Spin Doctors: An Experiment on Vague Disclosure

Marvin Deversi (LMU Munich)
Alessandro Ispano (CY Cergy Paris Université, CNRS and THEMA)
Peter Schwardmann (LMU Munich)

Discussion Paper No. 304

December 1, 2021

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Georg Weizsäcker, Ph.D., Humboldt University Berlin, 10117 Berlin, Germany
info@rationality-and-competition.de

Spin Doctors: An Experiment on Vague Disclosure

Marvin Deversi Alessandro Ispano Peter Schwardmann

July 2021

Abstract

Unfavorable news are often delivered under the disguise of vagueness. Our theory-driven laboratory experiment investigates this strategic use of vagueness in voluntary disclosure and asks whether there is scope for policy to improve information transmission. We find that vagueness is profitably deployed by senders to fool those receivers that lack strategic sophistication. Imposing precise disclosure leads to more easily interpretable messages, but results in fewer sender types disclosing at all. Since non-disclosure also systematically misleads naive receivers, the welfare implications of imposing precision are not obvious. However, our model and experiment show that information transmission and the welfare of naive receivers are improved by policies that impose precision. Our results speak to the rules governing firms' disclosure of quality-relevant information, the disclosure of research findings, and testimonies in a court of law.

Keywords: communication, naivete, flexibility, regulation
JEL classifications: D82, D83, C72, C92, L15, D04

Deversi: University of Munich (LMU), marvin.deversi@econ.lmu.de; Ispano: CY Cergy Paris Université, CNRS, THEMA, alessandro.ispano@gmail.com; Schwardmann: University of Munich (LMU), peter.schwardmann@econ.lmu.de. We thank Ron Berman, Milo Bianchi, Jeanne Hagenbach, Martin Kocher, Matthias Lang, David Laibson, George Loewenstein, John List, Michael Luca, Daniel Martin, Klaus Schmidt, Simeon Schudy, Stefan Trautmann, Carolin Wagner and seminar audiences at the U Amsterdam, U Vienna, TU Munich, SITE Experimental Economics workshop, EEA Congress in Cologne, 11th Workshop on the Economics of Advertising and Marketing at Columbia U, CESifo Behavioral Eonomics conference, RGS Workshop at the U Essen, Incentives workshop U Osaka, SABE/IAREP Conference at Middlesex U, Workshop on Limited Attention at the U Copenhagen, ESA World Meeting in San Diego, AFSE Meeting at PSE for helpful comments. We acknowledge support through DFG grant CRC TR190 and ANR grant 19-CE26-0010-03. Deversi acknowledges funding through the doctoral program "Evidence-Based Economics" and DFG grant KO 4100/1-1. Ispano acknowledges support from Labex MME-DII.

1 Introduction

In many settings, informed parties not only decide whether to disclose verifiable private information, but also enjoy substantial flexibility in how information is disclosed. One way to exploit flexibility in disclosure is by means of vague messages. Vague messages are designed to inflate a receiver's perception of the sender's type by clearly separating from worse but not from better types. They are not outright lies, which may invite litigation, but merely put a positive spin on unfavorable news. Consider the following examples.

A college that ranks 10th in the latest US news ranking is likely to call itself a top 10 college rather than referring to itself as the 10th ranked college. A wine whose sole designation of origin is France is unlikely to come from the Bordeaux region, renowned for its superior wine. A wine whose sole designation of origin is Bordeaux is unlikely to come from Pomerol, an especially beloved subregion of Bordeaux. Researchers often refer to "significance at the 5 percent level" when a p-value is just below 0.05, while stating the exact p-value for a highly significant result. During legal proceedings, a defendant may try to convince a jury of her innocence by answering only those questions that are likely to exonerate her.

Sophisticated receivers understand and can correct for senders' strategic use of vagueness. But if these deceptive practices are deployed on naive receivers, then they result in systematic misperceptions. This paper presents a theory-founded experiment to investigate voluntary disclosure to receivers of heterogeneous strategic sophistication under both flexible language, which facilitates vague messages, and precise language. We seek to answer three main questions. How do senders optimally design messages to exploit receivers' naivete? Are (some) receivers systematically fooled by vague disclosure? And can restricting senders' flexibility in disclosure improve information transmission?

We present a model by Hagenbach and Koessler (2017) that builds on Milgrom and Roberts (1986) and Eyster and Rabin (2005) and provides behavioral predictions for our experiment. We then derive novel results on the effects of vagueness on information transmission and welfare. Consider a voluntary disclosure game in which a privately informed sender decides whether and how to disclose verifiable information about her type to a receiver. The sender's payoff is increasing in the receiver's belief about the sender's type, while the receiver's payoff is increasing in the accuracy of her belief. We distinguish between two language regimes: in the precise language regime, if the sender discloses, then the message has to reflect her exact type; in the flexible language regime a sender may send vague messages, i.e. a message that is an interval that contains the sender's true type.[1,2]

If all agents are rational, then the equilibrium features full information revelation in both the precise and the flexible language regimes (Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981).[3] However, the arguments for full information revelation and the irrelevance of language crucially depend on a high degree of strategic sophistication on behalf of the receiver. In reality, many receivers may be naive and struggle to be maximally skeptical in the face of nondisclosure or vague messages. We assume that when a naive receiver encounters nondisclosure, she estimates that the sender is the average type. When she encounters a vague message, she estimates that the sender's type is the average of the sent interval.

The presence of naive receivers in the model drives both nondisclosure (under precise language) and the exploitative deployment of vague messaging (under flexible language). Vague messages take the following simple form. Senders send an interval that spans their

[1]We use the term language to refer to the set of messages at the sender's disposal. The term vague message refers to a message that contains several possible states of the world, but leaves no uncertainty on what these possible states are.

[2]An alternative specification of flexibility might allow senders to disclose any set of types that includes their actual type rather than constraining them to disclose an interval. We favor our choice for two reasons. First, disclosure of a disconnected set of types is rarely observed in the field, presumably because it is less natural and would tip naive receivers off. Modeling this tipping off explicitly would require that receivers' degree of naivete depend on the message, with little guidance from previous work on how this dependence might look. Second, imposing that disclosure has to take the form of an interval in the experiment helps reduce differences in complexity between precise and flexible language and thereby facilitates the interpretation of our treatment comparison

[3]In the precise language regime, the highest type discloses because the disclosed information definitely exceeds receiver expectations. Because nondisclosure now cannot stem from the highest type, the second highest type is compelled to disclose. An iteration of this reasoning yields full disclosure. In the flexible language regime, the receiver's belief that a sender's type is the lower bound of the message sent is self-fulfilling.

actual type and the upper bound of the message space. As noted by Hagenbach and Koessler (2017), moving from the flexible language regime to the precise language regime then implies a trade-off. There is more frequent disclosure in the flexible language regime and more precise disclosure in the precise language regime. Sophisticated receivers, who are not fooled by vagueness, form more accurate beliefs under flexible than under precise language. Naive receivers form more accurate beliefs under precise language. Importantly, we additionally show how information transmission, i.e. the average accuracy of receivers' beliefs, is higher under precise language, irrespective of the proportion of naive receivers.[4]

Our experiment compares a FLEXIBLE and a PRECISE treatment that reflect the distinction between the two language regimes in the model. In both treatments, a sender's type is uniformly distributed over the integers from 0 to 5. A sender in the FLEXIBLE treatment can disclose any interval containing her actual type. For example, a sender with type 2 could disclose that her type belongs to the interval between 2 and 5. A sender in the PRECISE treatment can only disclose her exact type or nothing.

The theoretical predictions are borne out in the experimental data. Many senders are apt spin doctors. In FLEXIBLE, they use vague messages and the exact form of the modal message we observe is remarkably close to the one predicted by the model. In PRECISE, sender behavior reflects a threshold equilibrium in which only high types disclose. Senders disclose more in FLEXIBLE than in PRECISE and only a minority of senders in both treatments does not behave according to the theoretical predictions.

Validating the model's key assumption, we find evidence for the existence of two distinct

[4]Our main rationale for focusing on information transmission and for not also taking into account the sender's payoffs in our discussion of policy is that any surplus the sender obtains relative to the full-rationality benchmark derives entirely from deception as opposed to an underlying fundamental. In common applications of persuasion games, such as sales or financial disclosures by managers, the sender's payoff can be thought of as a price or salary and therefore constitutes a pure transfer. On a practical note, results on the sum of sender and receiver payoffs would be sensitive to the exact scaling of players' payoffs, which would introduce an element of arbitrariness. At the same time, while focusing on information transmission is sensible in our general setting, it is not always clear-cut how information transmission maps into welfare in specific applications. For instance, Ispano and Schwardmann (2021) show that when vertically differentiated firms compete for sophisticated and naive consumers through quality disclosure, inflated beliefs about low-quality products on behalf of naive consumers may improve welfare by exerting competitive pressure on the prices of high-quality products. Moreover, consumers might simply enjoy thinking of the Bordeaux region while drinking a blended wine from Roussillon.

receiver types, i.e. naives and sophisticates. We categorize receivers as either sophisticated or naive on the basis of their guesses and find that the average naive receiver makes smaller mistakes in Precise than in Flexible. Instead, the average sophisticated receiver makes larger or equally large mistakes in Precise.

When we consider all senders, the treatment effect of precise language on information transmission is positive but statistically insignificant because a very small number of observations in Precise feature a sender making the outlier mistake of not disclosing the highest type. If we restrict attention to the majority of senders that are rational or payoff-maximizing, i.e. those senders that probably best capture the behavior of firms, then average information transmission is significantly higher in Precise. The experimental data also match our theoretical predictions on information transmission conditional on the sender's realized type, which should be higher under precise language for intermediate types. Finally, as we provided payoff information to our subjects after each of the experiment's 15 rounds, we find that receivers are able to learn and become more sophisticated over time, with senders adjusting their disclosure strategies accordingly.

The experimental laboratory allows us to strip the decision-making environment of any confounding drivers of vagueness that may be present in the field, to cleanly identify a bias in receivers' inference, and to use experimental variation to test for the policy implications of a model that accounts for this bias in receiver inference. We find that policies that impose precise language on senders are likely to improve information transmission, especially to naive receivers. To appreciate the applications of this result, consider the following real-world settings where precise language is already being imposed.

Germany's main certifier of consumer products, Stiftung Warentest, gives products and services a precise mark and a vague summary category like "very good" and imposes that the latter cannot be disclosed without mention of the former. Similarly, the NGO Consumer Reports in the United States does not allow companies to excerpt content from reports selectively. Several scientific journals require authors to disclose the exact p-value, effect size, degrees of freedom, and statistical test underlying a given result, thereby significantly decreasing the flexibility with which results may be presented. Finally, the self-incrimination clause of the fifth amendment of the United States constitution affords defendants the right not to testify against themselves in criminal cases. However, a majority of US courts take the position that voluntarily waiving the right against self-incrimination opens a defendant up to cross-examination on all issues relevant to the trial.[5] The right not to self-incriminate therefore imposes precise voluntary disclosure.[6]

Our paper contributes to a small literature that studies information disclosure in the experimental laboratory. Our main contribution over previous work in this literature lies in the comparison of different language regimes that have clear empirical counterparts in the policy domain. Our precise language treatment follows Jin et al. (2021), who provide evidence for both incomplete unraveling and receiver naivete. Earlier studies by Forsythe et al. (1989), King and Wallin (1991) and Dickhaut et al. (2003) find evidence for full unraveling after a sufficiently high number of repetitions, albeit in a setting that features several receivers and auctioning mechanisms that potentially permit other explanations for players' behavior (Jin et al., 2021).[7]

Related to our flexible language treatment, Jin et al. (2019) study a mandatory disclosure game in which senders can complexify their disclosure by revealing their type as the sum of a string of numbers. They find that low sender types make use of complexity and that some receivers are fooled by it because they are overconfident in their ability to interpret complex messages.[8] Li and Schipper (2020) study a voluntary disclosure game in

[5]See Yale Law Journal (1952) and Stanford Law Review (1962) for discussions of the waiver and how it has and should be interpreted by the law.

[6]In an influential court ruling, the majority opinion argues against allowing the defendant to "decide how far he will disclose what he has chosen to tell in part [...]" because "it must be conceded that the privilege is to suppress the truth, but that does not mean that it is a privilege to garble it." See the opinion by judge Hand in United States v. St. Pierre, 132 F.2d 837 (2d circuit 1942). Our results highlight a key distinction between suppressing the truth and garbling it by means of partial or vague disclosure and speak to the wisdom in prohibiting the latter.

