Vague Messages in Biased Information Transmission: Experiments and Theory

Daniel H. Wood*

April 20, 2015

Preliminary and incomplete.

Abstract

Spoken language allows for rich communication, but the message spaces used in most cheap talk models and experiments are usually quite restrictive. I show theoretically that introducing vague messages into a strategic information transmission game *a la* Crawford and Sobel (1982) increases communication between boundedly rational players if some senders are moderately honest. My model treats vague messages as explicitly imprecise messages, e.g., "the state is 1, 2, or 3" in contrast to a precise message, which might say "the state is 2". Senders would like to bias the receivers' beliefs upwards. Theoretically the introduction of vague messages causes more honest senders in some cases to send a truthful but vague message rather a precise lie. These message switches replace low-information lies with more informative message and have an additional indirect effect of making the remaining precise messages more informative as well. increasing how informative the average message is about the state. I test this prediction experimentally and find that messages are more likely to be truthful and to be believed credulously in a treatment where both kinds of messages can be sent. Finally I structurally estimate the parameters of my model, and find that about half of subjects get utility from truth-telling equal to around \$0.30 (average earnings per round are around \$0.70).

Keywords: overcommunication, cheap talk, strategic information transmission, level-k, honesty, imprecise messages, vagueness.

*Department of Economics, Clemson University, email: dwood2@clemson.edu.

1 Introduction

A salesman might say that he receives "few complaints" about an appliance he is trying to sell, even if eight of the twelve customers who bought it in the past had complained. A friend might explain her absence from a Sunday afternoon gathering she did not wish to attend by saying she was "sick all weekend", even if her temperature never went above 99.5 and she felt well enough to go to a movie. Like these examples, the majority of biased communication in the real world probably is non-numeric in nature or at least not precise, even if the subject is quantitative.[1] In contrast, the majority of strategic information transmission experiments involve precise communication of quantitative information.

The salesman or friend chooses from a richer (infinite) set of messages than the participant, and the salesman or friend can choose ambiguous messages, messages which lack a definite meaning.[2] I investigate experimentally and theoretically how allowing players to send explicitly imprecise messages affects strategic information transmission in a standard one-dimensional biased cheap talk setting (Crawford and Sobel 1982).

In my experiment, subjects communicate about a state s, an integer between 1 and 5. The sender observes s and chooses freely from a menu of English-language messages to send about it.[3] In the precise language treatment, the menu consists of every message that specifies the state precisely, such as "The state is 3." In the vague language treatment, these five precise messages are still on the menu, but in addition there are three messages that specify an interval the state is within,

[1] I am unaware of attempts to measure the forms that biased communication take, but non-precise speech is prevalent in every-day life. Notable exceptions to this claim are wine recommendations and stock analysts' summary ratings.

[2] In addition, the salesman can use messages which lack a definite referent, such as a "few", a description whose boundaries are fuzzy. In the philosophy literature this is the defining characteristic of vague messages. Lacking a definite referrent necessarily implies imprecision, though, and arguably imprecision is more important to the nature of communication in practice.

[3] Therefore my messages have exogenously given meaning. Much research has focused on fully endogenous meanings instead.

either 1 to 3, 2 to 4, or 3 to 5; for example, "The state is 1, 2, or 3." After hearing whatever message the sender selected, the receiver then chooses an action a. The receiver would like to choose $a = s$ while the sender would like her to choose $a = s + 2$, so they play a partial common interest game (Blume et al. 2001). Subjects alternate roles and are randomly rematched each period in order to avoid repeated game effects. Section 2 describes the experimental setup in detail and Appendix A contains the experimental instructions.

In both treatments, every Nash equilibrium with payoff-maximizing players is a no-communication one. In these equilibria, no m conveys any information about s. As a result a is independent of m. This prediction is not borne out in my data, as a higher m predicts a higher s ($corr(m, s) = 0.311$) and receivers choose higher a in response to higher m ($corr(a, m) = 0.491$). These results are consistent with previous experiments on strategic information transmission, which have typically found that people "overcommunicate": messages convey more information than optimal for senders, and receivers are more trusting of messages than would be optimal if messages were uninformative (e.g., Cai and Wang (2006) or Duffy et al. (2014)).

The prior literature has identified two sources for this excess communication. First, people prefer truth-telling (Gneezy 2005; Sanchez-Pages and Vorsatz 2007, 2009). Second, experimental participants exhibit boundedly rational behaviors which can lead to both more accurate messages and more trusting responses than the most sophisticated participants would employ (Cai and Wang 2006; Wang et al. 2010). No papers that I am aware of attempt to compare the importance of these factors. Instead experiments cleanly test one factor and different experiments find support for both. The first contribution of this paper is to show that both factors are important and that, if the language available is rich, both can combine together to magnify overcommunication.

The second contribution of this paper is to consider how the option of imprecise messages affects strategic information transmission given these factors. If players prefer telling the truth, they may want to choose a different kind of message rather than telling a precise lie. Subjects in some cheap talk games choose non-communication over deceptive communication (Sanchez-Pages and Vorsatz 2007, 2009). However, in these experiments, the sender's direction of bias is unknown to the receiver and not communicating does not reveal any information about the true state. When the sender has a known upward bias and the state space is bounded, in high states the sender would want to truthfully communicate about and so silence in fact conveys information.

Prior experiments also provide limited guidance on how boundedly rational players respond to a richer message space, which makes both composing a message and interpreting a message more complex. Increased complexity could make players who are boundedly rational rely on simpler strategies, perhaps making their messages more truthful and their actions more trusting. Alternatively, it might introduce more noise into their decision-making, producing the opposite effect on accuracy and trust.

In fact, communication is more effective in the vague treatment. 46% (195/424) of messages are truthful in the vague treatment, compared to 24% (65/268) in the precise treatment. Receivers are also more trusting in the vague treatment, choosing the same action as the the state specified in the message 34% (146/424) of the time, as compared to 23% (61/207) of the time in the precise treatment. Section 3 documents the qualitative differences in communication between the two treatments.

My reduced-form evidence does not speak to why communication is better in the vague treatment. Therefore in Section 4 I develop a model combining level-k strategic reasoning and heterogenous tastes for honesty that teases out distinct implications of honesty and bounded rationality in strategic information transmission games with imprecise and precise messages. Level-k models are structural non-equilibrium models of how people reason strategically (Crawford et al. 2013; Stahl and Wilson 1994; Nagel 1995). In these models players have a (typically incorrect) belief about how other players are acting, and optimize given their belief. Level 0 ($L0$) players act non-strategically – in this context by either being truthful or by being completely credulous – $L1$ players believe all players are $L0$, $L2$ players believe all players are $L1$, and so on. Crawford (2003) first applied level-k to cheap talk and his model was subsequently modified and used to explain overcommunication in strategic information transmission games by Cai and Wang (2006).

I introduce a second dimension to players' types: each has an honesty type in addition to a cognitive type. An honesty type is simply how strong a preference for truth-telling that player has. If cognitive types and honesty types are independent, my behavioral model has two vectors of parameters of interest: λ – the distribution of cognitive types – and τ – the distribution of honesty types. In Section 5 I use maximum likelihood to estimate these parameters. Most of the subjects in my experiment employed one or two levels of reasoning. About half of the subjects showed moderate preferences for honesty – getting a utility benefit equivalent to a payment of 30 additional cents per message from being truthful (where the maximum monetary payoff is one dollar per message). These parameter estimates suggest that honesty is the primary cause of the increase in communication between treatments; in the absence of honest types, communication is predicted to be equally effective in both treatments.

About half (192/424) of the messages in the vague treatment were imprecise. This pattern of use allows me to disentangle honesty from bounded rationality. Extensive use of imprecise messages is consistent with weakly honest players sending a message "The state is 1, 2, or 3" when the state really is 1, rather than lying outright and sending the message "The state is 3". My model predicts that these players believe that precise lies are more effective deceptions, and non-honest players would use imprecise messages in fewer circumstances. My other major source of identification is the pattern of deviations from the predicted behavior of level-k self-interested players. In my parameter estimation I use a discrete choice model with logit errors, so even if truthful messages are not optimal, they are more frequently sent by relatively more honest players.

Three experimental papers are closely related in subject or design to this paper. My experimental setup in the precise treatment mirrors Cai and Wang (2006).[4] They use level-k to explain over-communication in their experiment, but do not allow preferences for truth-telling in their model. Agranov and Schotter (2012) is less related in design but more related in subject: a social planner uses vague language to limit subjects' information about payoffs in an asymmetric coordination game, which leads to less coordination failure.[5] Finally, Serra-Garcia, van Damme and Potters (2011) studies experimentally the use of vague messages in a leader-follower public goods game and finds that the leader uses vagueness strategically to hide information that she would hide by an outright lie when only precise messages are allowed. Their structure is similar to mine in having a precise and imprecise treatment, but otherwise the game in their experiment is quite different from that in mine; for instance, the leader takes an action which is observed by the follower before she makes her decision. My study complements theirs by showing that vague language is used to avoid outright lies in a classic strategic information transmission context.

