aligned) incentives. Column 2 uses the same specification, but uses $\ln(\text{Language})$ as the dependent variable, where *Language* represents the number that the word used was benchmarked to. Column 3 regresses the difference in (1) and (2) on incentives. Subject controls include indicators for whether the researcher is tenure-track, male, and does empirical research, and study fixed effects are included. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Policymakers’ Behavior

We see evidence that researchers’ incentives have an effect on policymakers’ predictions. As shown in Columns 1 and 2 of Appendix Table A10, policymakers predict larger effects (2.2 percentage points [$p = 0.030$], or 22% [$p = 0.042$], larger) when researchers are incentivized to directionally persuade policymakers. In other words, policymakers are at least partially naive to researchers’ slant.³⁰

Columns 3 and 4 of Appendix Table A10 restrict to cases where researchers are incentivized to persuade policymakers that the study effects were large. Here we see suggestive evidence

²⁹To get these percentages, we exponentiate the estimated coefficients in the tables.

³⁰Recall that in Study 2 policymakers do not know the incentives researchers face, so sophistication would involve “correcting for” unrealistic messages in general. Recall also that in Study 1 we observed that naivete persisted even in the case where incentives were communicated.

that policymakers were particularly persuaded by language compared to numeric messages (though the statistical power is lower): their predictions are 3.3 percentage points ($p = 0.095$), or 23% ($p = 0.307$), larger when receiving a language compared with a numeric message.

After each prediction, policymakers indicate whether they would be willing to extend the survey (i.e., a costly signal of effort) in order to receive an infographic about the study in question and also whether they plan to share information about the study with their colleagues. The majority (67% and 57%, respectively) indicate a willingness to meaningfully engage with the information outside of the experiment, pointing to its relevance for the policymakers involved.

4 Discussion and Conclusion

We provide evidence that, in many settings, people use language more when they want to persuade others. We first show this effect in a tightly-controlled abstract setting, and then show qualitatively- and quantitatively-similar effects in a setting where academic researchers communicate the effect size of a policy intervention reported in a research paper to policymakers. In addition, both experiments find that senders who prefer language choose to slant language more, a pattern that is not observed for numbers. These results indicate that imprecise language provides a way for some senders who wish to distort messages in the direction of their incentives to avoid overt lies.

Psychologically, this behavior has commonalities with moral wiggle room, in which people intentionally avoid information in order to excuse selfish actions (e.g. Dana, Weber, and Kuang 2007). Here, senders are avoiding sending precise messages in order to excuse the selfish action of slanting messages towards their incentives. Interestingly, the heterogeneities we observe in Study 1 suggest that many senders appear to be aware of their strategy, suggesting that it may not be as much about self-deception as about signaling to receivers that they are not behaving deceitfully.³¹

These results are consistent with senders having a greater psychological disutility of misreporting precise messages than imprecise messages (as in Serra-Garcia, Van Damme, and Potters 2011), and primarily affect people who have “intermediate” preferences. Senders with a high cost of misreporting will not slant either numeric or language messages, and senders who have no cost of misreporting will slant both messages to maximize expected payoffs; both of these types would not have a clear reason to prefer language over numbers. However, an intermediate type who finds it psychologically costly to slant numbers, but less

³¹This observation relates to the findings of Serra-Garcia and Saccardo (2023) that many people are sophisticated about how they distort beliefs.

costly to slant language, would end up both preferring language and slanting language more, exhibiting the patterns we observe.

It is worth noting that all of these results presuppose that messages have inherent meaning to senders. For instance, the message “It is probable that you will draw a red ball.” is meaningful because the word “probable” has meaning outside of this context, but this meaning can plausibly reflect a probability of 51% or 81%. With the message “There is a 66% chance that you will draw a red ball.”, the “66%” has a very precise meaning outside of the experiment. This contrasts with some theories of cheap talk in which messages are only meaningful through their impact in equilibrium.³²

Study 2 poses another natural question: How much do readers consider researchers’ incentives when interpreting research findings? Here, we consider the case where it is difficult to know whether a researcher has aligned incentives, or is just trying to promote the findings of a paper. We see that policymakers are indeed persuaded by researchers in this case. Future research can explore whether this is true for other consumers of research evidence, including groups such as journal editors who might plausibly be more sophisticated about the impact of incentives on researchers.

We also caution against interpreting our results as indicating that researchers usually misrepresent their results. For one, Study 2 does not shed light on how often in practice researchers face strategic incentives to persuade that are strong enough to affect their communication strategies. It is also worth restating that in Study 2, even in the case where incentives to persuade are explicitly provided to academic researchers and they have an anonymous one-shot interaction, most researchers do not slant numbers upwards. That is, while we do see a small subset of researchers slant numbers upwards, most researchers seem to have an aversion to *numeric* misreporting. Instead, our results suggest that is important to be more cautious when interpreting imprecise descriptions of results, which may be distorted, relative to more precise statements of effect sizes.

We see several additional directions for future work. First, in many settings there are repeated interactions between senders and receivers; in such settings, it may be important for senders to build a reputation for not misreporting messages. This could scale up or down our effect; on the one hand, receivers may learn to infer senders’ strategies, making all distortions less beneficial. But it could also amplify senders’ tendency to use language, since distorted numeric predictions may be more transparently incorrect if the state is revealed. Relatedly, there are many other types of incentives that affect communication. When building

³²As an example, in our Study 1, senders could send messages in the range $\{0\%,1\%,2\%,\dots,100\%\}$. We think that the interpretation of a sender choosing “66%” is not wholly dependent on the menu of options. If they faced the set $\{0\%,2\%,4\% \dots,200\%\}$, we would not expect them to switch to sending “132%”. However, in other cases, we may expect the context to be of first-order importance, especially for language.

a reputation, researchers do not just wish to persuade audiences that their effects are large in one paper, but that they should be trusted to analyze additional effects credibly. If using numbers signals precision, a researcher may want to use numbers for precise effects and language for imprecise effects. As another example, in cases where receivers are biased, reputation incentives may push senders towards imprecisely confirming receivers' prior beliefs using language rather than giving unbiased estimates of the truth.³³

Second, it is important to document the degree to which persuasion impacts real-world communication. For instance, when incentives for journalists, companies, or researchers vary, do they shift to describing numeric effects more with language, and is there a way to quantify the slant that they use? For instance, Gentzkow and Shapiro (2010) uses the similarity of text data with congressional speeches to quantify political slant of newspapers, and Raymond and Taylor (2021) uses variation in incentives induced by the timing of baseball games to study how local newspapers numerically report weather forecasts. Testing incentive variation on the extensive margin of reporting in numbers versus language, and on the intensive margin of the particular numbers and words used, could justify the external validity of the mechanisms we discuss.

Finally, our results indicate a substantial level of naivete among receivers. From a policy perspective, it is important to understand how much of a role receivers' sophistication can play in affecting senders' strategies. If receivers became more aware that observing language messages was a signal that senders were distorting more, this may make them form more accurate beliefs, and be persuaded less, by language. In turn, this could lead senders to send more precise messages in the first place, as using imprecise messages to persuade would be less beneficial. Policies that increase awareness of the strategic use of language can potentially improve the efficiency of communication and the accuracy of people's beliefs. Future work can also shed light on the efficacy of language use guidelines or even requirements for reporting numbers (such as the CONSORT guidelines for reporting on medical RCTs) as alternative policy tools.

³³In a political context, Thaler (2023) shows that incentives to be perceived as truthful can lead senders to send more directly-false messages when receivers engage in motivated reasoning. This argument may extend to message format choices as well.