[7]Also see Benndorf et al. (2015) for an unraveling failure that is driven by senders' bounded rationality and Brown and Fragiadakis (2019) for a receiver misinference that is not based on a lack of strategic sophistication.

[8]Hagenbach and Perez-Richet (2018) conduct an experiment that allows for vague messages. Instead of varying the language at a sender's disposal, they vary the sender's incentive structure. Like us, they find that types who wish to be perceived as another type are more likely to use partial- or nondisclosure. They also find that receivers are better off under acyclical incentive structures, i.e. games in which masquerading

which senders can disclose any set of types that contains their actual type and in which subjects play in the role of both sender and receiver as the experiment progresses. Senders use vagueness in a way that is reminiscent of senders' strategy in our experiment and results are indicative of a high average level of iterative strategic reasoning on behalf of both receivers and senders.

Our results point to a specific form of receiver naivete that is distinct from either overconfidence (Jin et al., 2019) or a failure to engage in higher order strategic reasoning (Li and Schipper, 2020).[9] Our model assumes that a fraction of receivers interpret a given message as stemming from the average sender type who could have sent this message. This bias is akin to fully cursed behavior (Eyster and Rabin, 2005), which in turn is based on receivers not understanding that senders condition their strategy on their private information. The fact that the model's point predictions do well in describing receiver behavior in the two treatments and that qualitative predictions are borne out in treatment comparisons of behavior and aggregate as well as type-specific information transmission lends much credence to the assumed form of naivete.

The exploitation of flexibility in voluntary information disclosure has been documented for car sellers describing their cars on ebay (Lewis, 2011), business schools referring to third-party rankings (Luca and Smith, 2015), and researchers presenting their findings (Krawczyk, 2015; Brodeur et al., 2016). Relatedly, there is evidence that firms shroud (Brown et al., 2010), obfuscate (Ellison and Ellison, 2009; Ferman, 2016), or complexify (Ru and Schoar, 2016) unfavorable information about their products. In markets where voluntary disclosure is necessarily precise, nondisclosure often ensues. For example, producers of salad dressings do not voluntarily disclose fat content if it is high (Mathios, 2000), poor health maintenance organizations do not obtain independent accreditations (Jin, 2005), and movie studios avoid pre-release screenings to critics if a movie's quality is low (Brown et al., 2012, 2013). However, data limitations in the field have thus far kept

incentives are not circular.

[9]Of course, models of iterative strategic reasoning can capture our results if the behavior of level-zero individuals is assumed to follow the behavior of the naive types in our setting.

researchers from studying the causal impact of different language regimes on information transmission and from characterizing the exact nature of receivers' misinference.[10]

In the next section, we present the simple model that guides our experimental design. Section 3 describes experimental design and results and section 4 discusses policy implications.

2 Theoretical predictions

We now turn to presenting the model in Hagenbach and Koessler (2017), adapted to a discrete type space and parametrized to match our experiment.[11] We refer to them for formal proofs of predictions other than our novel result on the effect of the language regime on average information transmission. A sender (S, he) privately observes his type ω, drawn uniformly from $\Omega = \{0, 1, ..., 5\}$, and sends a message $m \in M(\omega)$ to the receiver (R, she), who then makes a guess $g \in [0, 5]$. While S aims at maximizing R's guess, R wants her guess to be as accurate as possible, i.e. minimize $(\omega - g)^2$, which means she finds it optimal to choose as guess her expectation $\mathbb{E}[\omega|m]$. The set of messages $M(\omega)$ available to each type ω depends on the communication regime. Under precise language, $M(\omega) = \{\omega, \Omega\}$, i.e. S can either disclose his type exactly ($m = \omega$) or remain silent ($m = \Omega$). Under flexible language, $M(\omega)$ consists of all closed intervals in Ω containing ω, which hence also includes the option to disclose his type exactly or to remain silent.

If R is rational, then in any (sequential) equilibrium of this game $g = \omega$ under both precise and flexible language since, upon any given m, R correctly infers it sent by the lowest type for which m is available. We suppose there is a probability $\chi \in (0, 1)$ that R is (fully) naive rather than (fully) sophisticated. The posterior distribution of a naive R upon message m coincides with the prior truncated over types for which m is available. Upon

[10] See Dranove and Jin (2010) for a review of the theory and empirics of disclosure in economic applications and Loewenstein et al. (2014) for the psychological subtleties surrounding the analysis of disclosure games.

[11] Our precise and flexible communication regimes correspond respectively to simple and rich language in the terminology of Hagenbach and Koessler (2017).

message $m = \Omega$, her guess hence coincides with the prior mean $5/2$. Upon message $m = \omega$, her guess is ω. And upon message $m = [\underline{\omega}, \bar{\omega}] \subset \Omega$ with $\underline{\omega} < \bar{\omega}$, her guess is $(\underline{\omega} + \bar{\omega})/2$. Parameter χ can be interpreted as the proportion of naive receivers in the population, and it is known by S and a sophisticated R.

When language is flexible, S elects to disclose an interval that spans his type and the highest type (hence only type $\omega = 0$ does not disclose). Indeed, as the equilibrium is necessarily fully separating, this strategy is optimal in that it maximally inflates the guesses of a naive R.[12]

Prediction 1 (Behavior under flexible language)**.** *Under flexible language, S sends $m = [\omega, 5]$, a sophisticated R guesses $g = \omega$ and a naive R guesses $g = (\omega + 5)/2$.*

When language is precise, instead, S elects to disclose only if his type is higher than some cutoff, since his payoffs from disclosing and not disclosing are respectively increasing and constant in ω. Indeed, if S discloses, both a naive and a sophisticate will guess S's type, while if he does not disclose, a naive will guess the average type and a sophisticate will guess her expectation of S's type conditional on nondisclosure. Thus, in a pure strategy equilibrium, the highest non-disclosing type, denoted by ω^*, must solve

$$\omega^* \leq \chi\, 5/2 + (1-\chi)\omega^*/2 \leq \omega^* + 1,$$

which states that type ω^* prefers not to disclose and type $\omega^* + 1$ prefers to do so when the expectation of a sophisticate conditional on nondisclosure ($\omega^*/2$) is correct. Rearranging the inequalities yields $\omega^* = 0$ for $\chi \leq 2/5$, $\omega^* = 1$ for $\chi \in [1/4, 3/4]$ and $\omega^* = 2$ for $\chi \geq 2/3$.[13]

[12] The reason for why full separation necessarily obtains is that, for any candidate equilibrium pooling message, the highest type in the pool always has access to another message that would strictly raise the guess of both a sophisticated and a naive R.

[13] The cutoff ω^* is not always unique since a sophisticated R's belief about ω^* may affect S's incentives to disclose. As described in the appendix (section A.1), in the interior of each multiplicity region there also exists a mixed strategy equilibrium in which the cutoff type randomizes. Since all predictions apply no matter the equilibrium, for ease of exposition we here focus on some pure selection.

Prediction 2 (Behavior under precise language). *Under precise language, there exists a cutoff type $\omega^*(\chi) \in \{0, 1, 2\}$, with $\omega^*(\chi)$ weakly increasing in χ, such that types $\omega > \omega^*$ send $m = \omega$ and types $\omega \leq \omega^*$ send $m = \Omega$. Upon receiving $m = \omega$, both a naive and a sophisticated R guess $g = \omega$, while upon $m = \Omega$ a naive R guesses $g = 5/2$ and a sophisticated R guesses $g = \omega^*/2$.*

Since all types $\omega > 0$ disclose under flexible language while only sufficiently high types disclose under precise language, the following prediction on differences in S's disclosure behavior obtains (all inequalities are strict provided $\omega^* > 0$).

Prediction 3 (Comparison of disclosure behavior). *For any given $\chi \in (0, 1)$*

(a) S discloses more often under flexible than under precise language;

(b) the average disclosing type is higher under precise than under flexible language.

Similarly, the following prediction considers R's expected guess, which coincides with S's expected payoff.

Prediction 4 (Comparison of guesses). *For any given $\chi \in (0, 1)$*

(a) R's expected guess exceeds S's expected type;

(b) R's expected guess increases with χ;

(c) R's expected guess is lower under precise than under flexible language.

Predictions on R's expected guess are driven by the guesses of a naive R, since, given the Bayesian consistency of rational beliefs, the expected guess of a sophisticated R always coincides with the prior mean. For any realization of S's type the guess of a naive R is (weakly, and strictly for some realizations) higher than S's type, which explains prediction 4(a). The average guess of a naive is therefore higher than that of a sophisticate, which entirely drives prediction 4(b) under flexible language. Under precise language, it is also the case that the average guess of a naive increases with χ since, as $\omega^*(\chi)$ increases, the set of S's types that she overestimates increases. Besides, for any realization of S's type, the guess of a naive R is higher under flexible than under precise language (strictly so provided

$0 < \omega < 5$), which explains prediction 4(c). Intuitively, flexible language allows S to more strongly inflate the expectation of a naive R and it also offers more opportunities to do so.

Finally, the following prediction, whose proof is in the appendix (section A.2), considers R's expected utility, i.e. the expected accuracy of her guesses measured by the mean squared error, to which we refer to as information transmission.

Prediction 5 (Comparison of information transmission). *For any $\chi \in (0, 1)$*

(a) Information transmission to a sophisticated R is higher under flexible language and information transmission to a naive R is higher under precise language;

(b) overall information transmission is higher under precise language than under flexible language;

(c) conditional on S's realized type, information transmission is higher under precise language than under flexible language for $\omega = 1$, $\omega = 2$, $\omega = 3$ and $\omega = 4$, while the opposite relation holds for $\omega = 0$ (for $\omega = 5$, $g = 5$ irrespectively).

A naive R makes more accurate guesses under precise language, since it limits the scope for being deceived, while a sophisticated R makes more accurate guesses under flexible language, since it allows her to perfectly infer S's type. Once both effects are considered, the former dominates, so that overall information transmission is higher under precise language.

To gain intuition for this result, first consider the case where there are only a few naive receivers. Then, under precise language, all but the lowest sender type will disclose. As a result, sophisticated receivers are able to always infer the sender's type under either language regime and the comparison of overall information transmission hinges on the inference of the few naive receivers. Since naive receivers will only be fooled by type $\omega = 0$ under precise language, but by types $\omega \in \{0, 1, 2, 3, 4\}$ under flexible language, less information is transmitted in the latter. Now consider the case where almost all receivers are naive. It will again be the case that the comparison of overall information transmission will be driven by naives. And since naives make larger mistakes under flexible language for every type $\omega \in \{1, 2, 3, 4\}$, precise language improves information transmission.

Understanding the effect of language on information transmission for intermediate proportions of naive types requires weighing up the greater mistakes of sophisticates under precise language against the greater mistakes of naives under flexible language. Naturally, sophisticates' mistakes matter more for overall information transmission when there are many of them. But the presence of many sophisticates leads to a lower disclosure threshold and smaller average mistakes by sophisticates (and naives) under precise language. As a result, average mistakes under precise language never rise above average mistakes under flexible language.

In appendix A.3, we further develop these intuitions by breaking down average mistakes into mistakes by sophisticates and naives for each realization of the sender's type. In the process, we prove that the prediction also holds if information transmission is measured in terms of the mean absolute error (i.e. $\mathbb{E}|g - \omega|$), which we use when presenting our experimental data as it is more directly interpretable. Appendix B establishes the robustness of the prediction on further dimensions, namely more general distributions of receiver naivete in the population (section B.1) and non-uniform priors over sender types (section B.2).

3 The Experiment

3.1 Design

The experiment was programmed in zTree (Fischbacher, 2007). A total of 158 subjects participated in 8 sessions at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA) in the spring of 2017, after we piloted our design with 58 subjects in the winter of 2016. One session lasted for about 45 minutes and the average earnings (including a €4 show-up fee) were €15.05, with minimum earnings of €5.90 and maximum earnings of €23.50. The instructions were read aloud by the experimenter. Screenshots of the decision screens are gathered in Appendix E and instructions and payoff tables can be found in Appendix F.

The experiment featured a between subject design that compared the two variants of the disclosure game described in section 2.[14] At the beginning of the experiment, subjects in both treatments were randomly assigned to the role of a sender or the role of a receiver. A subject remained in her assigned role for the duration of the experiment. All subjects played 15 rounds of the disclosure game. In each round, a subject played the game with a randomly selected anonymous partner in the opposite role.

It was common knowledge that a sender's type ω was drawn in each round from the set $\{0, 1, 2, 3, 4, 5\}$ and that each type was equally likely. After privately observing her type, a sender decided on a message to send to the receiver. Our two treatments differed only in the type of messages senders were able to send.