There is also a sizeable theoretical literature related to my subject. Kartik, Ottaviani and Squintani (2007) consider biased strategic information with senders who have costs of misrepresenting information, which could stem from preferences for truth-telling, and a fraction of receivers who are credulous. In their environment fully separating equilibria are possible, while behavior in neither of my treatments resembles a fully separating equilibrium. However, a crucial assumption of that paper is that the state and message space are unbounded above, while in my experiment the boundedness of the state space plays an important role (e.g., about half of messages that subjects

[4] While Cai and Wang allow senders to choose multiple states in their message, note that this was not observed by receivers. If a sender chose "1 or 2" in Cai and Wang's experiment, this meant sending each message "1" and "2" with 50% probability. The receiver did not observe the set $\{1, 2\}$.

[5] Interestingly, Agranov and Schotter have both an imprecise ("$x \in \{1, 2, 3\}$") and vague ("x is low") treatment and find that either form of coarse information is similar.

sent were that "the state is 5", which was the highest possible state).[6] Chen (2011) similarly assumes a fraction of players are honest senders or naive receivers and characterizes equilibria for bounded state spaces in their presence. Some of Chen's qualitative results are consistent with my findings. However, in my work the cognitive types that perform one level of best-response reasoning are the most important both for my theoretical results and in my estimates of the distribution of cognitive types. In Chen's model only L0 types and equilibrium types exist. Neither of these papers considers imprecise messages. In contrast, Blume and Board (2014) model intentional use of vague messages, which is partly my focus, in their model vague messages are noisy messages rather than set-valued as in my framework. Nonetheless, my experimental results are consistent with some of the characteristics of the equilibrium in their paper: namely that vague messages are not sent in the highest state but are sent in lower states, and that vague messages are sent in equilibrium. More generally, it is well-known in this literature that noise in the communication process can actually improve information transmission in equilibrium (Myerson 1997; Blume et al. 2007). If imprecise messages represent a form of noise, my paper supports this conclusion.

2 Experiment

Before introducing and estimating my structural model, I provide an atheoretical description of the experiment – this section – and subjects' behavior – Section 3. The experiment was similar to Cai and Wang (2006). Each period subjects played a discrete version of the Crawford and Sobel (1982) cheap talk game. Subjects were paired, with one member of each pair being the sender and one being the receiver. They communicated about the state of the world, $s \in S \equiv \{1, 2, \ldots, 5\}$ in the following manner:

[6]It seems natural to think of communication about probabilities, reliability of appliances, or illness as having bounded spaces (the upper bounds being 1, totally reliable, and dead respectively. For other contexts, such as stock returns, unbounded spaces may be more natural instead.

1. Nature draws s. Each element of S is equally likely to be drawn.

2. The sender learns s and chooses $m \in M$ to send to the receiver.

3. Finally, the receiver learns m and chooses an $a \in A = S$.

Both the sender and receiver monetary payoffs (in pennies) that were functions of a and s:

$$
\begin{aligned}
u_s(s,a) &= 100 - 25 * (s - a - 2)^{1.2} \\
u_r(s,a) &= 100 - 25 * (s - a)^{1.2}
\end{aligned}
$$

Subjects saw payoffs presented in table form for both roles similarly to Figure 1 below, so payoffs were common knowledge.

	Sender payoffs					*Receiver payoffs*				
Receiver action choice	**State**					**State**				
	s=1	s=2	s=3	s=4	s=5	s=1	s=2	s=3	s=4	s=5
a=1	43	7	-32	-72	-115	100	75	43	7	-32
a=2	75	43	7	-32	-72	75	100	75	43	7
a=3	100	75	43	7	-32	43	75	100	75	43
a=4	75	100	75	43	7	7	43	75	100	75
a=5	43	75	100	75	43	-32	7	43	75	100

Figure 1: Payoff Tables

There were three conditions. In the *precise* message condition, the sender could send only precise messages of the form "The state is x.". i.e., $M_P = \{1, 2, 3, 4, 5\}$. In the *vague* message condition, the sender could choose to send either precise messages or imprecise messages. Imprecise messages specified an interval of three numbers that s could be within, such as “The state is 2, 3, or 4.”. $M_V = \{1, 2, 3, 4, 5, \{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}\}$.

Finally in the third condition, the *noise* condition, the sender observed s imperfectly and could send precise messages or imprecise messages ($M_N = M_V$). Instead of observing s, the sender observed $\sigma \in \Sigma$, where

$$\Sigma = \{\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}\},$$

and s was equally likely to be any element of σ. Therefore in the noise treatment, $\Pr(s = 1) = \Pr(s = 5) = 0.155$, $\Pr(s = 2) = \Pr(s = 4) = 0.211$, and $\Pr(s = 3) = 0.266$.

In practice behavior in the noise and vague treatments were similar, and so I add the noise data to the vague data, setting $s = \mathrm{E}[s \mid \sigma]$.

Instructions for each treatment were handed out and read aloud at the start of that treatment (Appendix A contains the experimental instructions). All subjects participated in all conditions, and the order of the conditions was randomized for each session. There were 8 to 12 rounds for each condition. Subjects alternated roles and were randomly matched each round. Average earnings were around \$24, which includes a \$6 show-up payment.

In total there were 42 subjects who collectively participated in 268 rounds of the precise treatment, 204 rounds of the vague treatment, and 220 rounds of the noise treatment. Note that a round yields one message choice observation and one action choice observation.

3 Results

The unique perfect Bayesian equilibrium if players are self-interested and rational is for no communication to occur: for messages to transmit no information about the state and for actions chosen to be independent of the messages received. Neither of these predictions is borne out.

See Appendix B for raw data.

3.1 Truthfulness

Consider first the fraction of truthful messages in each treatment. Precise messages are truthful if its literal meaning is true, i.e., $m = s$.[7] There is some ambiguity in which messages are truthful in the noise and vague message treatments. The definition $m = s$ is narrow, as it classifies all imprecise messages as false. A wider definition might also include imprecise messages for whicj $s \in m$. The second definition requires accuracy but not precision, so that the statement "The state is 1, 2, or 3." is true if the state is 2. As measured with this definition, senders were more likely to send true messages in the vague condition than the precise condition. Among the subset of messages that were precise, there was also more narrowly defined truth-telling in the vague treatment.

A related question is how informative each message was. Figure 2 shows Bayesian posterior probabilities $\Pr(s = x|m, \text{treatment})$ for each message in the precise and vague treatments. For $m = 1$ and $m = 2$ the modal state from which the message was sent was the true state, while for $m > 2$ with the exception of precise $m = 4$ in the vague treatment, the modal state was *never* the true state. While $m < 3$ were more accurate, they were less common than $m \geq 3$ ($\Pr(m < 3) = 0.172$ and $\Pr(m < 3) = 0.125$ in the precise treatment and in the vague treatment

[7] While Sutter (2009) finds evidence that in some games senders tell the truth deceptively, that seems implausible in my context with sender bias, and indeed in my model no sender has beliefs that would lead to deceptive truth-telling.

Table 1: Truthfulness

	Treatment					
	Precise		**Vague**		**Total**	
	No.	%	No.	%	No.	%
Truthfulness (strict definition)						
False message	203	76%	354	83%	557	80%
Truthful message	65	24%	70	17%	135	20%
Total	268	100%	424	100%	692	100%
Truthfulness (weaker definition)						
False message	203	76%	229	54%	432	62%
Truthful message	65	24%	195	46%	260	38%
Total	268	100%	424	100%	692	100%
Truthfulness (strict definition, precise m)						
False message	203	76%	162	70%	365	73%
Truthful message	65	24%	70	30%	135	27%
Total	268	100%	232	100%	500	100%
Truthfulness (weaker definition, vague m)						
False message	–	–	67	35%	67	35%
Truthful message	–	–	125	65%	125	65%
Total	–	–	192	100%	192	100%

Source: novembexp.dta

respectively).

Messages $m = 3$ or $m = \{2, 3, 4\}$ were accurate in neither treatment, being most likely to have been sent when $s = 1$. In general, higher messages ($m = 4$, $m = 5$, and $m = \{3, 4, 5\}$) were closer to the truth in the vague treatment than corresponding messages in the precise treatment. Visual comparison of the series shows that the CDF of $\Pr(s \mid m)$ in the vague treatment first-order stochastically dominate their counterparts in the precise. There is a single crossing point $\bar{s}$ above which $\Pr(s \mid m, \text{precise}) < \Pr(s \mid m, \text{vague})$ and below which $\Pr(s \mid m, \text{precise}) > \Pr(s \mid m, \text{vague})$. Because high messages make up the majority of messages sent, this regularity is evidence that senders transmitted more information in the vague treatment.