References

- Abeler, Johannes, Daniele Nosenzo, and Collin Raymond (2019). “Preferences for Truth-Telling.” In: *Econometrica* 87(4), pp. 1115–1153.
- Alempaki, Despoina, Valeria Burdea, and Daniel Read (2023). *Deceptive Communication: Direct Lies vs. Ignorance, Partial-Truth and Silence*. Discussion Paper 444. University of Munich (LMU), Humboldt University Berlin, Collaborative Research Center Transregio 190: Rationality, and Competition. URL: <https://www.econstor.eu/bitstream/10419/282135/1/444.pdf>.
- Blume, Andreas, Ernest K Lai, and Wooyoung Lim (2020). “Strategic information transmission: A survey of experiments and theoretical foundations.” In: *Handbook of experimental game theory*. Ed. by C. M. Capra, Rachel T.A. Croson, Mary L. Rigdon, and Tanya S. Rosenblat. Edward Elgar Publishing. Chap. 13, pp. 311–347.
- Cai, Hongbin and Joseph Tao-Yi Wang (2006). “Overcommunication in strategic information transmission games.” In: *Games and Economic Behavior* 56(1), pp. 7–36.
- Charness, Gary and Martin Dufwenberg (2006). “Promises and Partnership.” In: *Econometrica* 74(6), pp. 1579–1601.
- Choe, Kyoung Whan, Jalisha Jenifer, Christopher Rozek, Marc Berman, and Sian Beilock (2019). “Calculated avoidance: Math anxiety predicts math avoidance in effort-based decision-making.” In: *Science Advances* 5(11), eaay1062.
- Crawford, Vincent (1998). “A Survey of Experiments on Communication via Cheap Talk.” In: *Journal of Economic Theory* 78(2), pp. 286–298.
- Crawford, Vincent and Joel Sobel (1982). “Strategic Information Transmission.” In: *Econometrica* 50(6), pp. 1431–1451.
- Dana, Jason, Roberto Weber, and Jason Xi Kuang (2007). “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness.” In: *Economic Theory* 33.
- DellaVigna, Stefano, Woojin Kim, and Elizabeth Linos (2024). “Bottlenecks for Evidence Adoption.” In: *Journal of Political Economy* 132(8).
- Deversi, Marvin, Alessandro Ispano, and Peter Schwardmann (2021). “Spin doctors: An experiment on vague disclosure.” In: *European Economic Review* 139.
- Dhami, Mandeep K. and David R. Mandel (2021). “Words or numbers? Communicating probability in intelligence analysis.” In: *American Psychologist* 76(3), pp. 549–560.
- Dhami, Mandeep K. and David R. Mandel (2022). “Communicating uncertainty using words and numbers.” In: *Trends in Cognitive Sciences* 26(6), pp. 514–526.
- Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji (1995). “An experimental study of strategic information transmission.” In: *Economic Theory* 6(3), pp. 389–403.
- Edlin, Aaron S. and Michael Love (2022). *Needed in Empirical Social Science: Numbers*. Working Paper. SSRN. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4195779.
- Fischbacher, Urs and Franziska Föllmi-Heusi (2013). “Lies in Disguise - An Experimental Study on Cheating.” In: *Journal of the European Economic Association* 11(3), pp. 525–547.
- Gentzkow, Matthew and Jesse Shapiro (2010). “What Drives Media Slant? Evidence From U.S. Daily Newspapers.” In: *Econometrica* 78 (1).
- Gneezy, Uri (2005). “Deception: The Role of Consequences.” In: *American Economic Review* 95(1), pp. 384–394.
- Graeber, Thomas, Shakked Noy, and Christopher Roth (2024a). *Explanations*. CESifo Working Paper 2024. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4856950.
- Graeber, Thomas, Shakked Noy, and Christopher Roth (2024b). *Lost in Transmission*. CESifo Working Paper 2024. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4711260.

- Graeber, Thomas, Christopher Roth, and Florian Zimmerman (2024). “Stories, Statistics and Memory.” In: *Quarterly Journal of Economics*, pp. 1–45. DOI: <https://doi.org/10.1093/qje/qjae020>.
- Hjort, J., D. Moreira, G. Rao, and J.F. Santini (2021). “How research affects policy: Experimental evidence from 2,150 Brazilian municipalities.” In: *American Economic Review* 111 (5), pp. 1442–1480.
- Ho, Emily H., David V. Budescu, Mandeep K. Dhimi, and David R. Mandel (2015). “Improving the communication of uncertainty in climate science and intelligence analysis.” In: *Behavioral Science and Policy* 1(2), pp. 43–55.
- Holtgraves, Thomas and Audrey Perdeu (2016). “Politeness and the communication of uncertainty.” In: *Cognition* 154, pp. 1–10.
- Hopster, Jeroen (2023). “Climate Change, Uncertainty, and Policy.” In: *Handbook of the Philosophy of Climate Change*. Ed. by Gianfranco Pellegrino and Marcello Di Paola. Springer, pp. 977–1000.
- Hossain, Tanjim and Ryo Okui (2013). “The binarized scoring rule.” In: *The Review of Economic Studies* 80(3), pp. 984–1001.
- Jin, Ginger Zhe, Michael Luca, and Daniel Martin (2022). “Complex Disclosure.” In: *Management Science* 68(5), pp. 3236–3261.
- Juanchich, Marie, Miroslav Sirota, and Christina Lea Butler (2012). “Politeness and the communication of uncertainty.” In: *Organizational Behavior and Human Decision Processes* 118(1), pp. 72–81.
- Kahan, D., E. Peters, M. Wittlin, P. Slovic, L.L. Ouellette, D. Braman, and G. Mandel (2012). “The polarizing impact of science literacy and numeracy on perceived climate change risks.” In: *Nature Climate Change* 2, pp. 732–735.
- Kartik, Navin (2009). “Strategic Communication with Lying Costs.” In: *The Review of Economic Studies* 76(4), pp. 1359–1395.
- Krupka, E. and R. Weber (2013). “Identifying social norms using coordination games: Why does dictator game sharing vary?” In: *Journal of the European Economic Association* 11 (3), pp. 495–524.
- Lipman, Barton L. (2009). *Why is Language Vague?* Working Paper. URL: <https://open.bu.edu/handle/2144/40222>.
- Mastrandrea, Michael D., Katharine J. Mach, Gian-Kasper Plattner, Ottmar Edenhofer, Thomas F. Stocker, Christopher B. Field, Kristie L. Ebi, and Patrick R. Matschoss (2011). “The IPCC AR5 guidance note on consistent treatment of uncertainties: a common approach across the working groups.” In: *Climatic Change* 108(675).
- Mehmood, S., S. Naseer, and D. Chen (2023). “Training policymakers in econometrics.” In: *Working Paper*.
- Mosteller, Frederick and Cleo Youtz (1990). “Quantifying Probabilistic Expressions.” In: *Statistical Science* 5(1), pp. 2–12.
- Ott, Douglas E. (2021). “Words Representing Numeric Probabilities in Medical Writing Are Ambiguous and Misinterpreted.” In: *Journal of The Society of Laparoscopic and Robotic Surgeons* 25(3).
- Prelec, Drazen (1998). “The probability weighting function.” In: *Econometrica* 66(3), pp. 497–527.
- Raymond, Collin and Sarah Taylor (2021). ““Tell all the truth, but tell it slant”: Documenting media bias.” In: *Journal of Economic Behavior & Organizations* 184.
- Schulz, K F, D G Altman, and D Moher (2010). “CONSORT 2010 Statement: updated guidelines for reporting parallel group randomised trials.” In: *BMJ* 340.