In the FLEXIBLE language treatment (80 subjects), the sender was allowed to send any interval containing her type.[15] In the PRECISE language treatment (78 subjects), the sender could either disclose her precise type or not disclose. In FLEXIBLE, senders were therefore able to send vague messages and while any feasible message in PRECISE was also feasible in FLEXIBLE, the reverse was not true. In the case of nondisclosure, the receiver was notified that "the sender did not send a message" in both treatments. Figure 1 depicts two messages a sender of type 2 might send in the different treatments.

(a) Precise message **(b)** Vague message

Figure 1 Examples of messages

After seeing the sender's message, the receiver stated her guess about the sender's type, i.e., $g \in \{0, 0.5, \ldots, 4.5, 5\}$. The receiver's action space was coarsened so that both sender

[14]As detailed below, the only differences lied in the receiver's action set, which was coarsened, and the exact incentive structure. These differences are theoretically unimportant in that, at least in pure strategy equilibria, the receiver's optimal guess is her expectation of the sender's type and such a guess is available for both naive and sophisticated beliefs.

[15]While sending an interval that contains all possible types was not allowed, the equivalent strategy of nondisclosure was always at a sender's disposal.

and receiver payoffs could be communicated in the form of digestible payoff tables rather than relying on subjects calculating their payoffs by themselves. While the sender was incentivized to induce the highest possible guess from the receiver, the receiver was paid for accuracy. Subjects were paid in probability points and for a single randomly selected round. After each round, subjects received information about the receiver's guess, the sender's type and the probability points they earned.

A receiver's points depended on her guess and the sender's type as follows

$$p_R = \frac{110 - 20|\frac{\omega - g}{1.37}|^{1.4}}{110}.$$

A sender's points depended only on the receiver's guess

$$p_S = \frac{110 - 20|\frac{5 - g}{1.37}|^{1.4}}{110}.$$

The probability points p a subject earned in the payoff-relevant round then determined the likelihood of winning a €8 prize. For example, a subject in the receiver role was paid according to a lottery that yielded a relatively high prize of €8 with probability p_R and a lower prize of €1 with the complementary probability $1 - p_R$. Paying subjects in probability points makes them less liable to the influence of risk preferences (Roth and Malouf, 1979; Hossain and Okui, 2013; Harrison et al., 2014). To make sure that subjects understood the incentive structure we provided them with payoff tables that mapped any constellation of receiver guess and sender type into the relevant probability points and let them solve comprehension tasks before the experiment.

After the main part of the experiment, we elicited subjects' "out-of-sample" beliefs about behavior in the pilot experiment. Senders stated the distribution of receiver guesses upon nondisclosure and receivers stated their belief distribution over non-disclosed sender types.[16] Subjects were paid for being close to a variable's empirical distribution in the pilot

[16] In additional unincentivized elicitations in FLEXIBLE, we asked senders about the average receiver guess in the pilot session after receiving the messages $\{1, 2, 3, 4, 5\}, \{2, 3, 4, 5\}, \{3, 4, 5\}$, and $\{4, 5\}$ and

sessions (see Appendix F.4 for details). These elicitations facilitated a rationality check for senders (see Appendix D) and a consistency check for our naivete classification of receivers. Finally, a very short post-experiment survey collected some additional sociodemographic data.

3.2 Results

We first describe participants' behavior in the two treatments and then analyze information transmission. Our analysis before section 3.2.3 is based on data that pools observations across rounds. For all statistical tests we report p-values from two-sided t-tests that come from a OLS regression with round fixed effects and standard errors clustered at the session level using bootstrapping.

3.2.1 Behavior

Flexible Language Treatment. According to the theory, a sender in the Flexible treatment discloses an interval that spans her type and the upper bound of the type space. Figure 2a depicts the average lower and upper bounds of the messages sent by different sender types. Observed messages are in line with the predictions of the model. Upper bounds are close to the highest type and lower bounds increase with the type. Modal messages, also depicted in the figure, almost perfectly coincide with the theory's predictions. The only exception is provided by senders of type 1, who remain silent more often than they send their predicted message.

As a first step toward analyzing receiver behavior, we normalize guesses. Given a guess g and a message with lower bound $\underline{\omega}$ and upper bound $\bar{\omega} > \underline{\omega}$, the normalized guess is

$$g_n = \frac{g - \underline{\omega}}{|\bar{\omega} - \underline{\omega}|}.$$

The normalization allows for the comparison of guesses induced by different messages.

receivers about the most likely message of all six possible sender types.

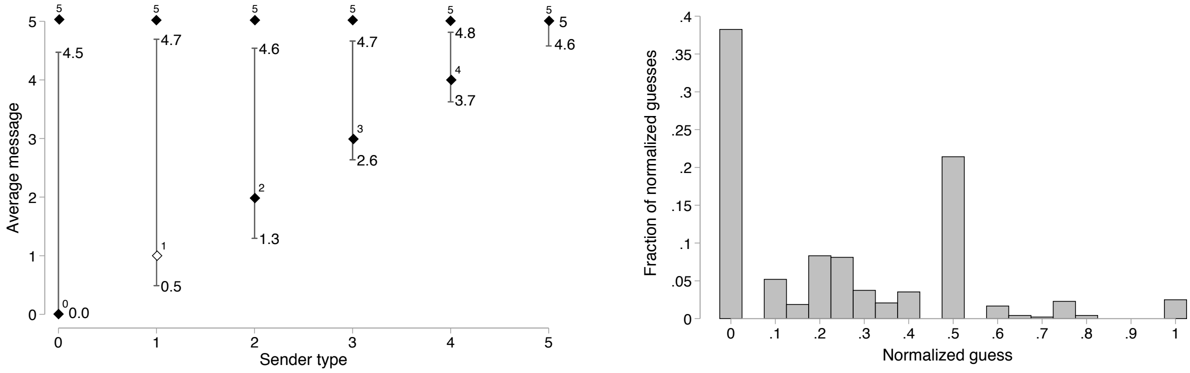

(a) Messages

(b) Normalized guesses

Figure 2 Behavior in Flexible: (a) Solid lines show the avg. lower and upper bounds of all messages sent. Diamonds show the model's predicted lower and upper bounds; diamonds are black when predictions coincide with the modal message in the experiment and hollow if not. Average and modal messages include nondisclosure. (b) Bars show the distribution of normalized guesses.

Normalized guesses range from 0 to 1 and are only defined for non or vague disclosure. A fully naive normalized guess always takes a value of 0.5. The theoretical prediction for a sophisticated normalized guess in Flexible is 0 for all messages.

Figure 2b shows the distribution of normalized guesses. The bimodal distribution with mass points at 0 and 0.5 vindicates our model's assumption that there are two distinct receiver types: sophisticates and fully naive receivers.[17] We find that receivers' average belief upon observing nondisclosure or receiving a vague message is upwardly biased. While the average normalized guess is at about 0.25, senders' average normalized type is significantly lower at 0.13 (p-value < 0.001).[18] Instead, all receivers are able to rationally interpret singleton intervals, i.e., a precisely disclosed type.

Precise Language Treatment. In the presence of naive receivers, our model predicts that precise language will give rise to a threshold equilibrium with nondisclosure by types below the threshold and disclosure by types above. Figure 3a depicts disclosure rates by sender type. In line with an equilibrium threshold of around 2, the disclosure rate is almost zero for the lowest two types, 40 percent for type 2, and above 80 percent for the highest

[17] We can reject the null hypothesis of unimodality using the Dip Test introduced by Hartigan and Hartigan (1985) (p_{Dip}-value < 0.001).

[18] The normalized type is the sender counterpart of the normalized guess and is given by $\omega_n = \frac{\omega - \underline{\omega}}{|\bar{\omega} - \underline{\omega}|}$.

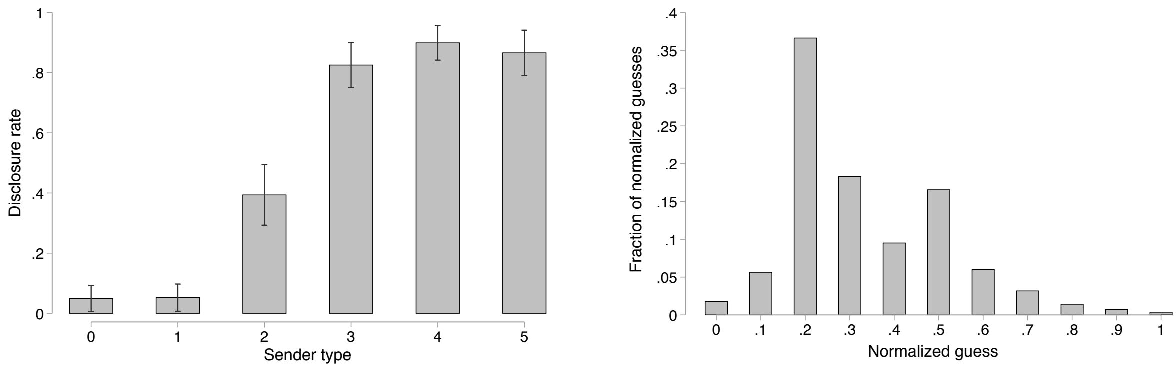

(a) Disclosure rates

(b) Normalized guesses

Figure 3 Behavior in PRECISE: Graph (a) shows 95% confidence intervals around the avg. disclosure rates. Graph (b) shows the distribution of normalized guesses.

three types. Note that disclosure rates of less than 100 percent for the highest types imply a slight departure from our hypothesis of sender rationality.[19]

All receivers are able to rationally interpret a precisely disclosed type. Figure 3b depicts the distribution of normalized receiver guesses upon nondisclosure. We observe a bimodal distribution with mass points around 0.2 and at 0.5.[20] Because of the threshold strategy, the sophisticated guess upon nondisclosure is now larger than zero. In particular, a receiver's empirical best response is equal to the average non-disclosing type, whose normalized value is equal to 0.25. Therefore, the histogram's first mode reflects the accurate beliefs of sophisticated receivers, whereas the second mode corresponds to the beliefs of a fully naive receiver.

Receivers' average normalized guess upon observing nondisclosure is 0.33, which reflects a significant overestimation of the average normalized non-disclosed type of 0.25 (p-value $= 0.007$).[21]

[19]However, after the initial five rounds the disclosure rate of high types increases markedly, e.g., for sender type 5, it increases from 70.8 percent to 92.9 percent.

[20]We can reject the null hypothesis of unimodality using a Dip Test (p_{Dip}-value < 0.001).

[21]Appendix D depicts senders' out-of-sample predictions of the pilot's receivers' guesses upon nondisclosure. Results are reflective of high average sender rationality and an unbiased understanding of receiver behavior. Matching actual receiver behavior in the experiment, senders' predictions feature a modal normalized guess of 0 in FLEXIBLE and 0.2 in PRECISE as well as substantial weight on high, naive guesses. Therefore, the average sender appears to be best-responding to unbiased beliefs about receiver behavior.

Treatment Comparison. Disclosure rates are higher in Flexible, where senders disclose 75 percent of the time, than in Precise, where they disclose 51.5 percent of the time (p-value < 0.001). This result seems to be driven by differences in disclosure strategies, as the average disclosing sender type is significantly higher in Precise than in Flexible (3.59 versus 3.08; p-value < 0.001).

We observe that the average receiver guess is only slightly lower under Precise than under Flexible (2.66 in Precise versus 2.85 in Flexible; p-value $= 0.258$). Therefore, the average sender was not significantly better off in Flexible. However, it is interesting to restrict attention to senders that chose the rational strategy, i.e. $m = [\omega, 5]$ in Flexible and disclosure if and only if $\omega > \omega^*$ in Precise, where the threshold ω^* is the empirical best response to receiver guesses. Looking only at these theory-conforming messages we see that the receivers' average guess is 2.87 under Precise and 3.40 under Flexible (p-value $= 0.001$). We can conclude that, as the model predicts, rational senders are better able to exploit the receivers' bias under Flexible.

3.2.2 Information transmission

We measure information transmission by receivers' *mistakes*, which themselves are given by the absolute difference between a receiver's guess and a sender's type. Perfect information transmission corresponds to a mistake of zero. Figure 4a displays the histograms of receiver mistakes by treatment. We see that, as theory predicts, receivers in Precise make more correct guesses and fewer mistakes of size 0.5 to 1.5, though the effect appears to be moderate. Figure 4b then shows how this pattern is substantially more pronounced for naive receivers. We will discuss our classification of receivers as naive and sophisticated in more detail below.