A final point to take from Figure 2 is that messages were quite noisy in both treatments. With the exception of $m = 1$ in the precise treatment, $\Pr(s \mid m) < 1/2$ for any s, and generally $\Pr(s \mid m) > 1/10$ for all s. This noisiness means that usually receivers maximize utility by adjusting their action choice inward towards $a = 3$, and in fact $a = 2$ and $a = 3$ maximize receivers' payoffs in every case except $m = 1$ (both conditions) or $m = 4$ or $m = 5$ (vague condition).[8]

[8] Optimal action choices are circled on the corresponding series – for instance in the $m = 4$ figure, $a^*(4)$ (precise treatment) and $a^*(\{3, 4, 5\})$ both equal 3, while $a^*(4)$ (vague treatment) is 4.

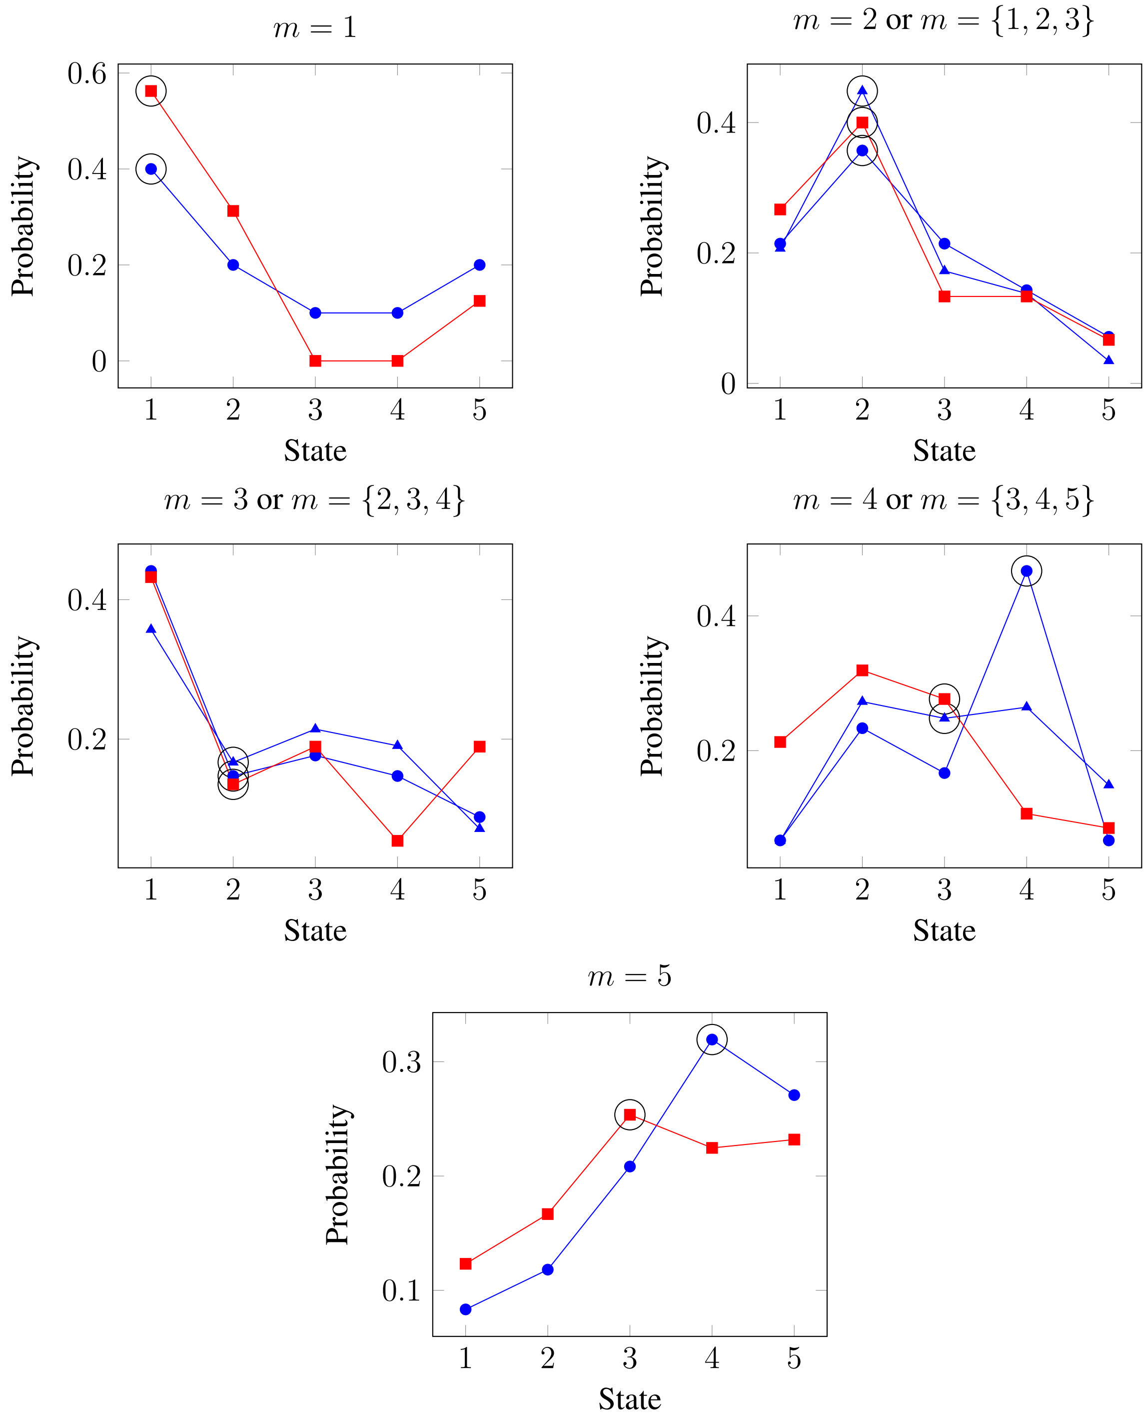

Figure 2: Posterior Probabilities $\Pr(s \mid m)$. Red squares are posteriors in precise treatment, blue circles are posteriors for precise message in vague treatment, and blue triangles are posteriors for imprecise message in vague treatment. Circled points are optimal action choices by Bayesian receiver.

3.2 Trust in Messages

Define a credulous a as an action equal to the message ($a = m$) for a point message or the average of the states in the message ($a = (\sum_{s \in m} s)/3$) for an imprecise message. Table 2 shows that credulity was more widespread in the vague and noise conditions than the precise condition.

Table 2: Credulity

	Treatment					
	Precise		**Vague**		**Total**	
	No.	%	No.	%	No.	%
Same a as m						
No	207	77%	278	66%	485	70%
Yes	61	23%	146	34%	207	30%
Total	268	100%	424	100%	692	100%
Same a as m (precise m only)						
No	207	77%	160	69%	367	73%
Yes	61	23%	72	31%	133	27%
Total	268	100%	232	100%	500	100%
Same a as expected m (imprecise m only)						
No	–	–	118	61%	118	61%
Yes	–	–	74	39%	74	39%
Total	–	–	192	100%	192	100%

Source: novembexp.dta

Figure 3 shows the frequency of each action for each message sent, $\Pr(a \mid m)$, in the precise and vague treatments. In the precise condition, there is widespread discounting of messages $m \geq 3$ by two units, such as choosing $a = 1$ in response to $m = 3$. While this phenomenon occurs in the vague treatment as well, subjects in that treatment more frequently discount by 0 or 1 instead. In addition, imprecise messages are discounted less than a corresponding precise messages in the vague treatment. The differences in receivers' responses between treatments are less pronounced, though, than the differences in sender behavior in Figure 2.