- Serra-Garcia, Marta and Silvia Saccardo (2023). “Enabling or Limiting Cognitive Flexibility? Evidence of Demand for Moral Commitment.” In: *American Economic Review* 113(2).
- Serra-Garcia, Marta, Eric Van Damme, and Jan Potters (2011). “Hiding an inconvenient truth: Lies and vagueness.” In: *Games and Economic Behavior* 73(1), pp. 244–261.
- Thaler, Michael (2023). “The Supply of Motivated Beliefs.” In: *Working Paper*. URL: <https://hq.ssrn.com/submissions/AbstractComments.cfm?AuthorID=2725209&AbstractID=4494688>.
- Toma, Mattie (2024). *Policymakers Lab*. <https://warwick.ac.uk/fac/soc/wbs/subjects/bsci/research/policymakers-lab>.
- Toma, Mattie and Elizabeth Bell (2024). “Understanding and increasing policymakers’ sensitivity to program impact.” In: *Journal of Public Economics* 234.
- Tversky, Amos and Daniel Kahneman (1992). “Advances in prospect theory: Cumulative representation of uncertainty.” In: *Journal of Risk and Uncertainty* 5(4), pp. 297–323.
- Vespa, Emanuel and Alistair Wilson (2016). “Communication with multiple senders: An experiment.” In: *Quantitative Economics* 7(1), pp. 1–36.
- Vivalt, Eva and Aidan Coville (2023). “How do policymakers update their beliefs?” In: *Journal of Development Economics* 165.
- Weizsäcker, Georg (2023). *Misunderstandings: False Beliefs in Communication*. Open Book Publishers: Cambridge, UK.
- Wood, Daniel H. (2022). “Communication-Enhancing Vagueness.” In: *Games* 13(4), p. 49.
- Zhang, Sookie Xue and Ralph-Christopher Bayer (2023). “Delegation based on cheap talk.” In: *Theory and Decision* 94, pp. 333–361.

Appendix

A Additional Results

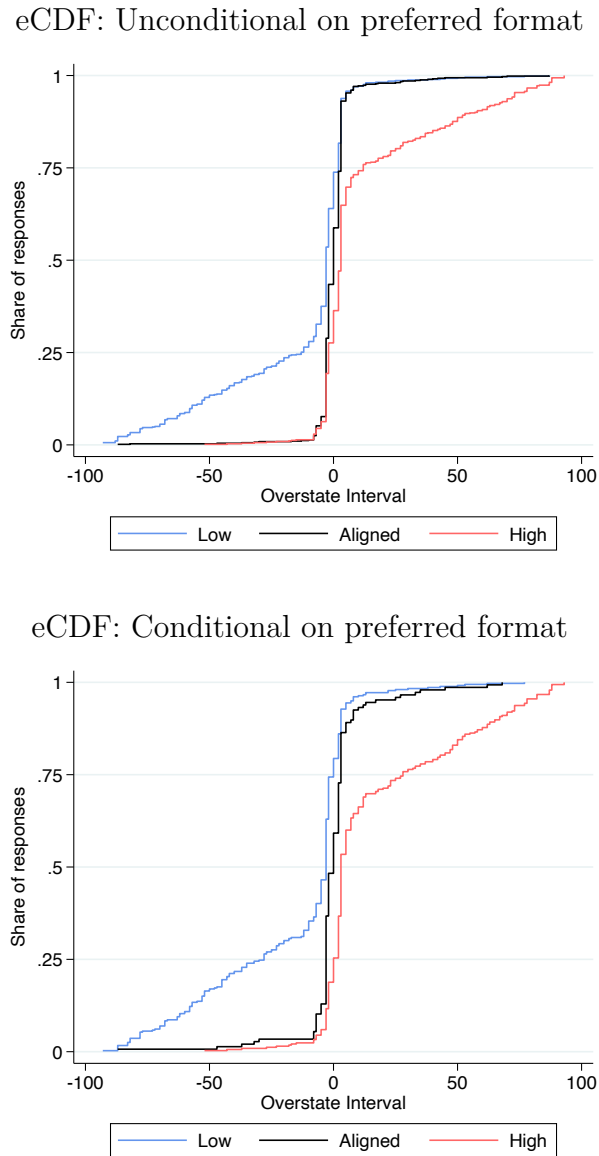
A.1 Study 1

Table A1: Language used in Study 1

Word (or Phrase)	Median response
Practically impossible	5%
Improbable	15%
Doubtful	20%
Unlikely	20%
Less likely than not	40%
About an even chance	50%
Possible	55%
A decent chance	65%
Probable	70%
Likely	75%
Expected	80%
Almost certain	90%
Practically guaranteed	95%


Notes: The left-hand column shows the complete list of words used in the drop-down list, as given to senders who were asked to select a language message. The right-hand column shows median responses to the benchmarking exercise among senders. For details on the benchmarking exercise see Section 2.2.2.

Figure A1: Cumulative Distribution Function of Slant by Incentives, for Interval Format Messages



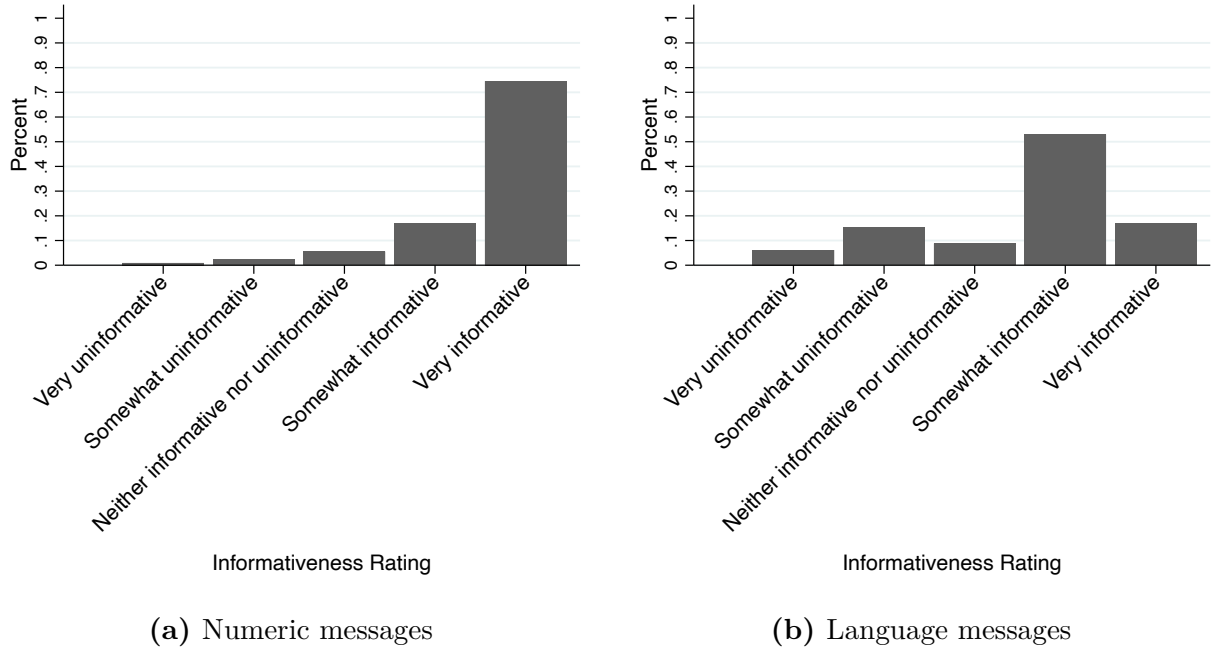
Notes: This figure plots the empirical cumulative distribution functions (eCDFs) for slant, by incentive conditions, for interval format messages. The x-axis, “Overstate interval” reflects the degree to which senders overstate their interval message compared to the truth. In particular, it takes the midpoint of a given interval message and then subtracts the true probability. The top panels show the distribution for all messages, while the bottom panels condition on the messages using the sender’s preferred format. The sample includes 1,992 choices across 249 senders. Similar eCDFs for language and numeric messages are shown in Figure 2.