Overall information transmission. The model predicts that average receiver mistakes are lower in Precise. Table 1 shows the determinants of receivers' mistakes. Column 1 depicts an OLS regression of receiver mistakes on the treatment and tells us that the treatment effect of PRECISE on average mistakes is negative, but given the comparatively low effect size it is rather imprecisely estimated. The related insignificance of the treatment effect is driven by the minority of sender choices that do not conform to our theoretical predictions. To test this ex-post hypothesis, we again restrict the sample to the 828 observations that feature theory-conforming sender behavior.[22] Column 2 focuses on the 70 percent of interactions in which sender behavior does not exhibit mistakes relative to the rational benchmark. These interactions better capture our applications, where senders usually have more experience or more at stake from finding their best strategy. In these cases, restricting senders to the use of PRECISE language leads to lower average receiver mistakes. The significant treatment effect emerges because in focusing on theory-conforming behavior, our data restriction eliminates a very small number of outlier observations driven by sender mistakes that disproportionately occurred in PRECISE. In particular, the treatment effect on average information transmission is also significant at the 10 percent level if we merely drop the 12 observations (1 percent of total observations) that feature a sender of type 5 who does not disclose and thereby generates a disproportionately large outlier receiver mistake. 11 of these observations occurred in PRECISE.

Imposed precision therefore improves information transmission in the absence of pronounced sender irrationality that is unbalanced across language regimes. Note that this result is not obvious ex ante. Even if we restrict senders to be theory-conforming types, the result requires that receivers' bias behaves in a way that does not stray too far from our model's assumptions.

A typology of players. The theory predicts that moving from FLEXIBLE to PRECISE decreases the average mistakes made by naive receivers and increases the average mistakes made by sophisticated receivers. A corollary of this prediction is that the interaction effect of imposed precision and a receiver's naivete on mistakes is negative. In order to test these

[22]In FLEXIBLE, such behavior takes the form of a message that spans the sender's type and 5. In PRECISE, it takes the form of a threshold strategy, whereby only types of 2 or higher disclose. Here, the threshold of 2 is the best response to the distribution of receiver guesses upon nondisclosure.

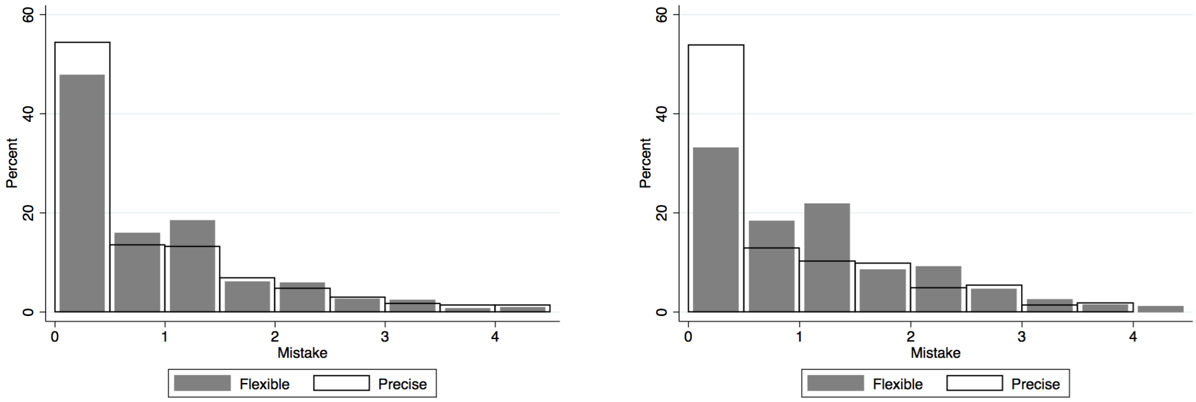

(a) All receivers

(b) Naive receivers

Figure 4 Histograms of receiver mistakes by treatment. Panel (a) contains all receivers. Panel (b) contains only naive receivers, i.e. receivers who made the fully naive guess in more than 85 percent of the rounds that did not feature precise disclosure.

predictions, we use our experimental data to classify receivers as naives and sophisticates.

A normalized guess is fully naive if it is equal to 0.5. We arrive at our measure of individual receiver naivete by dividing the number of rounds in which the receiver stated a fully naive guess by the number of rounds in which the receiver did *not* encounter precise disclosure. If this ratio is smaller than 0.15, we say that a receiver is "hardly ever naive" or sophisticated. Otherwise, a receiver is deemed naive. Applying this classification procedure, we find that 57 percent of the receivers in our sample are sophisticated.[23]

Columns 3 and 4 of Table 1 repeat the regression models of columns 1 and 2, but include only naive receivers. Regardless of whether or not we only include theory-conforming sender behavior, the treatment effect on naive receivers is negative and significant at the 5% level. In columns 5 and 6, we see that, as the model predicts, the treatment effect on sophisticated receiver's mistakes is positive, but these effects are not statistically significant.

[23] The fraction of naive receivers is higher in Precise (61.5 percent) than in Flexible (52.5 percent). However, this difference is not robust to different classification criteria. In general, the proportion of naives is slightly higher in Precise if the classification is based on the frequency of fully naive choices (as in our main classification) and slightly lower if the classification is based on the proximity to best-response behavior, as some classifications in Appendix C. Therefore, there is no reason to suspect that the treatment effect on overall information transmission is driven by differences in the proportion of naives across treatments. To validate our naivete classification, we may ask whether receivers' out-of-sample beliefs, elicited after the experiment, about the type of non-disclosers in another experiment vary systematically according to their classification. Indeed, we find that the average normalized guess of naives is 1.28 whereas the average guess of sophisticates is 0.86 (p-value $= 0.022$).

Dep. variable:	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	-0.0610	-0.159**	-0.252**	-0.231**	0.140	0.0250
	(0.113)	(0.0686)	(0.114)	(0.114)	(0.136)	(0.0824)
Round	-0.0298***	-0.0207***	-0.0331***	-0.0342***	-0.0257***	-0.00932**
	(0.00658)	(0.00493)	(0.0102)	(0.00984)	(0.00710)	(0.00363)
Constant	1.647***	1.742***	2.168***	2.322***	1.259***	1.277***
	(0.102)	(0.0887)	(0.167)	(0.144)	(0.161)	(0.105)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. sender choices	All	Theory-conforming	All	Theory-conforming	All	Theory-conforming
Incl. receivers	All	All	Naives	Naives	Soph.	Soph.
R^2	0.172	0.456	0.331	0.605	0.135	0.454
Observations	1185	828	510	360	675	468

Table 1 OLS regressions of the treatment effect on receivers' absolute mistakes; standard errors clustered at the session level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Appendix C we repeat the regressions in columns 3 through 6 for several alternative classifications of naivete and sophistication, including subjects' high school math grade and various notions of empirical best response. In the majority of specifications, naive receivers make significantly smaller mistakes under precise language and sophisticated receivers make insignificantly larger mistakes under precise language. We also show that the data bears out the corollary of the model's predictions: for all classifications, moving from flexible to precise language leads to relatively smaller mistakes for naive receivers, i.e., the interaction effect between the precise treatment and naivete on mistakes is negative.

Mistakes by sender type. One potential problem with restricting observations to theory-confirming sender behavior is that our exclusion of observations may not be balanced across sender types. Table 2 avoids this problem by looking at receiver mistakes conditional on sender type. According to the model (prediction 5(c)), mistakes in Precise should be lower for sender types $\omega \in \{1, 2, 3, 4\}$, higher for $\omega = 0$ and the same as in Flexible for $\omega = 5$. We see all of these predictions borne out in the data, although

Dep. variable:	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	0.192	-0.455*	-0.525***	-0.367***	-0.119***	0.0328
	(0.174)	(0.256)	(0.103)	(0.0743)	(0.0433)	(0.0376)
Round	-0.0415*	-0.0495***	-0.0109	-0.0202**	-0.000308	-0.00693
	(0.0225)	(0.0101)	(0.0149)	(0.00887)	(0.00313)	(0.00510)
Constant	1.695***	1.548***	0.669***	0.614***	0.199***	0.0850
	(0.149)	(0.182)	(0.106)	(0.119)	(0.0386)	(0.0652)
Sender type	$\omega = 0$	$\omega = 1$	$\omega = 2$	$\omega = 3$	$\omega = 4$	$\omega = 5$
R^2	0.043	0.135	0.201	0.190	0.051	0.014
Observations	155	123	68	159	172	151

Table 2 OLS regressions of the treatment effect on receivers' absolute mistakes considering only interactions with theory-conforming sender choices; standard errors clustered at the session level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the positive effect of Precise on mistakes when the sender's type is 0 is not significant at conventional levels.

3.2.3 Evolution of play and learning

After each round, each subject receives feedback about the sender's type and the receiver's guess in her pair. We may therefore ask how subjects learn from their payoff experience and whether learning differs by treatment. We first look at the senders. Table 3 summarizes how sender behavior evolves over rounds in Flexible and Precise, respectively. We divide the 15 rounds of the experiment into three phases: rounds 1 to 5, rounds 6 to 10 and rounds 11 to 15. In Flexible, the most frequent messages of types 3, 4, and 5 coincide with the theoretical predictions in all phases. In rounds 1-5 types 1 and 2 are more likely to not disclose than using their predicted messages of $\{1, 2, 3, 4, 5\}$ and $\{2, 3, 4, 5\}$ respectively. Over time their behavior gets closer to the theoretical predictions.[24] In Precise, disclosure rates increase for high types and decrease for low types as subjects learn. The sharp increase in the disclosure rate of types 2 is likely to reflect a shift in the

[24]In particular, in rounds 6-10 and 11-15 type 2 most frequently sends the predicted message. Likewise, in rounds 11-15, the predicted message of type 1 is almost as frequent as nondisclosure (14 subjects of type 1 do not disclose, 12 subjects send $\{1, 2, 3, 4, 5\}$ and 2 subjects send $\{1, 2, 3, 4\}$).

Sender type:	$\omega = 0$	$\omega = 1$	$\omega = 2$	$\omega = 3$	$\omega = 4$	$\omega = 5$
Modal sender messages over time in Flexible						
Rounds 1-5	nondisclosure	nondisclosure	nondisclosure	$\{3, 4, 5\}$	$\{4, 5\}$	$\{5\}$
Rounds 6-10	nondisclosure	nondisclosure	$\{2, 3, 4, 5\}$	$\{3, 4, 5\}$	$\{4, 5\}$	$\{5\}$
Rounds 11-15	nondisclosure	nondisclosure	$\{2, 3, 4, 5\}$	$\{3, 4, 5\}$	$\{4, 5\}$	$\{5\}$
Disclosure rates over time in Precise						
Rounds 1-5	6.3%	10%	20.6%	71.1%	83.8%	70.8%
Rounds 6-10	3.5%	3%	38.9%	92.1%	93.1%	93.3%
Rounds 11-15	5%	3%	66.7%	85.2%	93%	92.9%

Table 3 Disclosure in PRECISE and FLEXIBLE by sender type.

disclosure threshold, i.e. a strategic response to the increase in receiver sophistication that we now turn to.

Receivers' average normalized guess decreases over time, which, in line with prediction 4(b), implies that receivers become more sophisticated. In FLEXIBLE, the average normalized guess is 32% in rounds 1-5, 23.1% in rounds 6-10, and 19% in rounds 11-15. In PRECISE, it decreases from 38.2%, to 33.7% in rounds 6-10, and to 27% in rounds 11-15. Table 4 replicates the analysis of Table 1 with the addition of dummies for each phase of the experiment and their interaction with the treatments. Columns 1 and 2 show that there is a negative time trend in receiver mistakes. However, receivers keep significantly overestimating sender types in all phases (two-tailed t-tests, for all phases p-value< 0.001).