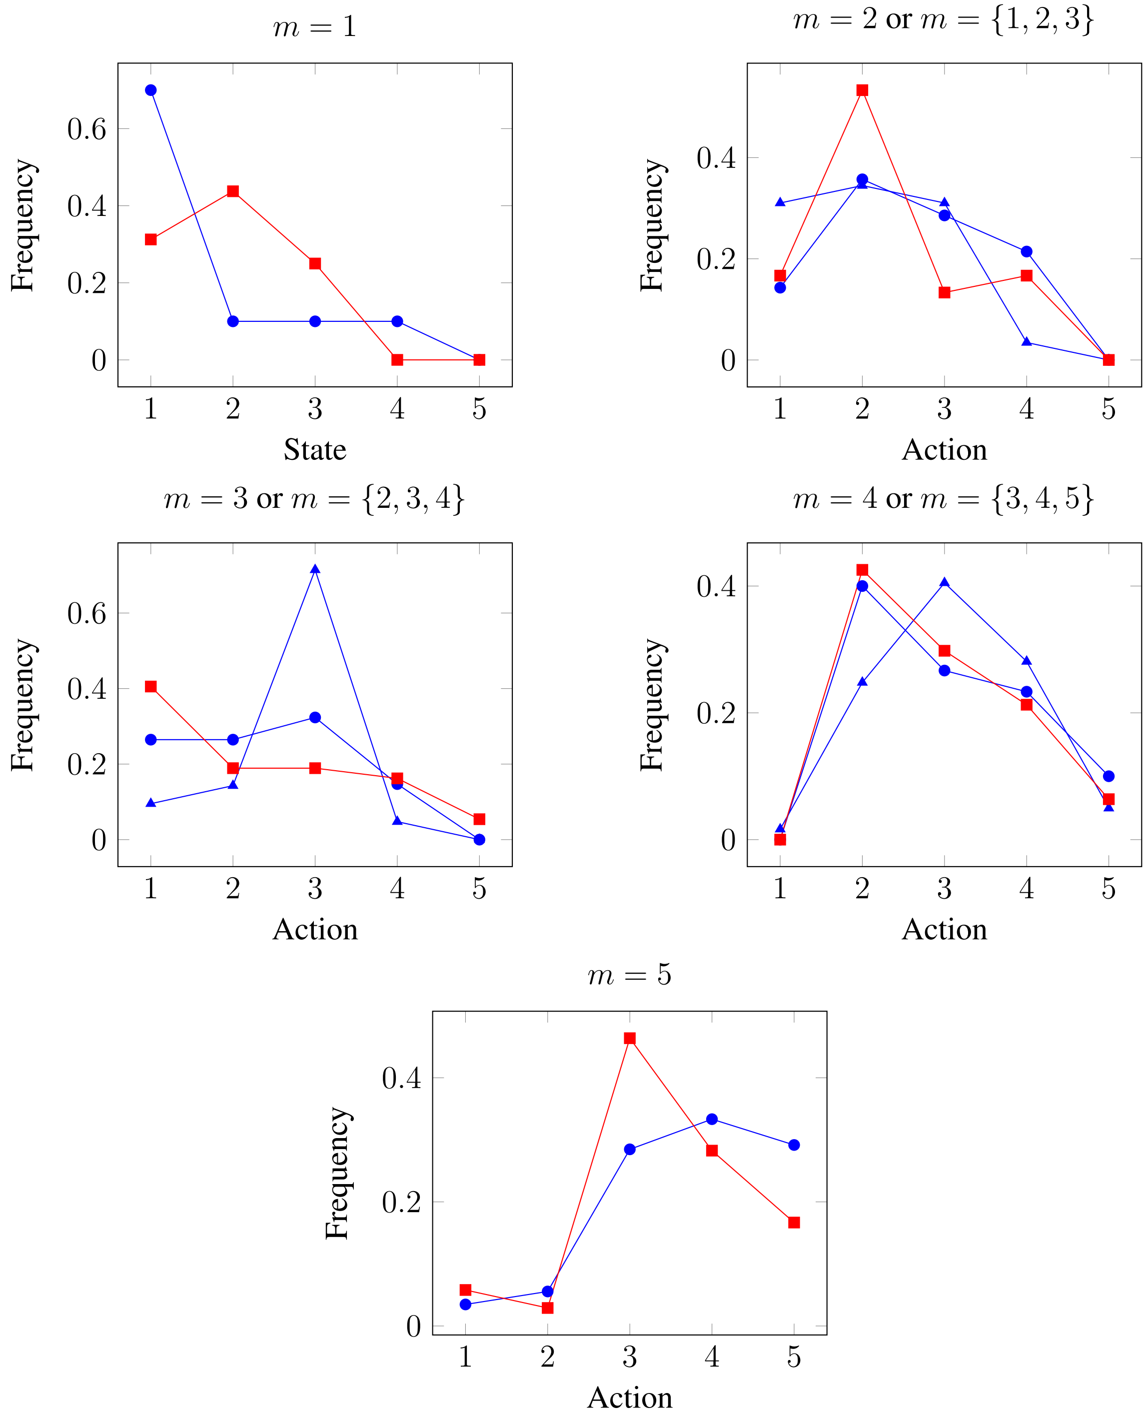

Figure 3: Action choice frequencies for each message received. Blue circles are responses to precise messages in vague treatment, blue triangles are responses to imprecise messages, and red squares are responses in precise treatment.

3.3 Correlations between message and action choices

Not done yet.

3.4 Dynamics

There is no evidence that on aggregate players' strategies changed during the sessions. In particular, the fraction of imprecise messages is roughly constant in early and late rounds, as is the share of truthful messages and the degree of discounting by receivers. (*More to come.*)

4 Model

Level-k models are structural models of how players reason strategically. Instead of strategy choices being in some form of equilibrium, players have a (typically incorrect) belief about how other players are acting, and optimize given their belief. Crawford (2003) first applied level-k to cheap talk and his model was adopted to a discrete version of the canonical Crawford and Sobel (1982) model of strategic information transmission by Cai and Wang (2006). My main addition to these earlier models is heterogenous preferences for truth-telling on the part of senders.

In these models, there is a hierarchy of cognitive types. Each type has different beliefs about the strategy choices of the rest of the population, generally based to how less sophisticated types behave. Each type reasons strategically based on their beliefs. For instance, in Crawford (2003)'s model, communication is about a binary state, and L0 senders are truthful ($m = s$) while L0 receivers are credulous ($a = m$); then Lk senders best-respond to $Lk - 1$ receivers while Lk receivers best-respond to $Lk - 1$ senders. In communication about a binary state, $L1$ receivers are also credulous, $L1$ senders are liars, $L2$ receivers invert the message ($a = 1$ if $m = 0$, $a = 0$ if $m = 1$), $L2$ senders are liars, $L3$ receivers invert m, and $L3$ senders deceptively tell the truth. Because senders in my experiment are biased, senders of higher types exaggerate their reports of s more and more, while receivers discount the senders' reports more if they are of a higher types. This exaggeration and discounting occasions a reduction in communication, as sufficiently sophisticated senders only send the highest message $m = 5$ and sophisticated recievers then choose $a(5) = 3$. In common with the Crawford model, however, is that $L0^R$ and $L1^R$ behavior is identical, as is $L1^S$ and $L2^S$ behavior, and $L2^R$ and $L3^R$ behavior.[9]

Formally, players differ in their cognitive type, honesty level, and role. A player's cognitive type determines her beliefs about the how players in the other role behave. A type Lk^s player has beliefs $B_k^s(a, m)$ over the conditional probabilities $\Pr(a \mid m)$ induced by receiver behavior, and a type Lk^r player has beliefs $B_k^r(m, s)$ over the the conditional probabilities $\Pr(m \mid s)$ that she faces.

A player's honesty type affects her relative utilities from sending dishonest and honest messages. Following the weak definition of truthfulness in Section 3.1, an honesty type h acts as if sending a truthful message is worth giving up h cents:

$$u_s^h(s, a, h) = \begin{cases} 100 - 25 * (s - a - 2)^{1.2} + h & \text{if } s = m \text{ or } s \in m \\ 100 - 25 * (s - a - 2)^{1.2} & \text{otherwise} \end{cases}.$$

A receiver's utility function is not affected by her honesty type.[10] I assume that there is a finite set of honesty types, with every $h \in H$. The vector τ denotes the distribution of honest types, and

[9] Because sender and receiver behaviora predictions differ, the entire distribution of types can still be estimated using message and action choices. If I estimate distributions using only one's roles choices, though, separate estimation of the identical levels is largely impossible.

Cai and Wang (2006) define types slightly differently: Lk senders best-respond to L$k-1$ receivers but Lk receivers best-respond to Lk senders. As Crawford et al. (2013) notes, these definitions are partly semantic. However, they are not entirely so for me, as I impose that $\lambda_k^R = \lambda_k^S$.

[10] While there is evidence that people get disutility from being tricked above and beyond the pecuniary costs, it is unclear how those preferences would translate into behavior in my experiment.

I write τ_h to indicate the fraction of players with honesty type h. Finally I assume H and τ are common knowledge.