Figure A2: Preferred Uses for Language



Notes: This figure plots the fraction of senders who preferred communicating likelihoods with language rather than numbers, for the following situations shown along the x axis: (1) wanting to communicate as accurately as possible; (2) wanting to signal to someone that one is well-informed; (3) making one's message easier to understand; (4) making it easier to persuade someone that the likelihood is higher than it actually is; (5) making it easier to lie to someone; (6) making it easier to withhold information.

Figure A3: Informativeness ratings by numeric and language messages



Notes: This figure plots the share of receivers selecting each respective option on a 5-point Likert scale, regarding the informativeness of the message they received from senders. (a) shows results for receivers who received numeric messages. (b) shows results for receivers who received language messages.

Table A2: Message slant by format, overall

	(1)	(2)	(3)
	Number - Prob	Language - Prob	Interval - Prob
High Incentives	10.712*** (0.942)	9.862*** (1.402)	11.541*** (1.321)
Low Incentives	-11.935*** (0.947)	-12.944*** (1.417)	-13.399*** (1.463)
Observations	4000	2008	1695
Aligned Mean	.13	.17	-.11
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes

Notes: This table reports the effect of incentives on message slant by format. In Column 1, the dependent variable *Number - Prob* is the number that the sender selected to communicate, minus the true probability they were given. In Column 2, the dependent variable *Language - Prob* is the language message selected for communication, mapped to numbers using the benchmarking exercise, minus the true probability. In Column 3, the dependent variable *Interval - Prob* is the the midpoint of the selected interval message, minus the true probability. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and incentive type. Column 1 is computed over the entire sample. Column 2 is computed over the sample of senders who had the option to communicate with language or numbers, and Column 3 is computed over the sample of senders who had the option to communicate with intervals or numbers. In the *High Incentives* row, the independent variable is an indicator equal to one when the sender faced directional-high incentives. In the *Low Incentives* row, the independent variable is an indicator equal to one when the sender faced directional-low incentives. Fixed effects for the respondent and true probability are included as controls. Aligned mean calculates the likelihood of using language among aligned individuals in the corresponding sample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A3: Slant of imprecise messages compared to numeric messages, by incentive

	(1)	(2)	(3)
	Language - Number	Interval - Number	Imprecise - Number
High Incentives	2.313 (1.446)	1.749*** (0.444)	2.331*** (0.771)
Observations	1329	1124	2242
Mean over Obs	-0.95	-0.03	-0.26
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes

Notes: This table considers slant between the sender’s chosen imprecise communication and their numeric message, and reports relative differences in slant by incentive type. In Column 1, the dependent variable is the language message selected for a communication, mapped to numbers using the benchmarking exercise, minus the numeric message. In Column 2, the dependent variable is the midpoint of the selected interval message minus the numeric message. Column 3 pools the differences across Columns 1 and 2. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and incentive type. In the row *High Incentives*, the independent variable is an indicator equal to one when senders faced directional-high incentives. These results were calculated over the subsample of senders who faced either a directional-high or a directional-low incentive (i.e. excluding aligned senders). We include fixed effects for the respondent and true probability as controls. Mean over obs reflects the average difference between the imprecise and numeric messages for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A4: Format choice by good state and bad state

	(1)	(2)	(3)
	Use Language	Use Intervals	Use Imprecise
High X Probability	-0.135*	-0.104	-0.118**
	(0.071)	(0.078)	(0.053)
Low X Probability	0.263***	0.170**	0.216***
	(0.074)	(0.083)	(0.055)
Observations	2008	1992	4000
Mean over Obs	0.30	0.42	0.36
Respondent FE	Yes	Yes	Yes

Notes: This table reports how format choice varies by being in a good or bad state of the world. For instance, if a sender has directional-high incentives, a good state is one where the true probability is actually high; a bad state is one where the true probability is actually low. In Column 1, the dependent variable is an indicator equal to one when the sender chose to communicate using language rather than numbers. In Column 2, the dependent variable is an indicator equal to one when the sender chose to communicate using intervals rather than numbers. Columns 1 is calculated over the sample of senders who had the option to communicate in language or numbers; Columns 2 is calculated over the sample of senders who had the option to communicate in intervals or numbers. Column 3 pools together Columns 1 and 2, with the dependent variable now being an indicator which equals one if the sender communicated with any imprecise message type, either language or intervals. In the row *High X Probability*, the independent variable is an interaction between the true probability, and an indicator which equals one if the sender had directional-high incentives. In the row *Low X Probability*, the independent variable is an interaction between the true probability, and an indicator which equals one if the sender had directional-low incentives. We control for respondent fixed effects. Mean over obs reflects the likelihood of using an imprecise message format in the corresponding subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A5: Effect of senders' incentives on receivers' guesses

	(1)	(2)	(3)
	Guess (Language)	Guess (Interval)	Guess (All)
High Incentives	14.814*** (2.199)	20.077*** (2.552)	17.045*** (1.682)
Observations	703	589	1292
Overall Mean	51.48	51.09	51.31
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes

Notes: This table shows that receivers are persuaded in the direction of senders' incentives. Our results are only calculated over conditions where the sender faces directional-high or directional-low incentives (i.e. we exclude conditions with aligned senders). The independent variable is an indicator which equals one if the sender faced directional-high incentives. Column 1 restricts to receivers who either receive language or number messages; Column 2 restricts to receivers who either receive interval or number messages. Column 3 pools together Columns 1 and 2. The dependent variable is the receiver's guess of the true probability. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and sender incentive type. We include fixed effects for the respondent and true probability as controls. *Overall Mean* reflects the mean guess among the receivers included in each column. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A6: Effect of senders' incentives on receivers' guesses by format

	(1)	(2)	(3)
	Guess (Language)	Guess (Interval)	Guess (All)
Imprecise Sent X High	12.497*** (2.370)	3.770* (1.990)	7.161*** (1.522)
Imprecise Sent X Low	-5.024* (2.554)	-3.378 (2.446)	-4.850*** (1.742)
High Incentives	7.308*** (1.522)	8.272*** (1.466)	8.043*** (1.047)
Low Incentives	-7.969*** (1.404)	-11.735*** (1.641)	-9.531*** (1.070)
Observations	1984	2016	4000
Overall Mean	51.97	50.82	51.37
Format Control	Yes	Yes	Yes
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes

Notes: This table shows that receivers are persuaded more in the direction of senders' incentives when they receive imprecise messages. Column 1 restricts to receivers who either receive language or number messages; Column 2 restricts to receivers who either receive interval or number messages. Column 3 pools together Columns 1 and 2. The independent variable of interest is the interaction between the sender's incentive, and an indicator which equals one if the sender sent an imprecise message (language in Column 1; intervals in Column 2; either language or intervals in Column 3). The dependent variable is the receiver's guess of the true probability. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and sender incentive type. We include fixed effects for the respondent and true probability as controls. *Overall Mean* reflects the mean guess among the receivers included in each column. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A7: Effect of senders’ incentives and receivers’ knowledge on receivers’ guesses

	(1)	(2)	(3)
	Guess (Language)	Guess (Interval)	Guess (All)
Know X High	-7.204** (3.325)	-4.288 (3.211)	-5.647** (2.315)
High Incentives	26.446*** (2.158)	26.154*** (1.985)	26.250*** (1.465)
Observations	1276	1308	2584
Overall Mean	55.96	48.84	51.97
Respondent FE	Yes	Yes	Yes
Probability FE	Yes	Yes	Yes