It is plausible that naive subjects would learn more quickly under FLEXIBLE because senders deploy deceptive disclosure strategies for a greater amount of realized types, so FLEXIBLE affords more opportunity to learn. As a result, information transmission under FLEXIBLE may not be lower in perpetuity. To test this hypothesis, we can look at column (4) in Table 4. Looking at the interaction between PRECISE and phase, we see that, in line with the hypothesis, experience with the game eradicates mistakes more slowly under PRECISE than under FLEXIBLE, but the effect is not statistically significant. We also acknowledge that the analysis of interactions between time effects and treatment is not

Dep. variable:	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
PRECISE (d)	-0.115	-0.234***	-0.227	-0.334***	-0.0265	-0.0393
	(0.137)	(0.0719)	(0.168)	(0.0799)	(0.169)	(0.0584)
Rounds 6-10 (d)	-0.258*	-0.197***	-0.301	-0.339**	-0.284**	-0.0926*
	(0.141)	(0.0690)	(0.205)	(0.147)	(0.129)	(0.0522)
Rounds 11-15 (d)	-0.349**	-0.277**	-0.328*	-0.447**	-0.381***	-0.0936
	(0.167)	(0.115)	(0.197)	(0.187)	(0.138)	(0.0833)
PRECISE × Rounds 6-10 (d)	0.0686	0.108	0.0248	0.128	0.208	0.138
	(0.142)	(0.0910)	(0.220)	(0.183)	(0.153)	(0.105)
PRECISE × Rounds 11-15 (d)	0.0919	0.0988	-0.103	0.165	0.294	0.0393
	(0.174)	(0.120)	(0.198)	(0.211)	(0.196)	(0.0946)
Constant	1.613***	1.740***	2.118***	2.316***	1.271***	1.271***
	(0.148)	(0.0922)	(0.147)	(0.181)	(0.183)	(0.0919)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. sender choices	All	Theory-conforming	All	Theory-conforming	All	Theory-conforming
Incl. receivers	All	All	Naives	Naives	Soph.	Soph.
R^2	0.173	0.457	0.338	0.608	0.136	0.454
Observations	1185	828	510	360	675	468

Table 4 OLS regressions of the treatment effect on receivers' absolute mistakes over time; robust standard errors clustered at the subject level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

powered well.

4 Discussion

Our model and experimental data suggest that information transmission can be increased by restricting senders' flexibility in disclosing private information to receivers. Moreover, we find that a move to precise voluntary disclosure is likely to disproportionally benefit naive receivers. Since sophisticated receivers are (weakly) harmed by restricting flexibility, while naive receivers benefit, it is tempting to think that the effect of restricting flexibility on average receiver welfare is generally negative when there are many sophisticated receivers. However, this intuition is wrong: restricting flexibility improves information transmission for a broad class of distributions of strategic sophistication. When there are many sophisticates, precise language features (almost) full disclosure and still beats out the flexible language regime.

We have analyzed the disclosure game through the lens of sender rationality. In terms of the applications we have in mind, it is plausible that professional marketers are able to make cunning disclosure decisions and that high-paid attorneys are able to advise their clients on optimal disclosure strategies. And while senders and receivers are often drawn from the same population in the case of research, authors of papers naturally devote substantially more time and cognitive resources to a paper than a paper's readership is able to. Our theoretical results can accommodate and are robust to some sender irrationality. However, as our experiment shows, noisy behavior on behalf of senders can make it difficult to detect the benefits of precise disclosure in the experimental laboratory.

In our simple framework, an easy way to facilitate information transmissions is to legislate the mandatory disclosure of information. Where mandatory disclosure is feasible and unproblematic, our results suggest that it is crucial to also legislate precise language. However, for a number of reasons mandatory precise disclosure may often be infeasible or undesirable where the mere imposition of precision is not. First, mandatory disclosure may be deemed unfair. For example, the mandatory disclosure of a college's rank may be deemed unfair because rankings contain an element of subjectivity and dimensions, like students' entertainment facilities, that a college may reasonably neglect. However, conditional on a college's voluntary disclosure, imposing precision by prohibiting disclosure in selectively broad categories (e.g. "top 30") is likely to be less contentious. Second, it may be prohibitively onerous for a regulator to determine whether a firm chose nondisclosure or simply lacked information. On the other hand, vague disclosure, an act of commission, is easier to determine. Third, mandatory disclosure may yield perverse incentives. For example, in the absence of a defendant's right not to self-incriminate, law enforcement may have an incentive to use coercion or even torture to extract an admission of guilt.[25]

The question of how the presence of naive receivers affects information transmission when senders are not exogenously endowed with private information about their type is

[25] In the case of markets, Matthews and Postlewaite (1985), Polinsky and Shavell (2012) and Ispano (2018) demonstrate that forcing firms to reveal their private information may ultimately hamper information transmission once firms' incentives to acquire information are taken into account.

an interesting avenue for future research. In particular, it is plausible that mandating precise language has a disincentive effect on information acquisition, given that it sets a limit on senders' ability to use information to deceive receivers. This would limit the benefits of imposing precision. At the same time, in other settings, flexible language may be even more harmful than our data suggests. Cain et al. (2005) show that the disclosure of a conflict of interest can lead advisors to give more biased advice by making them feel morally licensed to pursue their private goals. Because flexible language leads to both less information transmission and more disclosure (i.e., moral licensing), it may lead both to a greater underappreciation of an advisor's conflict of interest and to poorer advice.

Our results pertain to information transmission to an average receiver. But the ultimate desirability of precise language may hinge on the weight society attaches to different receiver types. For example, in the case of research, society may deem that information transmission to referees, who are mostly sophisticated, is initially more important than information transmission to the general public, who is more likely to be naive. Yet researchers may write up their findings in an attempt to persuade both of these audiences. It may then be the case that flexible language and its superior information transmission to sophisticated receivers ought to be favored. On a related note, in some settings vague messages may serve a more benevolent purpose than the exploitation of naive receivers. For example, an organization or policy maker may resort to vagueness to communicate uncertainty about the exact information. Then, if precision were imposed, this would lead to overprecise beliefs on behalf of receivers.

References

Benndorf, Volker, Dorothea Kübler, and Hans-Theo Normann, "Privacy concerns, voluntary disclosure of information, and unraveling: An experiment," *European Economic Review*, 2015, *75*, 43–59.

Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg, "Star wars: The empirics strike back," *American Economic Journal: Applied Economics*, 2016, *8* (1), 1–32.

Brown, Alexander and Daniel Fragiadakis, "Benign vs. Self-Serving Information Reduction: Do Individuals Understand the Difference?," *Working Paper*, 2019.

Brown, Alexander L, Colin F Camerer, and Dan Lovallo, "To review or not to review? Limited strategic thinking at the movie box office," *American Economic Journal: Microeconomics*, 2012, *4* (2), 1–26.

_ , _ , and _ , "Estimating structural models of equilibrium and cognitive hierarchy thinking in the field: The case of withheld movie critic reviews," *Management Science*, 2013, *59* (3), 733–747.

Brown, Jennifer, Tanjim Hossain, and John Morgan, "Shrouded attributes and information suppression: Evidence from the field," *Quarterly Journal of Economics*, 2010, *125* (2), 859–876.

Cain, Daylian M, George Loewenstein, and Don A Moore, "The dirt on coming clean: Perverse effects of disclosing conflicts of interest," *The Journal of Legal Studies*, 2005, *34* (1), 1–25.

Dickhaut, John, Margaret Ledyard, Arijit Mukherji, and Haresh Sapra, "Information management and valuation: an experimental investigation," *Games and Economic Behavior*, 2003, *44* (1), 26–53.

Dranove, David and Ginger Zhe Jin, "Quality disclosure and certification: Theory and practice," *Journal of Economic Literature*, 2010, *48* (4), 935–963.

Ellison, Glenn and Sara Fisher Ellison, "Search, obfuscation, and price elasticities on the internet," *Econometrica*, 2009, *77* (2), 427–452.

Eyster, Erik and Matthew Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Ferman, Bruno, "Reading the fine print: information disclosure in the Brazilian credit card market," *Management Science*, 2016, *62* (12), 3534–3548.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Forsythe, Robert, R Mark Isaac, and Thomas R Palfrey, "Theories and tests of "blind bidding" in sealed-bid auctions," *The Rand Journal of Economics*, 1989, *20* (2), 214–238.

Grossman, Sanford J, "The informational role of warranties and private disclosure about product quality," *The Journal of Law and Economics*, 1981, *24* (3), 461–483.

_ **and Oliver D Hart**, "Disclosure laws and takeover bids," *The Journal of Finance*, 1980, *35* (2), 323–334.

Hagenbach, Jeanne and Eduardo Perez-Richet, "Communication with Evidence in the Lab," *Games and Economic Behavior*, 2018, *112*, 139–165.

_ **and Frédéric Koessler**, "Simple versus rich language in disclosure games," *Review of Economic Design*, 2017, *21* (3), 163–175.

Harrison, Glenn W, Jimmy Martínez-Correa, and J Todd Swarthout, "Eliciting subjective probabilities with binary lotteries," *Journal of Economic Behavior & Organization*, 2014, *101*, 128–140.

Hartigan, John A and PM Hartigan, "The dip test of unimodality," *The Annals of Statistics*, 1985, *13* (1), 70–84.

Hossain, Tanjim and Ryo Okui, "The binarized scoring rule," *The Review of Economic Studies*, 2013, *80* (3), 984–1001.

Ispano, Alessandro, "Information acquisition and the value of bad news," *Games and Economic Behavior*, 2018, *110*, 165–173.

_ **and Peter Schwardmann**, "Cursed Consumers and the Effectiveness of Consumer Protection Policies," *Working paper*, 2021.

Jin, Ginger Zhe, "Competition and disclosure incentives: an empirical study of HMOs," *RAND Journal of Economics*, 2005, pp. 93–112.

_ **, Michael Luca, and Daniel J Martin**, "Complex Disclosure," Working Paper 2019.

_ **,** _ **, and Daniel Martin**, "Is no news (perceived as) bad news? An experimental investigation of information disclosure," *American Economic Journal: Microeconomics*, 2021, *13* (2), 141–73.

King, Ronald R and David E Wallin, "Voluntary disclosures when seller's level of information is unknown," *Journal of Accounting Research*, 1991, *29* (1), 96–108.

Krawczyk, Michał, "The search for significance: a few peculiarities in the distribution of P values in experimental psychology literature," *PloS one*, 2015, *10* (6).

Lewis, Gregory, "Asymmetric information, adverse selection and online disclosure: The case of eBay motors," *The American Economic Review*, 2011, *101* (4), 1535–1546.

Li, Ying Xue and Burkhard C Schipper, "Strategic reasoning in persuasion games: An Experiment," *Games and Economic Behavior*, 2020, *121*, 329–367.

Loewenstein, George, Cass R. Sunstein, and Russell Golman, "Disclosure: Psychology Changes Everything," *Annual Review of Economics*, 2014, *6* (1), 391–419.

Luca, Michael and Jonathan Smith, "Strategic disclosure: The case of business school rankings," *Journal of Economic Behavior & Organization*, 2015, *112*, 17–25.

Mathios, Alan D., "The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market," *Journal of Law and Economics*, 2000, *43* (2), pp. 651–678.

Matthews, Steven and Andrew Postlewaite, "Quality Testing and Disclosure," *RAND Journal of Economics*, 1985, *16* (3), 328–340.

Milgrom, Paul R, "Good news and bad news: Representation theorems and applications," *The Bell Journal of Economics*, 1981, *12(2)*, 380–391.

_ **and John Roberts**, "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 1986, *17* (1), 18–32.

Polinsky, A. Mitchell and Steven Shavell, "Mandatory Versus Voluntary Disclosure of Product Risks," *Journal of Law, Economics, and Organization*, 2012, *28* (2), 360–379.

Roth, Alvin E and Michael W Malouf, "Game-theoretic models and the role of information in bargaining.," *Psychological Review*, 1979, *86* (6), 574.

Ru, Hong and Antoinette Schoar, "Do credit card companies screen for behavioral biases?," Working Paper, National Bureau of Economic Research 2016.

Stanford Law Review, "Waiver of the Privilege against Self Incrimination," 1962, *14* (4), 811–826.

Yale Law Journal, "The Privilege against Self-Incrimination: The Doctrine of Waiver," 1952, *61* (1), 105–110.

SUPPLEMENTARY MATERIALS

A Additional material on theoretical predictions

A.1 Mixed-strategy equilibria under precise language

Since the payoff from disclosing is strictly increasing in ω at most one type can randomize in equilibrium. Also, this must be either type 1 or 2, since disclosing is clearly suboptimal for type 0 and not disclosing is clearly suboptimal for types $\omega \geq 3$. In a candidate equilibrium in which type 1 discloses with probability ϵ, his indifference condition dictates

$$1 = \chi \frac{5}{2} + (1 - \chi) \frac{(1 - \epsilon)}{1 + (1 - \epsilon)}.$$

Solving for ϵ yields $\epsilon^* = (8\chi - 2)/3\chi$. This equilibrium hence exists for $\chi \in (1/4, 2/5)$, since only then $\epsilon^* \in (0, 1)$. Likewise, in a candidate equilibrium in which type 2 discloses with probability ϵ, his indifference condition dictates

$$2 = \chi \frac{5}{2} + (1 - \chi) \frac{(1 - \epsilon)2 + 1}{2 + (1 - \epsilon)}.$$

Solving ϵ yields $\epsilon^* = (9\chi - 6)/\chi$, and hence this equilibrium exists for $\chi \in (2/3, 3/4)$.