A level k sender chooses for each state a (possibly) mixed strategy over messages $\sigma_{s,k}(m \mid s)$. If $k \geq 1$ $\sigma_{s,k}(s)$ is a best response to that type's beliefs $B_k^s(a,m)$ about receiver behavior. Likewise a level k receiver chooses $\sigma_{a,k}(a \mid m)$ which is a best response to her beliefs B_k^r about sender behavior if if $k \geq 1$. $k = 0$ types are non-strategic and primarily serve as conjectures about behavior by $L1$ players.

Definition 1. A type $L0^S$ player sends messages such that $\Pr(s = m \mid \sigma_{s,0}) = 1$ if m is precise, or $\Pr(s \in m \mid \sigma_{s,0}) = 1/3$, if m is imprecise. A type $L0^R$ player chooses actions such that $\Pr(a = m \mid \sigma_{r,0}) = 1$ if m is precise or $\Pr(a = \mathrm{E}[s \mid s \in m$ uniformly distributed$]) = 1$.

In other words $L0^s$ messages are truthful, either exactly or in expectation, and $L0^r$ are credulous.[11] These are standard assumptions about non-strategic behavior when applying Lk models to cheap talk, in contrast to the more typical assumption in other contexts that that $L0$ randomize uniformly over their possible actions.

Strategic Lk behavior is then defined inductively as responding optimally to the other role's $Lk-1$ behavior:

Definition 2. For $k \geq 1$, a type Lk^s player with honesty h has beliefs B_k^s and chooses $\sigma_{s,k}$ that

[11]For the vague treatment, truthfulness in expectation requires that appropriate fractions of $L0_s$ send imprecise messages so that $\mathrm{E}[s|m = \{1,2,3\}] = 2$ and so on. Generating this property for imprecise messages in fact requires slightly counterintuitive $L0_S$ behavior; for instance, these senders must be three times as likely to send an imprecise message in $s = 3$ as in $s = 1$ to guarantee $\Pr(s = 1|m = 123) = \Pr(s = 2|m = 123) = \Pr(s = 3|m = 123)$.

follow

$$B_k^s(a,m) = \sigma_{r,k-1}(a \mid m) \tag{1}$$

$$\sigma_{s,k}(m \mid s,h) > 0 \Rightarrow m \in \arg\max_m \mathrm{E}[u_s^h(s,a,h) \mid B_k^s, m] \tag{2}$$

$$= \arg\max_m \sum_a u_s^h(s,a) B_k^s(a,m) \tag{3}$$

and a type Lk^r player has beliefs B_k^r and chooses $\sigma_{r,k}$ that follow

$$B_k^r(m,s) = \sum_{h \in H} \tau_h \sigma_{s,k-1}(m \mid s,h) \tag{4}$$

$$\sigma_{r,k}(a \mid m) > 0 \Rightarrow a \in \arg\max_a \mathrm{E}[u_r(s,a) \mid m, B_k^r] \tag{5}$$

where $\mathrm{E}[u_r(s,a) \mid B_k^r]$ is calculated via Bayes' rule:

$$\mathrm{E}[u_r(s,a) \mid m, B_k^r] = \sum_s \Pr(s \mid m, B_k^r) u_r(s,a) \tag{6}$$

$$= \sum_s \frac{(1/5) B_k^r(m,s)}{\sum_{s'} B_k^r(m,s')}. \tag{7}$$

This definition of heterogenous sophistication implies that all $k \leq 3$ players engage in no more than two rounds of introspective strategic reasoning. $L1$ and $L2$ senders each use one round of strategic reasoning, best-responding to credulous receiver strategy $\sigma_{r,0} = \sigma_{r,1}$, as do $L2$ and $L3$ receivers, who best respond to $\sigma_{s,1} = \sigma_{s,2}$.[12] $L3$ senders engage in two rounds of strategic reasoning, best-responsing to $\sigma_{r,2}$ which is itself a best-response to $\sigma_{s,1}$.

Let λ be a vector denoting the frequency of each cognitive type, so λ_2 is the frequency of $L2$ players. λ and τ together predict frequencies of m and a choices and correlations between

[12]While $L1^R$'s best-respond to $L0^S$, this leads to effectively non-strategic credulous action choices.

m and s and a and m. In that sense applying the model to explaining my experimental data is straightforward. However, interpreting the resulting parameter estimates is made complicated by the fact that each cognitive type k both serves as a prediction about actual observed behavior, if Lk types exist in the population, and as a conjecture about behavior by more sophisticated types. For instance, $L1^s$ behavior serves as $L2^r$ players' model of sender behavior. Because $L3^s$ players use $L2^r$ behavior as their model of receiver behavior, the $L1^s$ behavior influences their conjectures about strategic behavior as well. However, I next turn to the the general predictions of this model, first in the absence of honesty and then given the presence of sufficiently honest. After outlining predictions, I return to the problem of interpretation.

State	*Sender's message*			**Message**	*Receiver's action*		
	$L0^s$	$L1^s$ and $L2^s$	$L3^s$		$L0^r$ and $L1^r$	$L2^r$ and $L3^r$	$L4^r$ and $L5^r$
$s = 1$	1	3	2, 4, and 5	$m = 1$	1	1	1
$s = 2$	2	4	5	$m = 2$	2	2	1
$s = 3$	3	5	5	$m = 3$	3	1	1
$s = 4$	4	5	5	$m = 4$	4	2	1
$s = 5$	5	5	5	$m = 5$	5	4	3

Table 3: Level-k Model – Precise Treatment if Limited Levels of Honesty ($h < 32$)

Table 3 shows the predictions of the Lk model for the precise treatment for low k if no honest type with $h > 32$ exists. Higher cognitive levels converge to playing $m = 5$ (if sender and $k \geq 5$) or continue using the $L4^r$ type's strategy (if receiver).

This model matches several aspects of the aggregate behavior in the precise condition: For $m < 3$, the messages are truthful and receivers trust them. For $m \geq 3$, many messages are exaggerations by two units and receivers discount them, usually by two units. Finally, most of the messages are $m = 5$, but a substantial number are $m < 5$. In contrast, level-k without honesty does a poor job in matching the observed behavior in the vague treatment. Table 4 shows predictions for that treatment.

State	*Sender's message*			Message	*Receiver's action*		
	$L0^S$	$L1^S$ and $L2^S$	$L3^S$		$L0^R$ and $L1^R$	$L2^R$ and $L3^R$	$L4^R$
$s = 1$	1	3 and "2, 3, 4"	2, 4, 5, "1, 2, 3", and "3, 4, 5"	$m = 1$	1	1	1
$s = 2$	2	4 and "3, 4, 5"	5	$m = 2$	2	2	1
$s = 3$	3	5	5	$m = 3$	3	1	1
$s = 4$	4	5	5	$m = 4$	4	2	1
$s = 5$	5	5	5	$m = 5$	5	4	3
				$m = \{1, 2, 3\}$	2	2	1
				$m = \{2, 3, 4\}$	3	1	1
				$m = \{3, 4, 5\}$	4	2	1

Table 4: Level-k Model – Vague Treatment if Honesty Low ($h < 25$)

There are three connected ways that the level-k model without honesty fails to predict observed behavior in the vague treatment. First, many fewer imprecise messages are predicted than are observed. In addition, imprecise messages are predicted to only be sent from $s < 3$. Finally, messages are no more informative than in the precise treatment.

Tables 3 and 4 derive predictions for the model where few players have high h. Adding honest types with intermediate levels of honesty ($25 < h < 57$) causes qualitatively significant changes to the predictions about behavior in the vague conditions without seriously affecting predicted behavior in the precise condition. Consider the opportunity cost of honesty, the loss in pecuniary payoff for sending an honest message versus the money-maximizing message. Table 5 shows the opportunity cost of honesty in the precise treatment given players using the strategies in Tables 3 and 4. On the right side is the opportunity cost in the vague treatment and also whether the optimal honest message would be imprecise (i) or precise (p). $L3^S$ opportunity costs depend on whether honest types affect $L2_R$ choices; the table reports opportunity costs given the same $L1_R$ behavior as in Tables 3 and 4.