Notes: This table shows that receivers are somewhat less persuaded by senders when receivers have knowledge of sender incentives. Our results are only calculated over conditions where the sender faces directional-high or directional-low incentives (i.e. we exclude conditions with aligned senders). Column 1 restricts to receivers who either receive language or number messages; Column 2 restricts to receivers who either receive interval or number messages. Column 3 pools together Columns 1 and 2. The independent variable of interest is the interaction between two indicators: an indicator which equals one if the sender had high incentives, and an indicator which equals one if receivers know the incentives of their paired sender. The dependent variable is the receiver’s guess of the true probability. Dependent variables are winsorized at the 2% and 98% level, with a separate winsorization calculated over each combination of true probability and sender incentive type. We include fixed effects for the respondent and true probability as controls. *Overall Mean* reflects the mean guess among the receivers included in each column. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

A.2 Study 2

Table A8: Relative slant by researchers' preferred format

	(1)
	log(Lang)-log(Num)
Use Language X Directional Incentives	0.494** (0.243)
Observations	289
Mean over obs	-0.22
Subject controls	Yes
Study FE	Yes

Notes: This table reports the relative difference in slant between the language and number messages, by incentive condition and whether the researcher preferred communicating using language. The dependent variable is the language message selected for a communication, mapped to numbers using the benchmarking exercise and logged, minus the logged numeric message. *Use Language X Directional Incentives* interacts an indicator equal to one when the sender preferred the language message with an indicator equal to one for directional-high incentives, controlling for the these indicators. Subject controls include indicators for whether the researcher is tenure-track, male, and does empirical research. We also control via fixed effects for each of the six real research studies whose results were being communicated in the experiment. *Mean over obs* reflects the average difference between the (logged) language and number messages for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A9: Lee bounds for effect of directional incentives on researchers’ use of language

	(1)	(2)	(3)
	Main specification	Lee lower bound	Lee upper bound
Directional Incentives	0.291*** (0.054)	0.284*** (0.055)	0.335*** (0.051)
Observations	290	278	278
Mean over obs	0.26	0.27	0.24
Subject controls	Yes	Yes	Yes
Study FE	Yes	Yes	Yes

Notes: Here we report Lee bounds for the effect of directional incentives on researchers’ use of language, an effect previously reported in Table 6 of the main paper. Our sample of researcher exhibited differential attrition. In the control group, 1 respondent attrited and 76 did not, for a response rate of 98.7%. In the treatment group, 7 respondents attrited and 69 did not. We assume that some respondents may have attrited after being assigned to the treatment group, because they did not want to proceed with the task of persuading policymakers. The difference in attrition rates between treatment and control was 7.9%. Thus we seek to trim $7.9/98.7 = 8\%$ of control respondents, i.e. 6 control respondents. To compute the lower bound, we randomly drop 6 control respondents who chose to communicate with numbers for both studies. Regarding the upper bound, there were only 3 control respondents who chose to communicate with language for both studies. We dropped these, and then on top of this randomly dropped a further 3 control respondents who communicated with language for one study and numbers for another. As a baseline, Column 1 displays the effect from the main specification reported in Table 6 of the main paper. Column 2 reports the Lee lower bound, and Column 3 reports the Lee upper bound. In all columns, the dependent variable is an indicator which equals one when the researcher chose to communicate with language rather than numbers. The independent variable is an indicator which equals one when researchers had directional-high incentives. Subject controls include indicators for whether the researcher is tenure-track, male, and does empirical research. We also control via fixed effects for each of the six real research studies whose results were being communicated in the experiment. *Mean over obs* reflects the likelihood of using language for that subsample. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

Table A10: Effect of researchers' messages on policymakers' guesses

	(1)	(2)	(3)	(4)
	Posterior	log(Posterior)	Posterior	log(Posterior)
Directional-High	2.200** (0.992)	0.200** (0.097)		
Language Message			3.316* (1.947)	0.206 (0.200)
Prior	0.374*** (0.118)		0.609*** (0.166)	
log(Prior)		0.441*** (0.076)		0.565*** (0.137)
Observations	129	129	57	57
Mean over obs	8.56	8.56	9.93	9.93
Study FE	Yes	Yes	Yes	Yes
Sample	All	All	Directional-High	Directional-High

Notes: In Columns 1 and 2, we show that policymakers' guesses are influenced upward when receiving messages from researchers with directional-high incentives. These two columns are calculated over the whole sample. In Column 1, the dependent variable *Posterior* is the policymaker's guess after seeing the researcher's message. The independent variable of interest is *Directional-High*, an indicator function that equals one when researchers had directional-high incentives. We also control for *Prior*, which is the policymaker's initial guess, elicited before they see the researcher's message. Column 2 is similar to Column 1, but instead using the natural log of the *Posterior* and *Prior* variables. Columns 3 and 4 are only computed over settings where researchers had directional-high incentives. Here we show that policymakers' guesses are more strongly influenced by language than numerical messages. In Columns 3 and 4, the respective dependent variables are *Posterior* and *log(Posterior)* as before, and we also respectively control for *Prior* and *log(Prior)* as before. However, the independent variable of interest is now *Language Message*, an indicator variable that equals one if the researcher sends a language message. As controls, we add fixed effects for each of the six real research studies whose results were being communicated in the experiment. The *Prior* and *Posterior* variables are winsorized at the 10% and 90% level. A separate winsorization was calculated for responses pertaining to each of the six real research studies. *Mean over obs* is the mean posterior over the senders included in each column. Standard errors, clustered at the individual level, are reported in parentheses. ***, **, * denote that estimates are statistically significant at the 1%, 5%, and 10% levels, respectively.

B Online Appendix: Study Materials

B.1 Study 1: Senders

Sender Instructions, version for language vs numbers

Survey overview

You will see several pages of questions in four parts. Parts 2 and 3 are the "main parts". You will learn about each part when you are at that point in the survey.

There will be an "attention check" question in the study. The answer to this question will be obvious to anyone paying attention. Participants who do not answer the attention check question correctly **will not be eligible for a bonus payment.**

Bonus payments

After all participants complete the study, 10 participants will be chosen at random to receive a bonus payment of up to \$50 based on their choices in Parts 2 and 3. In particular, **if you are randomly selected to receive a bonus payment, we will randomly choose one of the questions in Parts 2 and 3 to determine your payment.** The high bonus is because it is important for us that you take this study seriously.

In Part 2, you will be assigned the role of a "Sender" and play 8 rounds in which you will make decisions about how you would communicate information to different Prolific participants ("Receivers").

First, you will complete Part 1, which is a practice round for Part 2. You will see example questions for your role as a Sender and for Receivers so that you fully understand both roles before you proceed.

You may now click to proceed to Part 1.

Part 2: Sender Instructions

You are now ready for Part 2 of the survey.

In case you've forgotten the details, please refer below for the instructions for the role of Sender, to which you are assigned. You will find additional information about how bonus payments are determined.

In this role, you will be shown some information about the chance of an event occurring. The questions will then prompt you to communicate this information to your paired Receiver, who will be randomly matched with you after you submit your responses.

There will be two message formats through which to communicate your information to your paired Receiver:

1. **Number:** You will choose a message that says "The chance that you will draw a red ball is [X] percent." for your choice of X.
2. **Language:** You will choose a message that says "It is [WORD] that you will draw a red ball." for your choice of **WORD** from a dropdown menu.

You will first be asked how you would communicate to the Receiver, given each format. For instance, you may choose "64 percent" for **Number** and "probable" for **Language**. Your message can reflect the information you see in the slider, but it does not have to.

Then, you will answer a question about which of these **formats** you would prefer to send. Your preferred format is more likely (but not guaranteed) to be the one communicated to your paired Receiver.