A.2 Proof of predictions on average information transmission

S's expected loss under flexible language is

$$EL^{\text{flex}} \equiv \sum_{\omega=0}^{5} \frac{1}{6} \chi \left(\omega - \frac{\omega + 5}{2} \right)^2 = \frac{55\chi}{24}.$$

Clearly, S's expected loss under precise language is the highest in the (pure) equilibrium in which the non-disclosing type is as high as possible. Selecting that equilibrium, i.e. $\omega^* = 0$

for $\chi < 1/4$, $\omega^* = 1$ for $\chi \in [1/4, 2/3)$ and $\omega^* = 2$ for $\chi \geq 2/3$, yields

$$EL^{\text{prec}} \equiv \sum_{\omega=0}^{\omega^*} \frac{1}{6} \left((1-\chi) \left(\omega - \frac{\omega^*}{2} \right)^2 + \chi \left(\omega - \frac{5}{2} \right)^2 \right) = \begin{cases} \frac{25\chi}{24} & 0 \leq \chi < \frac{1}{4} \\ \frac{4\chi}{3} + \frac{1}{12} & \frac{1}{4} \leq \chi < \frac{2}{3} \\ \frac{9\chi}{8} + \frac{1}{3} & \frac{2}{3} \leq \chi \leq 1 \end{cases}.$$

By a pairwise comparison of each case, $EL^{\text{flex}} > EL^{\text{prec}}$, since the opposite inequality would require $\chi \leq 2/7$ when $\omega^* = 2$ and $\chi \leq 2/(23)$ when $\omega^* = 1$ (the result is already apparent when $\omega^* = 0$).

Once S's type has realized, it is clear that R's expected loss is always zero for type $\omega = 5$, strictly higher under flexible language than under precise language for all types $\omega > \omega^*$, and higher under precise language for type $\omega = 0$ (strictly so if $\omega^* > 0$). Also, letting the subscript denote the realized type throughout, $EL^{\text{flex}}|_{\omega=2} = \chi\frac{9}{4}$ and $EL^{\text{flex}}|_{\omega=1} = \chi 4$. Suppose first that under precise language $\omega^* = 2$. Then, $EL^{\text{prec}}|_{\omega=2} = (1-\chi) + \chi\frac{1}{4}$ and $EL^{\text{prec}}|_{\omega=1} = \chi\frac{9}{4}$. $EL^{\text{prec}}|_{\omega=2} > EL^{\text{flex}}|_{\omega=2}$ would require $\chi < 1/3$, but under such value $\omega^* < 2$. Likewise, it is clear that $EL^{\text{prec}}|_{\omega=1} < EL^{\text{flex}}|_{\omega=1}$ no matter the χ. Suppose instead that under precise language $\omega^* = 1$. Then, $EL^{\text{prec}}|_{\omega=1} = \chi\frac{9}{4} + (1-\chi)\frac{1}{4}$. $EL^{\text{prec}}|_{\omega=1} > EL^{\text{flex}}|_{\omega=1}$ would require $\chi < 1/8$, but under such value $\omega^* = 0$.

A.3 Mistakes by type

Let us define R's mistake as $|g - \omega|$, i.e. the distance between S's type and R's guess. Table 5 and 6 report the equilibrium mistakes respectively of a naive and sophisticated R for each realization of the sender's type, where under precise language we distinguished the three possibilities for S's equilibrium disclosure cutoff ω^*.

In each of the four cases, expected total absolute mistakes obtain by summing over realizations of R's status as sophisticate or naive, weighted respectively by $1-\chi$ and χ, and of S's type ω (the fact that each realization of ω has probability 1/6 is just a normalizing factor, which can hence be ignored). Let us denote expected total absolute

		$\omega=0$	$\omega=1$	$\omega=2$	$\omega=3$	$\omega=4$	$\omega=5$
Flexible		5/2	2	3/2	1	1/2	0
	$\omega^*=0$	5/2	0	0	0	0	0
Precise	$\omega^*=1$	5/2	3/2	0	0	0	0
	$\omega^*=2$	5/2	3/2	1/2	0	0	0

Table 5 Equilibrium mistakes of a naive receiver by sender's type

		$\omega=0$	$\omega=1$	$\omega=2$	$\omega=3$	$\omega=4$	$\omega=5$
Flexible		0	0	0	0	0	0
	$\omega^*=0$	0	0	0	0	0	0
Precise	$\omega^*=1$	1/2	1/2	0	0	0	0
	$\omega^*=2$	1	0	1	0	0	0

Table 6 Equilibrium mistakes of a sophisticated receiver by sender's type

mistakes under flexible language and under precise language with cutoff ω^* respectively AM^{flex} and $AM^{\text{prec}}_{\omega*}$. Then, AM^{flex} and $AM^{\text{prec}}_{\omega^*=2}$ can be decomposed as follows

$$\begin{aligned} AM^{\text{flex}} &= \underbrace{\chi\left(\frac{5}{2}+\frac{3}{2}+\frac{1}{2}\right)}_{\text{mistakes of naives}} + \underbrace{\chi(2+1)}_{\text{mistakes of naives}} \\ AM^{\text{prec}}_{\omega^*=2} &= \underbrace{\chi\left(\frac{5}{2}+\frac{3}{2}+\frac{1}{2}\right)}_{\text{mistakes of naives}} + \underbrace{(1-\chi)2}_{\text{mistakes of sophisticates}} \quad , \end{aligned}$$

i.e. in a common term of naive errors, and a term of errors respectively of naives and sophisticates only. $AM^{\text{prec}}_{\omega^*=2} > AM^{\text{flex}}$ would hence require $\chi < 2/5$. But as shown above, under such a low χ, $\omega^* < 2$ necessarily. Likewise, AM^{flex} and $AM^{\text{prec}}_{\omega^*=1}$ can be decomposed as

$$\begin{aligned} AM^{\text{flex}} &= \underbrace{\chi\left(\frac{5}{2}+\frac{3}{2}\right)}_{\text{mistakes of naives}} + \underbrace{\chi(2+1+\frac{1}{2})}_{\text{mistakes of naives}} \\ AM^{\text{prec}}_{\omega^*=1} &= \underbrace{\chi\left(\frac{5}{2}+\frac{3}{2}\right)}_{\text{mistakes of naives}} + \underbrace{(1-\chi)(\frac{1}{2}+\frac{1}{2})}_{\text{mistakes of sophisticates}} \quad . \end{aligned}$$

$AM^{\text{prec}}_{\omega^*=1} > AM^{\text{flex}}$ would hence require $\chi < 2/9$. But again, under such a low χ, $\omega^* < 1$.

Finally, it is apparent that $AM^{\text{flex}} > AM^{\text{prec}}_{\omega^*=0}$. The qualitative robustness of prediction 5(c) to the use of absolute mistakes can be checked by analogous arguments.[26]

B Robustness of predictions on average information transmission

Throughout, following Hagenbach and Koessler (2017), we suppose for convenience and tractability that S's type is drawn from a continuous distribution $f(\omega)$ and normalize its (full) support to $\Omega = [0,1]$. Let $F(\omega)$ denote its cumulative distribution and μ its prior mean. Under flexible language, the equilibrium is unaffected except that the guess of a naive R upon message $m = [\underline{\omega}, \bar{\omega}]$ with $\underline{\omega} < \bar{\omega}$ is now

$$g = \mathbb{E}\big[\omega \mid \omega \in [\underline{\omega}, \bar{\omega}]\big] = \frac{\int_{\underline{\omega}}^{\bar{\omega}} \omega f(\omega)\,\mathrm{d}\omega}{F(\bar{\omega}) - F(\underline{\omega})},$$

which, when $f(\omega)$ is uniform, is again $g = (\underline{\omega} + \bar{\omega})/2$.

Under precise language, the equilibrium is again characterized by a disclosure cutoff $\omega^* \in (0,1)$ such that types $\omega > \omega^*$ disclose and types $\omega < \omega^*$ stay silent (contrary to the discrete model, the disclosure decision of type ω^* can be arbitrary). For a given ω^*, the guess of a sophisticated and a naive R upon nondisclosure are now respectively $g = \mathbb{E}[\omega \mid \omega \leq \omega^*] = \frac{\int_0^{\omega^*} \omega f(\omega)\mathrm{d}\omega}{F(\omega^*)}$ and $g = \mu$, which when $f(\omega)$ is uniform are respectively $g = \omega^*/2$ and $g = 1/2$. Thus, the disclosure cutoff now solves

$$\omega^* = \chi\mu + (1-\chi)\frac{\int_0^{\omega^*} \omega f(\omega)\,\mathrm{d}\omega}{F(\omega^*)}, \tag{B.1}$$

which when $f(\omega)$ is uniform yields the unique solution $\omega^* = \frac{\chi}{1+\chi}$, strictly increasing in χ.

[26] The sole difference is that absolute mistakes are higher under precise language also for type $\omega = 1$ when $\chi \in (1/2, 2/3)$ and $\omega^* = 1$.

The expected losses of R under flexible and precise language are respectively

$$EL^{\text{flex}} = \chi \int_0^1 \left(\frac{\int_\omega^1 t\, f(t)\, \mathrm{d}t}{1 - F(\omega)} - \omega \right)^2 f(\omega) \mathrm{d}\omega$$

$$EL^{\text{prec}} = \chi \int_0^{\omega^*} \left(\int_0^1 t f(t)\, \mathrm{d}t - \omega \right)^2 f(\omega)\, \mathrm{d}\omega + (1 - \chi) \int_0^{\omega^*} \left(\frac{\int_0^{\omega^*} t f(t)\, \mathrm{d}t}{F(\omega^*)} - \omega \right)^2 f(\omega)\, \mathrm{d}\omega.$$

When $f(\omega)$ is uniform these become respectively $EL^{\text{flex}} = \frac{\chi}{12}$ and $EL^{\text{prec}} = \frac{\chi^2(3+\chi)}{12(1+\chi)^3}$. Note that prediction 5(b) hence also holds for the continuous uniform model, since then $EL^{\text{prec}} = cEL^{\text{flex}}$ with $c = \frac{\chi(3+\chi)}{(1+\chi)^3} < 1$.

B.1 General Distribution of Naivete

In this section, we fix $f(\omega)$ to be the uniform distribution. Instead, R's naivete χ is now drawn form a continuous distribution $h(\chi)$ with full support on $[0, 1]$, mean λ and variance σ^2. The belief of a χ-naive R upon any given message is a mixture of the posterior of a fully sophisticated receiver (with weight $1 - \chi$) and a fully naive receiver (with weight χ). Our main binary model obtains as a limit and special case when $h(\chi)$ puts weight only on 0 and 1. Likewise, the model in Eyster and Rabin (2005) corresponds to a degenerate $h(\chi)$ that puts all weight on a single value of χ.

S's behavior under flexible language is unaffected, so that the guess of a χ-naive R upon message $m = [\underline{\omega}, \bar{\omega}]$ with $\underline{\omega} \leq \bar{\omega}$ is $g = \chi(\underline{\omega} + \bar{\omega})/2 + (1 - \chi)\underline{\omega}$. As for S's behavior under precise language, S's disclosure cutoff now must solve

$$\omega^* = \int_0^1 \left(\chi \frac{1}{2} + (1 - \chi) \frac{\omega^*}{2} \right) h(\chi) \mathrm{d}\chi,$$

since the guess of χ-naive R upon nondisclosure is $g = \chi \frac{1}{2} + (1 - \chi) \frac{\omega^*}{2}$. The unique solution is $\omega^* = \frac{\lambda}{1+\lambda}$.

We formally establish (the generalization of) prediction 5(a) and prediction 5(b).

Prediction B.1. *For any distribution of naivete in the population* *(a) information transmission to a χ-naive R is higher under precise language than under flexible language if and only if χ is above some cutoff $\chi^* \in (0,1)$;*

(b) overall, information transmission is higher under precise than under flexible language.