Consider the distribution of messages sent by $L1^S$ and $L2^S$ if τ_h fraction of senders have $32 < h < 57$, and the remainder of senders have $h = 0$. Because the opportunity cost is too high, few of these senders choose honest messages in the precise treatment. However, all of the honest $L1^s$ and $L2^s$ choose exaggerated but weakly honest imprecise messages. This implies that both kinds of message become more informative as τ_h increases. Consider first $m = \{2, 3, 4\}$. $L1^s$ with $h < 25$ send $m = \{2, 3, 4\}$ with probability $1/2$ in state $s = 1$ (see Table 4). $L1^s$ with $h > 25$ do not send that message in $s = 1$ (they will send $m = \{1, 2, 3\}$ instead) but do send it in state $s = 2$. Hence

$$\mathrm{E}[s \mid m = \{2,3,4\}] = \frac{1 \cdot (1-\tau_h)/2 + 2\tau_h}{(1-\tau_h)/2 + \tau_h} = \frac{1+3\tau_h}{1+\tau_h}$$

All $1 - \tau_h$ fraction of the $h = 0$ senders in the vague treatment send dishonest messages $m \geq 3$, so messages $m \geq 3$ also become more informative as τ_h increases: the number of $m = 3$ and $m = 4$ drops, and $m = 5$ messages become more accurate:

$$\begin{aligned}
\Pr(s = 3 \mid m = 5) = \Pr(s = 4 \mid m = 5) &= \frac{(1/5)(1-\tau_h)}{(1/5) + (2/5)(1-\tau_h)} = \frac{1-\tau_h}{3-2\tau_h} \\
\Pr(s = 5 \mid m = 5) &= \frac{(1/5)}{(1/5) + (2/5)(1-\tau_h)} = \frac{1}{3-2\tau_h} \\
\mathrm{E}[s \mid m = 5] &= (3+4)\left(\frac{1-\tau_h}{3-2\tau_h}\right) + (5)\left(\frac{1}{3-2\tau_h}\right) \\
&= \frac{12-7\tau_h}{3-2\tau_h}.
\end{aligned}$$

e.g., if $\tau_h = 1/2$, $\mathrm{E}[s \mid m = 5] = 4.25$ compared to $\mathrm{E}[s \mid m = 5] = 4$ if no $h \geq 25$ types exist in the vague treatment or $h \geq 57$ types exist in precise treatment.

State	*Cognitive type*		
	$L0^S$	$L1^S$ and $L2^S$	$L3^S$
1	N/A	57	32
2	N/A	57	57
3	N/A	57	107
4	N/A	32	75
5	N/A	0	0

State	*Cognitive type*		
	$L0^S$	$L1^S$ and $L2^S$	$L3^S$
1	N/A	25 (i)	57 (i)
2	N/A	25 (i)	57 (i, p)
3	N/A	25 (i)	32 (i)
4	N/A	32 (i, p)	36 (i)
5	N/A	0	0

Table 5: Opportunity cost of honesty in precise treatment (left) and vague treatment (right). For instance, $L1^S$ players with $s = 1$ in the vague condition believe $\Pr(a = 2 \mid m = \{1, 2, 3\}) = 1$ and $\Pr(a = 3 \mid m = 3) = 1$, so their perceived utility from the optimal non-honest message is $\mathrm{E}[u(s, a) \mid m = 3] = 100$ and from the optimal honest message is $\mathrm{E}[u(s, a) \mid m = \{1, 2, 3\}] = 75$. If $h \geq 25$, they choose $m = \{1, 2, 3\}$.

In my structural model, the fraction of honest subjects is partially identified through sufficiently honest $L1^s$ and $L2^s$ players' greater tendency to choose imprecise messages. These players believe that sending an appropriate precise message can produce any action, yielding a payoff of 100. However, they also believe that sending an imprecise message that is honest but exaggerates somewhat, such as "1, 2, or 3" when $s = 1$, is almost as good in pecuniary terms, yielding a monetary payoff of 75, but also produces h utils from honesty. If $h > 25$, honest and self-interested $L1^S$ and $L2^S$ mostly separate. Because $L3^S$ usually believe $a(m) < \mathrm{E}[m]$ for imprecise messages and $a(m) < m$ for precise messages, they face higher opportunity costs and so not use truthful imprecise or precise messages unless h is much higher than 25. The presence of highly honest players will also affect sender behavior in the precise treatment. Players with honesty type $h \geq 57$ send wholly truthful messages, but all less honest players act like $h = 0$ players.

Now I return to the issues that arise because cognitive types both serve as a prediction about actual observed behavior and as conjectural inputs into predicting the behavior of more sophisticated types. There are two main implications of this dual nature of cognitive type.

The first implication is that any auxiliary assumptions beyond those in the definitions above are also assumptions both about behaviors and beliefs. Unfortunately, I need to make auxiliary assumptions because problems arise in applying level-k models to games with more than two messages that are reminiscent of classical problems with Nash equilibrium concepts. One problem is that senders are sometimes indifferent between multiple m. In these cases (2) does not determine a unique strategy, and I assume that senders randomize uniformly over their best messages.[13] Another problem is that receivers sometimes observe m that they believe are probability 0, so that Bayes' rule cannot be used in calculating their expected utility in (7). For example, $L2_R$ typically believe no $L1_S$ would ever send $m = 1$. If a Lk_R observes an m of this nature, I assume that she intreprets it as coming from the most sophisticated sender type that would send it, frequently an $L0_S$. I return briefly to these assumptions in the robustness section.

The second implication is that the presence of honesty types in the population represents both behavior and beliefs. For instance if $\tau_h = 1/2$ and h is high enough, then (i) 50% of senders choose more honest messages but also (ii) all receivers believe that 50% of senders choose more honest messages.

[13] This assumption seems trivial for sender behavior. All players in my empirical analysis make logit errors, which mechanically imply that senders place equal probability on messages that they believe produce the highest payoff. If one is willing to accept that logit errors are reasonable, uniform randomization for senders follows. However, players best-respond not to choices with logit errors but to the exact strategy choices.

5 Structural Estimates

Many papers have structurally estimated level-k models before, but to my knowledge none has jointly estimated varying cognition levels and varying preferences. One benefit of jointly estimated λ and τ is that I can answer the counterfactual question of what communication would be like if no honest types existed.

I assume that subjects choose a or m influenced by logistic errors of precision β, independent across rounds and treatments. These errors capture that level-k reasoning is subject to error and that unobserved utility shocks may occur. Hence a level k sender or receiver chooses each m or a with probability

$$\Pr(m \mid s, k, h, \beta, T) = \frac{\exp(\beta\,\mathrm{E}[u_s(a, s, h) \mid m, B_k^s, T])}{\sum_{m'=1}^{5} \exp(\beta\,\mathrm{E}[u_s(a, s, h) \mid m', B_k^s, T])} \tag{8}$$

$$\Pr(a \mid m, k, \beta, T) = \frac{\exp(\beta\,\mathrm{E}[u_r(a, s) \mid B_k^r, T])}{\sum_{a'=1}^{5} \exp(\beta\,\mathrm{E}[u_r(a', s) \mid B_k^r, T])} \tag{9}$$

where B_k^r is a level-k sender's beliefs, B_k^r is a level-k receiver's beliefs, T is the treatment, and h is a sender's honesty type. β measures the noisiness of player choices. As $\beta \to 0$ each choice becomes equally likely, while as $\beta \to \infty$ players best-respond exactly to their beliefs.[14]

It could be that reasoning about optimal choices by more sophisticated players is more error-prone as it requires reasoning through multiple rounds of posterior beliefs. Alternatively, less sophisticated players may be worse at figuring out optimal behavior in addition to being more limited in strategic reasoning. Therefore I allow choice precision to vary by cognitive type, with β_k being the precision of level k players in both roles.

[14]Note that $L0$ players have no beliefs; instead I assume they exactly follow $L0^r$ or $L0^s$ strategies specified in Definition 1.

In addition to the vector β, I also estimate λ, where λ_k is the frequency of level k players. My preferred estimates assume λ is the same for both roles. In addition, I assume no $L4$ or higher players exist. In my level-k model, the predicted behavior of senders above $L3$ is only slightly different from $L3$. For instance, $L4^s$ are behaviorally identical to $L3^s$ while $L5^s$ differ only in their m choices in state 1. While $L4^R$ behavior is more distinct from that of $L3^R$'s, I find limited evidence for $L3^R$ receivers in my base models already.

Finally, for my preferred estimates I restrict honest types to either having $h = 0$ (the "not honest" type) or $h = 30$ (the "weakly honest" type). Most honest types are behaviorally similar and so it is hard to precisely identify the frequencies of more two types of honest player. Section 5.2 focuses separately on both these identification issues and what conclusions can be drawn about the distribution of honesty in my experiment.

Given A_i action choices by subject i, $a_{i,1}, \ldots, a_{i,A_i}$, and M_i messages choices by subject i $m_{i,1}, \ldots, m_{i,M_i}$, the probability of i's observed choices for a given parameterization is

$$L_i(\lambda, h, \beta) = \sum_{k \in K} \lambda_k \sum_{h \in H} h_h \left(\prod_{j=1}^{M_i} \Pr(m_{i,j} \mid s, k, h, \beta) \right) \left(\prod_{j=1}^{A_i} \Pr(a_{i,j} \mid m, k, \beta) \right). \tag{10}$$

Given a set of subjects N, the log likehihood of the observed behavior is

$$LL(\lambda, h, \beta) = \sum_{i \in N} \log L_i(\lambda, h, \beta) \tag{11}$$

I estimate the parameters by maximizing LL.