Your bonus payment will depend on what your Receiver predicts after they see your message. On each question, you will learn whether your bonus payment will be higher if your Receiver:

- **Correctly predicts** the chance of drawing a **RED** ball; the probability you will earn a bonus will be equal to the receiver's probability of earning a bonus.
- Predicts that the share of **RED** balls is **HIGH**; the probability you will earn a bonus will be equal to the percent chance your Receiver assigns to drawing a **RED** ball.
- Predicts that the share of **RED** balls is **LOW**; the probability you will earn a bonus will be equal to the percent chance your Receiver assigns to drawing a **BLUE** ball.

Sender Main Decisions, version with language vs numbers

Note that the comprehension check, and 'explain your reasoning' prompts, are only displayed for the first of eight questions.

The share of red balls in the box can take the following possible values: 2, 8, 17, 25, 33, 42, 50, 58, 67, 75, 83, 92, 98.

The information about the payment structure can take three different forms, based on which incentives are selected:

- “If this question is selected for payment, you will be more likely to earn the bonus if your Receiver predicts that the chance of drawing a **RED** ball is **HIGH**”
- “If this question is selected for payment, you will be more likely to earn the bonus if your Receiver predicts that the chance of drawing a **RED** ball is **LOW**”
- “If this question is selected for payment, you will be more likely to earn the bonus if your Receiver’s prediction of the chance of drawing a **RED** ball is **MORE ACCURATE**”

Question 1 of 8

Your payment: If this question is selected for payment, you will be more likely to earn the bonus if your Receiver predicts that the chance of drawing a **RED** ball is **LOW**.

Here is the slider showing the share of red balls:

0 10 20 30 40 50 60 70 80 90 100

Share of red balls in box (out of 100)



COMPREHENSION CHECK: Which of the following reflects how your bonus payment will be determined?

I will be more likely to win the bonus if my Receiver correctly predicts the chance of drawing a red ball

I will be more likely to win the bonus if my Receiver predicts that the chance of drawing a red ball is high

I will be more likely to win the bonus if my Receiver predicts that the chance of drawing a red ball is low

Number: Suppose you are assigned to send a number message to your paired Receiver. You can choose one number between 0-100%. Which message would you like to send?

"The chance that you will draw a red ball is _____ percent."

Language: Suppose you are assigned to send a natural language message to your paired Receiver. Which message would you like to send?


"It is _____ that you will draw a red ball."

Why did you make your number and language choices in this way? (You will only be asked to explain your reasoning for this first question.)

Recall, this is the slider that you saw on the previous page:

0 10 20 30 40 50 60 70 80 90 100

Share of red balls in box (out of 100)



Format

We will now ask you a question about your preferred message **format**: Number or Language.

Remember that: **If this question is selected for payment, you will be more likely to earn the bonus if your Receiver predicts that the chance of drawing a RED ball is LOW.**

Which of the following would you prefer to send to your paired receiver?

Number: "The chance that you will draw a red ball is 2 percent."

Language: "It is probable that you will draw a red ball."

Why did you choose between number and language messages in this way? (You will only be asked to explain your reasoning for this first question.)

Benchmarking Exercise

Part 3: Words and probabilities

You are now ready for Part 3 of the survey. The following table shows 13 words or phrases. We want to know what numeric probabilities people think each word or phrase corresponds to when describing the chance of an event happening.

Specifically, you will guess what percent chance people thought each word or phrase corresponded to (between 0 and 100 percent). At the end of the survey, we will compare each of your answers to the average answers given by all other participants. If a question is randomly selected for payment, you will earn the bonus if your answer is within is **within 3 percentage points of the average response. So, it is in your best interest to give guesses that you think will be similar to those of other participants.**

	Percent chance
practically impossible	<input type="text"/>
practically guaranteed	<input type="text"/>
a decent chance	<input type="text"/>
about an even chance	<input type="text"/>
less likely than not	<input type="text"/>
improbable	<input type="text"/>
unlikely	<input type="text"/>
almost certain	<input type="text"/>
likely	<input type="text"/>
doubtful	<input type="text"/>
expected	<input type="text"/>
probable	<input type="text"/>
possible	<input type="text"/>

B.2 Study 1: Receivers

Receiver Instructions, version for language vs numbers

Part 2: Receiver Instructions

You are now ready for Part 2 of the survey. Your answers to this section may affect your bonus payment.

In case you've forgotten the details, please refer below for the instructions for the role of Receiver, to which you are assigned. You will also find some additional information below.

In this role, in each question that is selected for a bonus you will be randomly paired with another Prolific worker (your Sender). A computer will simulate a random draw of one ball from a box that contains 100 balls in total, some of which are red and the rest of which are blue. **Your objective is to correctly predict the chance of drawing a red (vs. blue) ball from the box.**

In each question, the box will have a random number of red and blue balls. Some boxes have more red balls, and some have more blue balls. The chance of drawing a red ball in one box does not affect the chance of drawing a red ball in another box.

You may use information sent to you by the Sender to impact your prediction.

You will receive one of two types of messages from your Sender about the chance of drawing the red ball:

1. **Number:** "The chance that you will draw a red ball is **[X] percent.**"
2. **Language:** "It is **[WORD]** that you will draw a red ball."

You will only see one message. For instance, you may see "The chance that you will draw a red ball is 64 percent." or "It is probable that you will draw a red ball.", but not both.

After receiving the message from your Sender, you will predict the percent chance of drawing a red ball from a box. This is the question that determines your bonus payment. If you would like details on the specific payment rule, see below: the important thing to know is that **the payment rule is carefully designed so that you will be most likely to earn the \$50 bonus payment if you report an accurate guess about the chance of drawing a red ball.**

Recall, your Sender's bonus payment is determined by the prediction you give. You will either be matched with a Sender who is more likely to win the bonus when you:

- **Correctly predict** the chance of drawing a **RED** ball;
- Predict that the chance of drawing a **RED** ball is **HIGH**; OR
- Predict that the chance of drawing a **RED** ball is **LOW**.

You **will not be informed** about how the Sender's bonus payment is determined for each message you see.

Note that when receivers are informed (rather than uninformed) about sender incentives, the final line of instructions above will read: “You **will be informed** about how the Sender’s bonus payment is determined for each message you see.”

Payment rule:

If you are selected to receive a bonus payment, to determine your payment the computer will randomly choose a question and then randomly draw two numbers. For each draw, all numbers between 0 and 100 (including decimal numbers) are equally likely to be selected. Draws are independent in the sense that the outcome of the first draw in no way affects the outcome of the second draw. If this question is chosen for payment, then:

- If the ball from the chosen question is red and the number you picked for red is *larger* than either of the two draws, you will get the \$50 bonus payment.
- If the ball from the chosen question is blue and the number you picked for red is *smaller* than either of the two draws, you will get the \$50 bonus payment.
- Otherwise, you will not get the bonus payment.

Next, you will take on the role of Receiver and predict the share of red balls in 8 different questions.

When you are ready to begin, please click to continue.

Receiver Main Decisions

Note that when receivers are informed about sender incentives, the line beginning “**Your Sender was informed about their bonus payment:**” can continue in three different ways, depending on the actual sender’s incentives:

- “Your Sender is more likely to receive the bonus if you predict that the chance of drawing a **RED** ball is **HIGH**.”
- “Your Sender is more likely to receive the bonus if you predict that the chance of drawing a **RED** ball is **LOW**.”
- “Your Sender is more likely to receive the bonus if your prediction of the chance of drawing a **RED** ball is **MORE ACCURATE**.”

When receivers are uninformed (rather than informed) about sender incentives, the line beginning “**Your Sender was informed about their bonus payment:**” will be omitted.

Question 1 of 8

Your bonus payment: You are more likely to earn the bonus if your prediction of the chance of drawing a **RED** ball is **MORE ACCURATE**.

Your Sender was informed about their bonus payment: Your Sender is more likely to receive the bonus if you predict that the chance of drawing a **RED** ball is **HIGH**.