The expected loss of R under flexible and precise language are now

$$EL^{\text{flex}} = \int_0^1 \int_0^1 \left(\chi\frac{\omega+1}{2} + (1-\chi)\omega - \omega\right)^2 \mathrm{d}\omega\, h(\chi)\mathrm{d}\chi = \int_0^1 \underbrace{\frac{\chi^2}{12}}_{EL^{\text{flex}}_\chi(\chi)} h(\chi)\mathrm{d}\chi$$

$$EL^{\text{prec}} = \int_0^1 \int_0^{\omega^*} \left(\chi\frac{1}{2} + (1-\chi)\frac{\omega^*}{2} - \omega\right)^2 \mathrm{d}\omega\, h(\chi)\mathrm{d}\chi = \int_0^1 \underbrace{\frac{1}{12}\omega^*\left((\omega^*)^2 + 3(1-\omega^*)^2\chi^2\right)}_{EL^{\text{prec}}_\chi(\chi)} h(\chi)\mathrm{d}\chi,$$

where EL^{flex}_χ and EL^{prec}_χ denote the expected loss of a χ-naive R in the respective language regime. Prediction B.1(a) follows from

$$\begin{aligned} EL^{\text{flex}}_\chi(0) &= 0 < \frac{(\omega^*)^3}{12} = EL^{\text{prec}}_\chi(0), \\ EL^{\text{flex}}_\chi(1) &= \frac{1}{12} > \frac{1}{2}\omega^*(3 - 6\omega^* + 4(\omega^*)^2) = EL^{\text{prec}}_\chi(1), \\ \frac{dEa^{\text{flex}}}{d\chi} &= \frac{\chi}{6} > \frac{1}{2}\omega^*(1-\omega^*)^2\chi = \frac{dEa^{\text{prec}}}{d\chi}. \end{aligned}$$

As for prediction B.1(b), we may write

$$\begin{aligned} EL^{\text{prec}} - EL^{\text{flex}} &= \frac{1}{12}\int_0^1 \left((\omega^*)^3 - (1 - 3(1-\omega^*)^2\omega^*)\chi^2\right) h(\chi)\mathrm{d}\chi \\ &\propto \frac{\lambda^3 - (1+\lambda^2(3+\lambda))\mathbb{E}\left[\chi^2\right]}{(1+\lambda)^3}. \end{aligned}$$

Thus, $EL^{\text{prec}} \geq EL^{\text{flex}}$ if and only if

$$\mathbb{E}\left[\chi^2\right] \leq \frac{\lambda^3}{1+3\lambda^2+\lambda^3}.$$

Using $\mathbb{E}\left[\chi^2\right] \equiv \lambda^2 + \sigma^2$, one sees that this is impossible as $\sigma^2 > 0$ and $\lambda^2 > \frac{\lambda^3}{1+3\lambda^2+\lambda^3}$.

B.2 General Distribution of Types

We now revert to a binary distribution of sophistication as in our main model, and investigate prediction 5(b) for non-uniform prior distributions over the sender's type by simulation. In detail, we use as family of priors the beta distribution, which is defined on $[0, 1]$ and can take a wide range of shapes (u-shaped, hill-shaped, increasing, decreasing) depending on its parameters $\alpha > 0$ and $\beta > 0$.[27] We numerically solve the model for different values of α and β, each ranging from 1/10 to 10, and of χ, ranging from 1/20 to 19/20, and check whether $EL^{\text{flex}} - EL^{\text{prec}} > 0$.[28]

The inequality is verified for 2226 out of 2250 parameter combinations. The 24 counterexamples occur when $\alpha \geq 3$ and $\beta = 1/10$, i.e., when the prior mean ($\frac{\alpha}{\alpha+\beta}$) is very large and the probability mass concentrated around 1.[29] The sender's ability to make upwardly vague claims under flexible language is then somehow limited by construction, while the set of types who disclose under precise language can be very small. This explains why information transmission may eventually be higher under flexible language. This occurs for intermediate levels of naivete in the population (in all counterexamples $\chi \in [13/20, 17/20]$), so that the disclosure cutoff under precise language remains large while at the same time information transmission to sophisticates has non-negligible weight in determining overall information transmission. Notice, however, that in all counterexamples the percentage reduction in information transmission that imposing precise language entails is small, i.e., $\frac{EL^{\text{flex}} - EL^{\text{prec}}}{EL^{flex}} < -4\%$, while in regular instances the correspondent percentage gain is typically larger (larger than 20% in 95% of the regular instances, and as high as 99%).

[27]The density of a beta distribution with shape parameters $\alpha > 0$ and $\beta > 0$ is

$$f(\omega) = \frac{\omega^{\alpha-1}(1-\omega)^{\beta-1}}{\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\mathrm{d}t}.$$

[28]When $\alpha < 1$ or $\beta < 1$, equation (B.1) can in principle have multiple solutions. Since EL^{prec} is increasing in the disclosure cutoff, we programmed both a more stringent test which uses the largest solution and a weaker test which uses the smallest one. This precaution proved unnecessary as in all instances ω^* turned out to be unique.

[29]When $\alpha > 1$ and $\beta < 1$, the density of the beta distribution is hyperbolically increasing with a vertical asymptote at 1 and, as α/β increases, the distribution gets steeper at high values of ω and flatter elsewhere.

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	-0.186	0.0294	-0.0856	0.188**	-0.343***	0.0265
	(0.151)	(0.125)	(0.122)	(0.0914)	(0.116)	(0.146)
Round	-0.0202	-0.0319***	-0.0333***	-0.0207***	-0.0491***	-0.0250***
	(0.0177)	(0.00653)	(0.00937)	(0.00567)	(0.0114)	(0.00838)
Constant	2.189***	1.445***	1.977***	1.008***	2.411***	1.432***
	(0.261)	(0.130)	(0.118)	(0.0686)	(0.107)	(0.135)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. receivers	Naives	Soph.	Naives	Soph.	Naives	Soph.
Criterion	**Rarely naive**		**Never naive**		**Math grade**	
R^2	0.361	0.138	0.265	0.116	0.375	0.130
Observations	300	885	750	435	330	855

Table 7 OLS regressions of the treatment effect on receivers' absolute mistakes; standard errors clustered at the session level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Alternative classification of naives and sophisticates

This appendix investigates the robustness of results in columns 3 through 6 of Table 1 by repeating the analysis using different classifications of sophistication and naivete. In columns 1 and 2 of Table 7, receivers are classified as sophisticated if they are "rarely naive", i.e., if they make the fully naive choice in less than 30 percent of the rounds in which they face either vague disclosure or nondisclosure. In columns 3 and 4 of Table 7, receivers are classified as sophisticated if they are "never naive", i.e., if they never make a fully naive choice in the rounds in which they face either vague disclosure or nondisclosure. In columns 5 and 6 of Table 7, we use a measure that is exogenous to receiver's choices in the experiment for the classification: we classify receivers with a high school math grade (Abitur) of 1 or 2 as sophisticated and receivers with a math grade of 3, 4, 5 or 6 as naives. This classification is equivalent to a median split.

When we use the "rarely naive" criterion, we find that naives make insignificantly smaller mistakes under precise language and that sophisticates make insignificantly larger mistakes under precise language. When we use the "never naive" criterion (columns 3 and 4), naives make insignificantly smaller mistakes under precise language, while sophisticates make significantly larger mistakes. The "never naive" criterion results in a more selective pool of sophisticates who are hurt by moving from flexible to precise language. When we classify receivers based on their high school math grade (columns 5 and 6), we find that naives do significantly worse and that sophisticates do insignificantly better under precise language.

The above criteria, except for the math grade, are based on the incidence of naive choices and therefore pool all other choices under the label of sophisticated behavior. Alternatively, we may call a receiver sophisticated if her choices line up well with empirical best response behavior. Table 8 uses three notions of empirical best response behavior to classify receivers. Consider the criterion "best response 1". As in section 3.2.3, we divide our experiment into phase 1 (rounds 1 to 5), phase 2 (rounds 6 to 10) and phase 3 (rounds 11 to 15). For each phase and each possible message, including nondisclosure, we calculate the average type that actually sent this message. The use of phases allows us to arrive at a more precise measure of average behavior. We call a receiver's guess a noisy empirical best response if it lies less than 0.5 above and less than 0.5 below the average sender type conditional on a given message. The criterion "best response 1" then classifies a receiver as sophisticated if her guess is a noisy best response in more than 75 percent of rounds that featured either vague disclosure or nondisclosure. The criterion "best response 2" is laxer and classifies an individual as sophisticated if her guess is a noisy best response in more than 50 percent of rounds that featured either vague disclosure or nondisclosure. The criterion "best response 3" is defined like "best response 1" except that is allows for a 1-unit deviation from the true average type in defining the empirical best response. In Table 8, columns 1 and 2 feature the criterion "best response 1" and columns 3 and 4 the same criterion, but only theory-conforming sender behavior. Columns 5 and 6 feature best response 2, whereas columns 7 and 8 feature best response 3. In all cases, naives are found to make significantly smaller mistakes under precise language, while there is no treatment effect on sophisticates.

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake	(7) Mistake	(8) Mistake
Precise (d)	-0.257**	0.125	-0.359***	0.0297	-0.311**	-0.0775	-0.257***	0.00789
	(0.108)	(0.129)	(0.0908)	(0.0443)	(0.122)	(0.128)	(0.0802)	(0.0333)
Round	-0.0280***	-0.0339***	-0.0198***	-0.0206***	-0.0178***	-0.0361***	-0.0219***	-0.0185***
	(0.00661)	(0.0116)	(0.00555)	(0.00633)	(0.0119)	(0.00840)	(0.00694)	(0.00464)
Constant	2.050***	1.220***	2.179***	1.219***	2.240***	1.458***	2.286***	1.238***
	(0.108)	(0.150)	(0.118)	(0.105)	(0.147)	(0.119)	(0.0857)	(0.0896)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Incl. receivers	Naives	Soph.	Naives	Soph.	Naives	Soph.	Naives	Soph.
Criterion	**Best response 1**		**Best response 1**		**Best response 2**		**Best response 3**	
Sender choices	All		Theory-conforming		All		All	
R^2	0.264	0.134	0.532	0.569	0.359	0.124	0.550	0.592
Observations	720	465	519	309	420	765	410	418

Table 8 OLS regressions of the treatment effect on receivers' absolute mistakes; standard errors clustered at the session level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Sender's out-of-sample beliefs

Figure 5 reports senders' out-of-sample predictions of receivers' distribution of guesses conditional on observing non-disclosure. Senders made these predictions on the behavior of receivers in a pilot experiment.

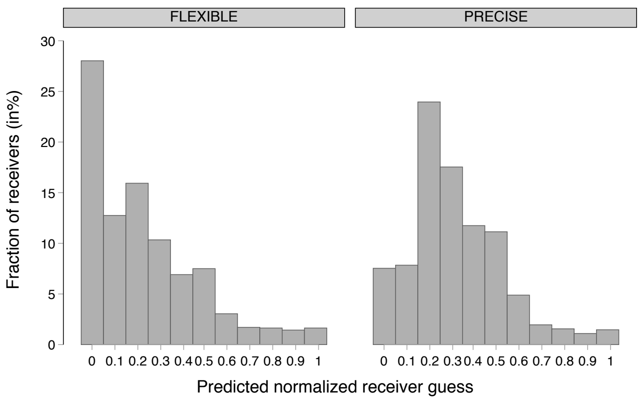

Figure 5 Senders' out-of-sample predictions of receivers' distribution of guesses conditional on observing non-disclosure (by treatment). Senders made predictions about receiver behavior in a pilot experiment.

E Decision Screens

Figure 6 shows the decision screen in Flexible. The sender could freely specify the interval to send by clicking on and herewith selecting the respective types to be included. A preview window showed how the message would appear on the receiver's screen.

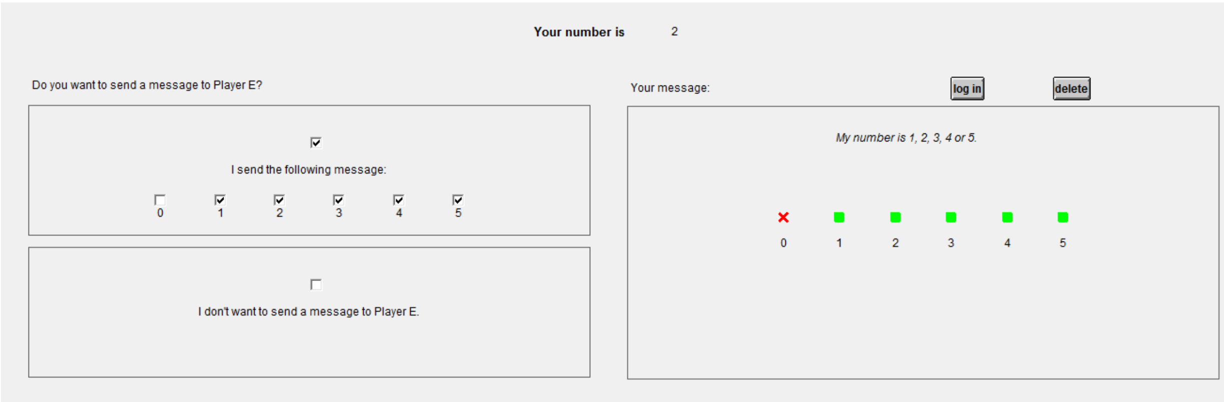

Figure 6 Senders' decision screen in Flexible

Figure 7 shows an example if a sender decision screen in Precise. Here, the senders were provided with the two options in random order.