I derive standard errors and confidence intervals through non-parametric bootstrap. Taking my set of 42 subjects, I create $R = 1000$ resampled datasets of 42 subjects by sampling 42 times with replacement from my dataset. For each resampled dataset, I re-estimate the model parameters. The actual standard deviation of these estimates is my estimate of the true standard errors (Efron

1979). Many of my distributions of parameter estimates are not symmetric, so I also calculate bias-corrected accelerated confidence intervals (DiCiccio and Efron 1996).

Table 6: Base Parameter Estimates.

	(1) base	(2) precise only	(3) vague only	(4) s only	(5) r only
λ_0	0.0	0.006	0.0	0.0	0.115 ***
	[0.0, 0.048]	[0.0, 0.082]	[0.0, 0.083]	[0.0, 0.1]	[0.032, 0.205]
λ_1	0.12 ***	0.524 ***	0.167 ***	0.398 ***	0.441 ***
	[0.034, 0.474]	[0.297, 0.722]	[0.03, 0.342]	[0.19, 0.812]	[0.084, 0.716]
λ_2	0.436 ***	0.106	0.404 ***	0.138	0.38 **
	[0.247, 0.669]	[0.0, 0.451]	[0.151, 0.724]	[0.0, 0.325]	[0.206, 0.786]
λ_3	0.444 ***	0.364 ***	0.428 ***	0.464 ***	0.064
	[0.282, 0.624]	[0.127, 0.542]	[0.208, 0.723]	[0.299, 0.634]	[0.0, 0.406]
τ_0	0.48 ***	0.484 ***	0.482 ***	0.416 ***	0.783 ***
	[0.02, 0.669]	[0.01, 1.0]	[0.021, 0.669]	[0.021, 0.669]	[0.013, 0.931]
τ_{30}	0.52 ***	0.516	0.518 ***	0.584 ***	0.217 ***
	[0.331, 0.98]	[0.0, 0.99]	[0.331, 0.979]	[0.331, 0.979]	[0.069, 0.987]
n	1384	536	848	692	692
$logl$	-2046.145	-652.668	-1383.155	-1061.797	-976.386
BIC	4142.919	1368.177	2826.995	2182.451	2018.168

Maximum likelihood estimates of share of Lk players (λ_k) and share of players with utility from honesty h (τ_h). 95% BCa confidence intervals are reported beneath estimates. * denotes significance at 0.1 level, ** significance at 0.05 level, and *** significance at 0.01 level.

5.1 Basic estimates

Table 6 column 1 reports the maximum likelihood estimates of model parameters for my entire dataset. Estimates of λ are similar to others in the literature: no $L0$ players, 12% $L1$ players, and roughly equal fractions of $L2$ and $L3$ players. Given my definitions of Lk behavior for each role, these estimates imply that 56% of senders engage in one iteration of strategic reasoning and 44% of senders engage in two iterations, while 12% of receivers do not reason strategically (they are credulous) and the remaining 88% of receivers engage in one iteration of strategic reasoning. About half of the subjects have $h = 30$ and are believed to have $h = 30$.

The next columns report my model estimated separately on the data from the precise treatment only (column 2) and the vague treatment only (column 3). Comparing the estimates, reasoning is more strategically sophisticated in the vague treatment, perhaps because subjects have to think harder given the multi-dimensional message space and cannot apply a simple heuristic like exaggerating and discounting by two units.

While point estimates for τ_{30} in the single treatment models (2) and (3) are similar to the combined model (1), the estimates of τ_{30} in model (2) cover the entire interval $[0, 1]$. This imprecision is not surprising, because type $h = 30$ players are behaviorally identical to type $h = 0$ players in the precise treatment. The τ_{30} estimates in model (3) are quite similar to model (1) because the estimates are identified off the same patterns in the data.

Taken together, the increased strategic sophistication in (3) combined with constant honesty estimates in the first three models shows that the increase in communication between treatment stems from prevalent weak honesty combined with the richer message space, rather than increased confusion or adoption of simpler strategies on the part of subjects. In Section 5.3 I use the parameter estimates from model (1) to decompose the importance of honesty and bounded rationality via a series of counterfactuals such as how much communication would occur if no subject cared about honesty ($\tau_{30} = 0$).

The final two colums in Table 6 report model estimates using only sender choice data (column 4) or receiver choice data (column 5). In other words, model (4) maximizes the product of the probabilities in (10) ignoring the second parenthetical term, the probability of the observed receiver choices for each subject.[15] Model (5) lends credence to the somewhat strange result of model (1) that receivers are less strategically sophisticated than senders, despite messages and action choices coming from the same set of subjects. Model (5) estimates imply even less sophistication by receivers: 56% behave unstrategically ($\lambda_0 + \lambda 1 = 0.556$)) and 44% engage in one round of strategic calculation ($\lambda_2 + \lambda_3 = 0.444$). Sender behavior in (4) is similar to estimates in (1) on the other hand (λ_0 estimates are comparable, $\lambda_1 + \lambda_2$ estimates are similar, and λ_3 estimates are alike).

An interesting difference that arises between the sender-only and receiver-only models is in the estimated τ_{30}. In model (4), 58% of players are estimated to be weakly honest, while in model (5), the corresponding estimate is 22% of players. τ_{30} in model (5) is a measure of receiver beliefs about honesty since it is identified off indirect effects of honesty on receiver choices. Hence models (4) and (5) suggest that many subjects are weakly honest but are hesitant to believe that other subjects are honest.

The model pair (2) and (3) and the model pair (4) and (5) can be thought of as alternate models that relax one of the assumptions of my level-k model. (2) and (3) allow λ to vary by treatment instead of forcing strategic sophistication to be constant. Hence the more general specification (2/3) nests (1) as a special case, and a likelihood-ratio test can be used to compare the null hypothesis (1) to the alternate (2/3). The test fails to reject the null ($D = 20.64, \chi^2(7) = 0.170$). A similar likelihood-ratio test comparing the null hypothesis (1) to the alternate model allowing role-specific λ and τ (4/5) also fails to reject the null ($D = 15.94, \chi^2(7) = 0.336$). Comparisons using the

[15] While there are pairs of behaviorally identical types in each case, e.g., $L1^s$ has the same optimal choices as $L2^s$, the distribution λ can be fully estimated because $\beta_1 \neq \beta_2$.

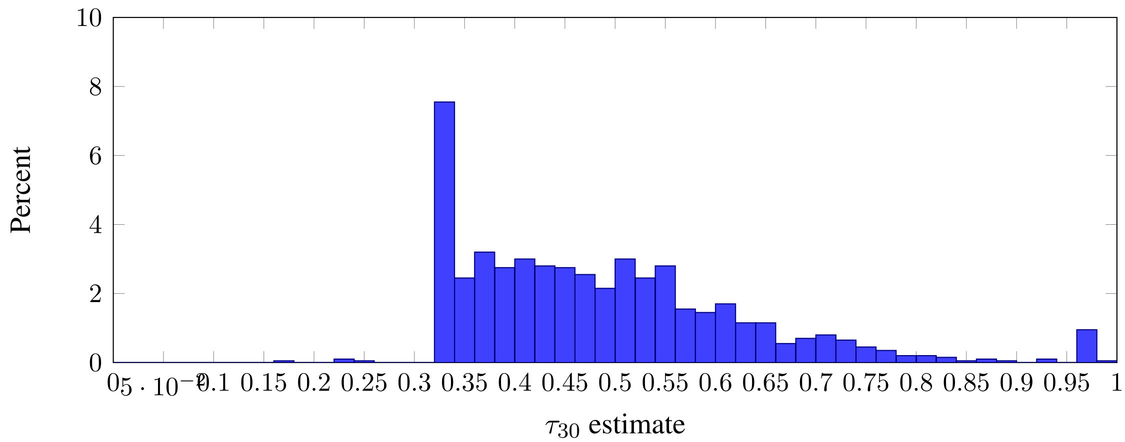

Figure 4: Distribution of τ_{30} estimates for both roles (preferred model – column (1))

Bayesian Information Criterion, which assigns more importance to model parsimony, also strongly support model (1) ($BIC = 4142.9$) over either (2/3) ($BIC = 4172.9$) or (4/5) ($BIC = 4177.6$).[16]

5.2 Honesty

Figure 4 shows the distribution of the fraction of weakly honest players from the bootstrap for my preferred model. There is truncation at $\tau_{30} = 1/3$, suggesting that the effect of $\tau_{30} \geq 1/3$ on receiver beliefs is an important constraint on thes estimates. Aside from the truncation spike, the parameter estimates are unimodal and bell-shaped.