Your Sender shared the following message with you:

“The chance that you will draw a red ball is 50 percent.”

Given this information, what do you think the percent chance is that you will draw a red ball?

0 10 20 30 40 50 60 70 80 90 100

Percent chance of drawing a red ball

Why did you make this prediction? (You will only be asked to explain your reasoning for this first question.)

How **informative** do you think your Sender's message was on this question?

- Very informative
- Somewhat informative
- Neither informative nor uninformative
- Somewhat uninformative
- Very uninformative

B.3 Study 2: Researchers

B.3.1 Recruitment Email

Survey on communicating research results

From Toma, Mattie <Mattie.Toma@wbs.ac.uk>

Date Thu 18/07/2024 22:13

[REDACTED]

[REDACTED]

I hope you're doing well!

I'm running a survey (with Michael Thaler and Victor Wang) to investigate how researchers communicate the results of research studies to policymakers. **Would you be willing to help with this research by filling out this survey, which should take about 10 minutes to complete?**

This survey link is unique to you and so should not be shared with others:

https://wbs.qualtrics.com/jfe/form/SV_a5bxZLIeZ15IQO?

[Q_CHL=gl&Q_DL=EMD_yfMUICrPqU7exVs_a5bxZLIeZ15IQO_CGC_dyikdZ75XtnYy02&g_=g](https://wbs.qualtrics.com/jfe/form/SV_a5bxZLIeZ15IQO?Q_CHL=gl&Q_DL=EMD_yfMUICrPqU7exVs_a5bxZLIeZ15IQO_CGC_dyikdZ75XtnYy02&g_=g)

We very much appreciate your help. As a token of our appreciation, everyone who completes the survey will be sent a **\$10 Amazon gift card**, and you will have the opportunity to earn an additional \$10 depending on your responses in the survey.

All data will be kept as confidential as possible, in accordance with Warwick ethics protocol HSSREC 219.23-24.

We will close the survey on Tuesday, July 23.

Thanks so much in advance for taking the time to read this email and for considering this!
Mattie

Mattie Toma

Assistant Professor | Warwick Business School |
Behavioural Science Group | [University of Warwick](https://www.warwick.ac.uk) | Coventry | CV4 7AL |
Mattie.Toma@wbs.ac.uk

<https://www.mattietoma.com/>



Instructions

Survey overview

In this survey, you will make decisions about how you, as a researcher, want to communicate the results of two research studies in which a policy intervention was tested. You are deciding how you want to communicate with real policymakers. Policymakers are recruited from a sample of high-ranking public servants worldwide involved in policy decision-making.

Each policymaker will be randomly paired with a researcher. We expect that about half of researchers will be paired with policymakers.

Bonus payments

After all participants complete the survey, all researchers who are paired with a policymaker will receive a bonus payment of up to \$/£10 based on their choices in the survey. In particular, **if you are selected to receive a bonus payment, we will randomly choose one question to determine your payment.** We will describe details of these incentives later.

Payment will be provided in the form of either a UK Amazon gift card (in £) or a US Amazon gift card (in \$). At the end of the survey, you will have the option to choose which gift card you would rather receive.

Instructions

You are tasked with communicating the results of research studies testing a policy intervention to your paired policymaker. Your paired policymaker will also learn about the studies, but not the results.

First, you will learn more about each study and its results on the following pages. Then, you will choose a message about the results to send to your paired policymaker.

There will be two message formats you can use to communicate:

1. **Number:** You will choose a message that uses numbers to describe the research results.
2. **Language:** You will choose a message that uses a word or phrase to describe the research results.

You will first be asked how you would communicate with the policymaker, given each format. For instance, you may say “The intervention led to a 3 percentage point increase” for **Number** and “The intervention led to a moderate increase” for **Language**. Your messages can reflect what you learned about the results of the research studies, but they do not have to.

Then, you will answer a question about which of these two formats you would prefer to send. We will always send one of your two messages to the policymaker, and will send the message using your preferred format more often, so it is in your interest to answer both of these questions honestly and carefully.

Note that the final part of the instructions page, about bonus payments, reads differently depending on whether senders are assigned to the aligned or directional incentive treatment.

Bonus payment instructions for aligned treatment below:

Your bonus payments

If you are selected to receive a bonus payment on a given question, your bonus payment will depend on the decisions your paired policymaker makes based on the message you send them.

In particular, during this survey we are asking you to imagine you are the author of the study testing the policy intervention and **you are trying to give the policymaker the best understanding of the data to inform their policy decisions.**

To ensure that your incentive is consistent with this hypothetical objective, **you are more likely to earn the bonus payment if your policymaker correctly predicts the intervention's effect size.** Specifically, if the true effect size is X percentage points, you will win the bonus payment if your paired policymaker predicts the effect size is between $X-1$ and $X+1$ percentage points. Your paired policymaker's bonus payment will be the same as yours: Their bonus payment will depend on correctly guessing the results of the study.

Note that the policymakers in this survey do not know whether you are incentivized to have them form accurate beliefs or to persuade them that the effects are large.

Bonus payment instructions for directional treatment below:

Your bonus payments

If you are selected to receive a bonus payment on a given question, your bonus payment will depend on the decisions your paired policymaker makes based on the message you send them.

In particular, during this survey we are asking you to imagine you are the author of the study testing the policy intervention and **you are trying to persuade a policymaker that the research is promising to increase your chance of getting a government grant or policymaker attention.**

To ensure that your incentive is consistent with this hypothetical objective, **you are more likely to earn the bonus payment if your paired policymaker predicts the intervention's effect size is larger.** Specifically, if your paired policymaker predicts the effect size is X percentage points, the probability you will win the bonus payment will be equal to X%, with a minimum probability of 0%. Your paired policymaker's bonus payment will depend on the accuracy of their predictions.

Note that the policymakers in this survey do not know whether you are incentivized to have them form accurate beliefs or to persuade them that the effects are large.

Main decisions, Adams et al (2021) study

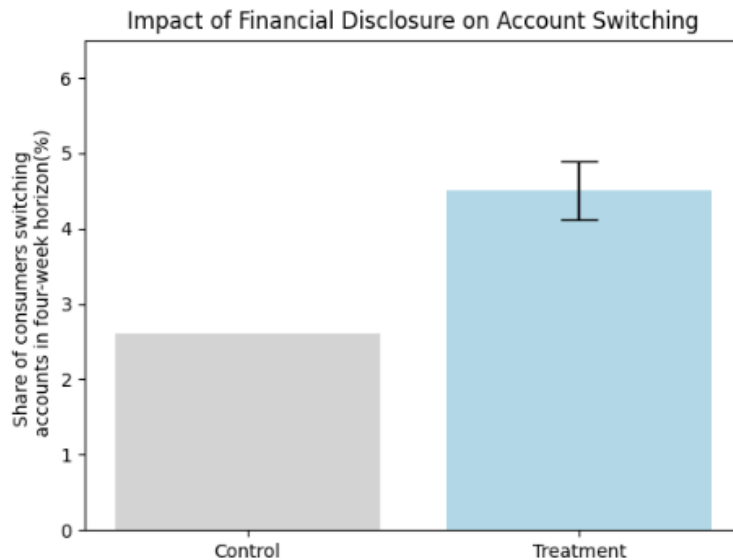
Note that the reminders about the bonus payment structure will display differently depending on whether the sender was assigned to the aligned or directional treatment arm. For senders in the aligned treatment, the reminder reads: “**Reminder:** When making your decisions, imagine you are the author of this study and you are trying to **give the policymaker the best understanding of the data.** In this decision you are more likely to earn the bonus payment if **your paired policymaker correctly predicts the intervention's effect size.**”

For senders in the directional treatment, the reminder reads: “**Reminder:** When making your decisions, imagine you are the author of this study and you are trying to **persuade a policymaker that the research is promising**. In this decision you are more likely to earn the bonus payment if **your paired policymaker predicts the intervention’s effect size is larger.**”

Research study information

In a study reported in the *Journal of Financial Economics* (Adams et al. 2021), consumers with savings accounts at five UK depositaries were randomly disclosed information about alternative financial products. This information was provided on the front page of their annual statements. The study examined whether this disclosure would make consumers more likely to switch between financial products.