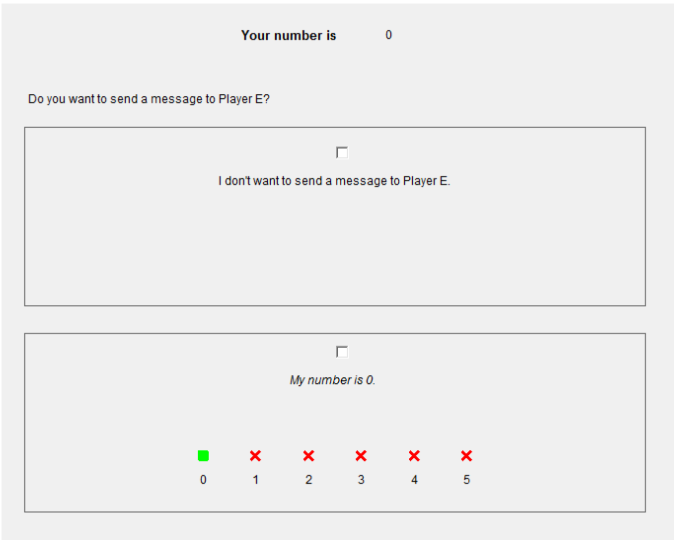

Figure 7 Senders' decision screen in Precise

F Instructions

F.1 Flexible treatment

This experiment is composed of 15 rounds. At the beginning of this experiment, it will be determined randomly whether you are **player S** or **player E**. You will keep this role in all 15 rounds. In each round you play a game with a randomly chosen participant in the opposite role. It is very unlikely, that you are paired up with the same participant in two consecutive rounds.

The Game

In each round, player S receives a number in the **range 0, 1, 2, 3, 4, 5** via the computer. All the numbers are equally likely. Player E does not see which number player S receives. However, player S can send a message regarding his or her number to player E. Player E must guess the number of player S. At the end of each round both players are informed about the number of player S and the guess of player E.

Decision of player S

After receiving the number, player S can decide about whether or not he or she would like to send a message to the recipient. Player S can decide which message he or she would like to send. In doing so, three rules must be complied with:

1. **The sent message must contain the true number of the sender**

   Example: If the sender receives number 3, he can only send messages that contain the number 3.

2. **The sent message must not contain gaps.**

   Example: The sender with number 3 must not send the numbers 2, 3, 5 as possible numbers because the 4 is missing in this row.

3. **The send message may contain maximum five numbers.**

Probability	Payoff
PP%	8 Euro
(100-PP)%	1 Euro

Example: The sender with the number 3 may only send 5 of his possible numbers in total. The sender may not send all six numbers (i.e., 0, 1, 2, 3, 4, 5).

When player S has received, for example, the number 3, he or she can send a message that contains the true number and no gaps or send no message at all. This, for example, applies to the message "My number is 3.". Graphically, the message "My number is 3" will be depicted by a green box above number 3 and red crosses above 0, 1, 2, 4 and 5:

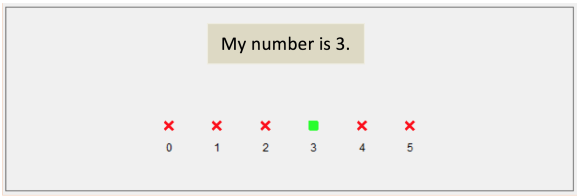

Decision of player E

Player E either sees the message sent by player S or he or she will see the note "Player S has not sent you a message." if player S has decided not to send a message. Then, player E must enter his or her guess about the actual number of player S. Here, every number can be entered in 0.5-intervals (0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5).

Payoff

The payment is determined by the following rules: The higher the guess of player E, the higher the payment of player S. And the closer the guess of player E is to the true number of player S, the higher the payment of player E.

Hereafter, the mechanism which determines the payment is explained in detail.

In each round you can earn between 0 and 100 **probability points (PP)**. The more probability points you earn, the higher the probability that you win the subsequent lottery:

If you gain 0 probability points you receive with certainty (with 100%) 1 Euro. If you gain 100 probability points you receive with certainty (with 100%) 8 Euro. If you gain e.g. 70 probability points, you receive, with the probability of 70%, 8 Euro and, with the probability of 30% 1 Euro. The more probability points you gain, the more probable it is that you receive 8 instead of 1.

Thus, you should try to gain as many probability points as possible.

The amount of your probability points in one round depends on both the number of player S and the **guess of player E**. The **payoff table**, which you can find at your spot, makes this clear. If player S e.g. receives the number 3 and player E guesses number 4.5, player E gains 79 probability points and player S 96 probability points. But, if player E guesses that the number of Player S is 1, player E gains 69 probability points and player S only 19 probability points.

Only one of the 15 rounds is chosen randomly and then is actually relevant to the payoff. Your probability points in this round determine the lottery that is played by the computer at the end of the experiment. Since you do not know, which of the 15 rounds is relevant to the payoff you should think carefully about your decisions in each round.

Summary

- Player S receives a random number that is unknown to player E.
- Player S decides whether or not to send a message to player E regarding the number. The message must contain the number of player S.
- What the message contains is determined by player S.
- Player E must guess the number of player S.
- The higher player E guesses the number of player S, the higher the chances of achieving a higher profit for player S.
- The more accurate the guess of player E for the number is, the higher the chances of profits for player E.

F.2 Precise treatment

This experiment is composed of 15 rounds. At the beginning of this experiment, it will be determined randomly whether you are **player S** or **player E**. You will keep this role in all 15 rounds. In each round you play a game with a randomly chosen participant in the opposite role. It is very unlikely, that you are paired up with the same participant in two consecutive rounds.

The Game

In each round, player S receives a number on the **range 0, 1, 2, 3, 4, 5** via the computer. All the numbers are equally probable. Player E does not see which number player S receives. However, player S can send a message regarding his or her number to player E. Player E must guess the number of player S.

At the end of each round both players are informed about the number of player S and the guess of player E.

Decision of player S

After receiving the number, player S can decide about whether or not he or she would like to send a message to the recipient. If player S does send a message, player E will be informed about the number. If player S does not send a message, player E will not be informed about the number.

When player S has received e.g. the number 3, he or she can send a message that contains the true number or send no message at all. This, for example, applies to the message "My number is 3". Graphically, the message "My number is 3" will be depicted by a green box above number 3 and red crosses above 0, 1, 2, 4 and 5:

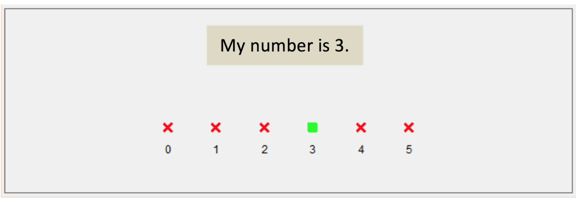

Decision of player E

Player E either sees the message sent by player S or sees the note "Player S has not sent you a message." if player S has decided not to send a message.

Then, player E must enter his or her guess about the actual number of player S. Here, every number can be entered in 0.5-intervals (0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5).

Payoff

The payment is determined by the following rules: The higher the guess of player E, the higher the payment of player S. And the closer the guess of player E is to the true number of player S, the higher the payment of player E.

Hereafter, the mechanism which determines the payment is explained in detail.

In each round you can earn between 0 and 100 **probability points (PP)**. The more probability points you earn, the higher the probability that you win the subsequent lottery:

Probability	**Payoff**
PP%	8 Euro
(100-PP)%	1 Euro

If you gain 0 probability points you receive with certainty (with 100%) 1 Euro. If you gain 100 probability points you receive with certainty (with 100%) 8 Euro. If you gain e.g. 70 probability points, you receive, with the probability of 70%, 8 Euro and, with the probability of 30% 1 Euro. The more probability points you gain, the more probable it is that you receive 8 Euro instead of 1 Euro.

Thus, you should try to gain as many probability points as possible.

The amount of your probability points in one round depends on both the number of player S and the **guess of player E**. The **payoff table**, which you can find at your spot, makes this clear. If player S e.g. receives the number 3 and player E guesses number 4.5, player E gains 79 probability points and player S 96 probability points. But, if player E guesses that the number of player S is 1, then player E gains 69 probability points and player S only 19 probability points.

Only one of the 15 rounds is chosen randomly and then is actually relevant to the payoff. Your probability points in this round determine the lottery which is played by the computer at the end of the experiment. Since you do not know, which of the 15 rounds is relevant to the payoff you should think about your decisions in each round.

Summary

- Player S receives a random number that is unknown to player E.
- Player S decides whether or not to send a message to player E regarding the number. The message must contain the number of player S.
- What the message contains is determined by player S.
- Player E must guess the number of player S.
- The higher player E guesses the number of player S, the higher the chances of profits for player S.
- The more accurate the guess of player E for the number is, the higher the chances of profits for player E.

F.3 Payoff tables

Table 9 Payoffs of Player E

		Guess of Player E										
		0	**0.5**	**1**	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**
Number of Player S	**0**	100	96	88	79	69	58	46	32	19	4	0
	1	88	96	100	96	88	79	69	58	46	32	19
	2	69	79	88	96	100	96	88	79	69	58	46
	3	46	58	69	79	88	96	100	96	88	79	69
	4	19	32	46	58	69	79	88	96	100	96	88
	5	0	4	19	32	46	58	69	79	88	96	100

Table 10 Payoffs of Player S

		Guess of Player E										
		0	**0.5**	**1**	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**
Number of Player S	**0**	0	4	19	32	46	58	69	79	88	96	100
	1	0	4	19	32	46	58	69	79	88	96	100
	2	0	4	19	32	46	58	69	79	88	96	100
	3	0	4	19	32	46	58	69	79	88	96	100
	4	0	4	19	32	46	58	69	79	88	96	100
	5	0	4	19	32	46	58	69	79	88	96	100

F.4 Out-of-sample belief elicitations

At the end of the experiment, subjects received the following questions based on their role and treatment condition. Some questions, as indicated below, were incentivized using the average behavior of subjects that participated in the pilot session as a benchmark.

F.4.1 Receiver beliefs in Precise

Please answer the following questions. You can earn additional money with your answers. In your answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 14 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players S who did not send a message to player E had the following number?

0:[ ]; 1:[ ]; 2:[ ]; 3:[ ]; 4:[ ]; 5:[ ]

F.4.2 Sender beliefs in Precise

Please answer the following questions. You can earn additional money with your answers. In your answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 8 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players E guessed the following number when they did not receive a message from player S?

0:[ ]; 0.5:[ ]; 1:[ ]; 1.5:[ ]; 2:[ ]; 2.5:[ ]; 3:[ ]; 3.5:[ ]; 4:[ ]; 4.5:[ ]; 5:[ ]

F.4.3 Receiver beliefs in Flexible

Please answer the following questions. You can earn additional money with your answers. In your answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 14 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players S had the following number when they did not send a message to player E?

0:[ ]; 1:[ ]; 2:[ ]; 3:[ ]; 4:[ ]; 5:[ ]

[On new screen:] Additionally, please answer the following questions. Refer again to the first round of the experiment that has already taken place at MELESSA.

What was the most common message sent to player E when player S had the following numbers?

(Please always state the upper and the lower number of a message. Example: For the message "My number is 3, 4, or 5", "3" is the lower number and "5" is the upper number. You should enter "3" in the left box and "5" in the right box. If a message only contains

one number, then this number should be entered as the lower as well as the upper number. If no message is sent, leave both boxes blank.)
0:[ ] to [ ]; 1:[ ] to [ ]; 2:[ ] to [ ]; 3:[ ] to [ ]; 4:[ ] to [ ]; 5:[ ] to [ ]

F.4.4 Sender beliefs in Flexible

Please answer the following questions. You can earn additional money with your answers. In your answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 8 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players E have guessed the following if they had not received a message from player S?
0:[ ]; 0.5:[ ]; 1:[ ]; 1.5:[ ]; 2:[ ]; 2.5:[ ]; 3:[ ]; 3.5:[ ]; 4:[ ]; 4.5:[ ]; 5:[ ]

[On new screen:] Additionally, please answer the following questions. Refer again to the first round of the experiment that has already taken place at MELESSA.

What was the average guess of player E when player S sent the following message?
"My number is 1, 2, 3, 4, or 5.":[ ] "My number is 2, 3, 4, or 5.":[ ] "My number is 3, 4, or 5.":[ ] "My number is 4 or 5.":[ ]