The graph below shows the estimated effect of consumer financial disclosure on the share of consumers switching accounts in a four-week horizon.



The graph includes point estimates for the control and treatment groups, and a 95% confidence interval for the treatment effect.

Reminder: When making your decisions, imagine you are the author of this study and you are trying to **persuade a policymaker that the research is promising**. In this decision you are more likely to earn the bonus payment if **your paired policymaker predicts the intervention's effect size is larger**.

Message choice

Number: Please choose a number message to communicate the treatment effect of the policy:

“In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to an increase in this share of ___ percentage points. That is, 2.6 + [Your Response]% of those receiving the financial disclosure treatment switched between financial products.”

percentage points

Language: Please choose a word or phrase to communicate the treatment effect of the policy:

“In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to a(n) ___ increase in the share of consumers switching between financial products.”

Note that the available options for the ‘choice of format’ question (below) will repeat the message choices made by the subject in the previous prompt.

Choice of format

We will now ask you to choose your preferred message format: Number or Language. Which of the following would you prefer to send?


Number format: "In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to an increase in this share of 2 percentage points. That is, 4.6% of those receiving the financial disclosure treatment switched between financial products."

Language format: "In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to a(n) modest increase in the share of consumers switching between financial products."

Reminder: When making your decisions, imagine you are the author of this study and you are trying to **persuade a policymaker that the research is promising**. In this decision you are more likely to earn the bonus payment if **your paired policymaker predicts the intervention's effect size is larger**.

Study reminder: In a study reported in the Journal of Financial Economics (Adams et al. 2021), consumers with savings accounts at five UK depositaries were randomly disclosed information about alternative financial products. This information was provided on the front page of their annual statements. The study examined whether this disclosure would make consumers more likely to switch between financial products.

The graph below shows the estimated effect of consumer financial disclosure on the share of consumers switching accounts in a four-week horizon.



The graph includes point estimates for the control and treatment groups, and a 95% confidence interval for the treatment effect.

Did you know the results from this research study (Adams et al, 2021) prior to this survey?

- Yes, I already knew the main effects quantitatively
- Yes, I was familiar with the results but not the precise quantitative effects
- No, I was familiar with the paper but did not know the main results
- No, I was not familiar with the paper

B.4 Study 2: Policymakers

Instructions

Survey overview

In this survey, you, as a **policymaker**, will be paired with a **researcher**. You will receive information about two research studies from your paired researcher, and you will be asked to predict the results of the studies. Researchers were recruited from a sample of academics in the social sciences, most of whom are tenured or tenure-track professors of economics.

Bonus payments

After everyone completes the survey, you may receive a bonus payment of \$/£5 based on your choices in the survey. **We will randomly choose one question to determine your bonus payment.**

Payment will be provided in the form of either a UK Amazon gift card (in £) or a US Amazon gift card (in \$). At the end of the survey, you will have the option to choose which gift card you would rather receive. Note that this survey is funded jointly by the British Academy and the University of Warwick.

Instructions

On the following pages, you will learn about two research studies in which a policy intervention was tested.

You will not immediately learn the results of the studies. Instead, for each study, you will be paired with an academic researcher who has already completed a version of this survey. Your paired researcher was told what the effect size of the policy intervention was, and chose how to communicate this information to you.

In particular, you will receive one of two types of messages from the researcher about the effect of the intervention:

1. Either **Number**: The message will use a number to describe the effect of the intervention (e.g., "the intervention led to a 3 percentage point increase")
2. Or **Language**: The message will use a word or phrase to describe the effect of the intervention (e.g., "the intervention led to a moderate increase")

Your task will be to predict the effect size of each policy intervention. You will make predictions both before and after you see the researcher's message.

The closer your predictions are to the true effect size, the more likely it is that you will receive the bonus payment of \$/£5.

We will tell you the actual results of the research at the end of the survey.

Half of the researchers were told to try to accurately communicate the effect size of the intervention, and the other half were told to try to persuade you that the effect size of the intervention was large. All researchers chose one number and language message to send, and you will see the format they preferred most of the time. **Researchers were told that their messages could reflect what they learned about the intervention's effects, but they did not have to.**

Main decisions overview

Main Decisions

Now you're ready to make the main decisions in this survey!

Which topics do you think best reflect your areas of expertise? **Please choose two options below.** We will show you research studies that are relevant to these areas.

Health insurance	<input type="checkbox"/>
Economic mobility and housing	<input type="checkbox"/>
Education policy	<input type="checkbox"/>
Consumer finance and savings	<input type="checkbox"/>
Criminal justice	<input type="checkbox"/>
Labor market and aging	<input type="checkbox"/>

After this page, subjects are taken to the following main decision sections depending on which policy areas they indicated:

- Health insurance: Domurat et al (2021)
- Economic mobility and housing: Bergman et al (2024)
- Education policy: Burland et al (2023)
- Consumer finance and savings: Adams et al (2021)
- Criminal justice: Fishbane et al (2020)

- Labor market and aging: Liebman and Luttmer (2015)

These decision sections are shown in the following pages.

Main decisions, Adams et al (2021)

Question 1 of 2: Research study information

In a study reported in the *Journal of Financial Economics* (Adams et al. 2021), consumers with savings accounts at five UK depositaries were randomly disclosed information about alternative financial products. This information was provided on the front page of their annual statements. The study examined whether this disclosure would make consumers more likely to switch between financial products.

Before we share the message from your paired researcher, we'd like to ask you for your **best guess of the treatment effect of the policy intervention?**

In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to an increase in this share of ___ percentage points. This means that $2.6 + [\text{Your Response}]$ % of those receiving the financial disclosure treatment switched between financial products.

Your researcher's message

Next, you will see the message from your paired researcher and have the opportunity to revise your prediction.

Your Researcher's message:

In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to a considerable increase in this share.

Study reminder: In a study reported in the *Journal of Financial Economics* (Adams et al. 2021), consumers with savings accounts at five UK depositaries were randomly disclosed information about alternative financial products. This information was provided on the front page of their annual statements. The study examined whether this disclosure would make consumers more likely to switch between financial products.

Now, what is your best guess of the treatment effect of the policy intervention?

In the control group, 2.6% of consumers switched between financial products in a four-week horizon. The financial disclosure treatment led to an increase in this share of ___ percentage points. This means that $2.6 + [\text{Your Response}]$ % of those receiving the financial disclosure treatment switched between financial products.

How likely would you be to share information about this study with your colleagues?

Very unlikely	<input type="radio"/>
Somewhat unlikely	<input type="radio"/>
Neither likely nor unlikely	<input type="radio"/>
Somewhat likely	<input type="radio"/>
Very likely	<input type="radio"/>

Did you know the results from this research study (Adams et al, 2021) prior to this survey?

Yes	<input type="radio"/>
No	<input type="radio"/>

Would you like to access an infographic designed for policymakers that provides more details on this study? If so, please click yes and we will share a link to a downloadable infographic at the end of the survey. (Note that clicking yes will slightly increase the survey length.)

Yes, I would like an infographic about this study	<input type="radio"/>
No, i would not like an infographic about this study	<input type="radio"/>

B.5 Infographics for Policymakers

This is an example of the infographics offered to policymakers in the survey “Study 2: Receivers”.