Figures 5 and 6 show the same distribution for the sender-only (column (4)) and receiver-only (column (5)) models. Comparing these to the previous figure, the sender-only distribution of τ_{30} closely resembles the combined roles distribution, while the receiver-only distribution does not resemble it. This is suggestive that while at least half of subjects are weakly honest ($h = 30$), they believe most other subjects are less honest than themselves.

[16]Note that each model effectively has 7 parameters: $\lambda_1, \lambda_2, \lambda_3, \tau_{30}, \beta_1, \beta_2$, and β_3 as $\lambda_0 = 1 - \lambda_1 - \lambda_2 - \lambda_3$ and $\tau_0 = 1 - \tau_{30}$.

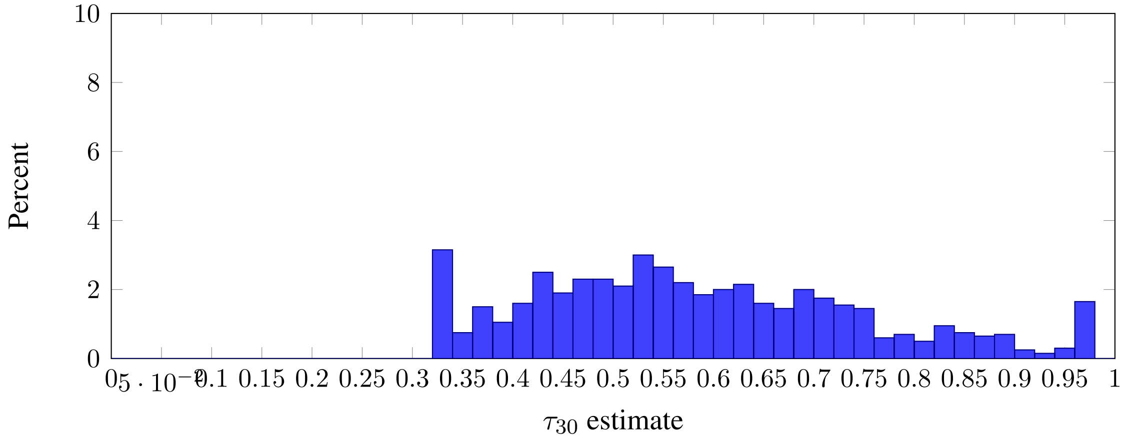

Figure 5: Distribution of τ_{30} estimates for senders (column (4))

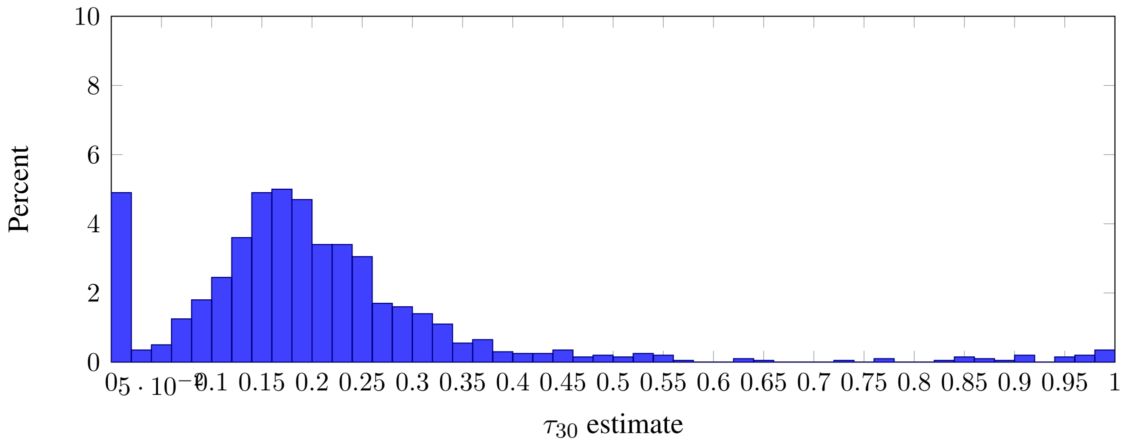

Figure 6: Distribution of τ_{30} estimates for receivers (column (5))

Honest but suspicious behavior by subjects is also consistent with the reduced-form evidence in Section 3, where receivers changed their behavior in the vague treatment less than senders. Nonetheless, my estimates suggest that half of senders are willing to sacrifice what they believe to be a 30 cents higher payoff per round in order to be honest.

Identification: The distribution of honesty is identified in two ways. First, there are several intervals of honesty in which senders with honesty in the interval behave qualitatively in the same way. Any $L1$ or $L2$ sender with $h < 25$ has the same optimal choices as an $h = 0$ sender, who send a mixture of precise and imprecise messages, while for $h \geq 25$ these senders send exclusively imprecise messages in the vague treatment. The other major effect of h on message choice is that $h \geq 57$ senders of all cognitive types send truthful messages exclusively in the precise treatment. There are a few other cutoffs as well, but none have as widespread an effect on message choices as these two.

Second, there are some minimum frequencies of minimum level of honesty that affect receiver behavior; for instance, if more than $1/3$ of senders have $H > 25$, B_2^r and B_2^r beliefs are more trusting of imprecise messages than if facing less frequent weakly honest types.

Because of the logit errors framework, changes in message frequency produced directly by variation in $u_s(s, a, h)$ are continuous in h. For instance, for $k = 1$ or $k = 2$ senders, $\mathrm{E}[u_s(a, s, h) \mid m, B_k^s]$ is continuous in h because B_k^s is not affected by h – these sender types believe that they face credulous senders – so the probability they choose a given m defined in equation (8) is also continuous. However, beliefs B_2^R, B_3^R, and B_3^S depend on the exact optimal choices of these sender types $\sigma_{s,1}$ and $\sigma_{s,2}$. These σ's are discontinuous in h, so that $L2_R$, $L3_R$, and $L3_S$ predictions are generally discontinuous in h. In practice these discontinuities only affect a responses to a few messages, and my numerical minimization routine seems to accomodate them.

Unfortunately, the log-likelihood function is relatively flat with respect to h within each the intervals where beliefs are constant, and in fact over the entire interval $25 < h < 57$ varies only from -2044 to -2047. Figure 8 shows the log-likelihood for each h. Because different levels of honesty have approximately the same explanatory power, pinning down the distribution of honesty exactly is difficult. I illustrate this problem by estimating my model assuming that people are either fully self-interested or have a single degree of honesty h. For each h value, Figure 7 shows the fraction of players estimated to be honest at this level and the 95% confidence intervals. For relatively weak honesty levels between 25 and 40 at least a fifth of playes can be classified as honest with 95% confidence.

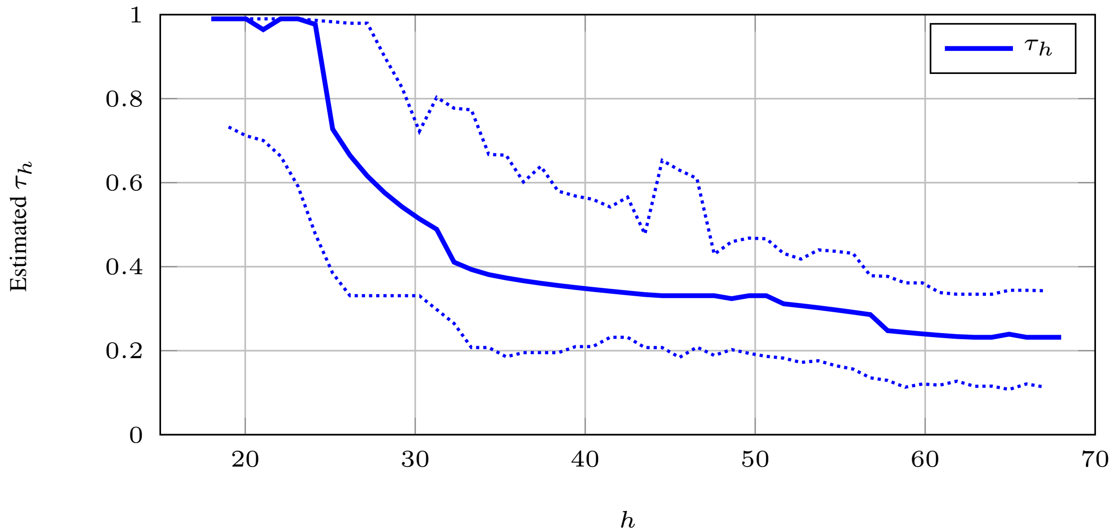

Figure 7: Estimated fraction of players who are weakly honest, τ_h, for varying levels of honesty h, with 95% confidence intervals.

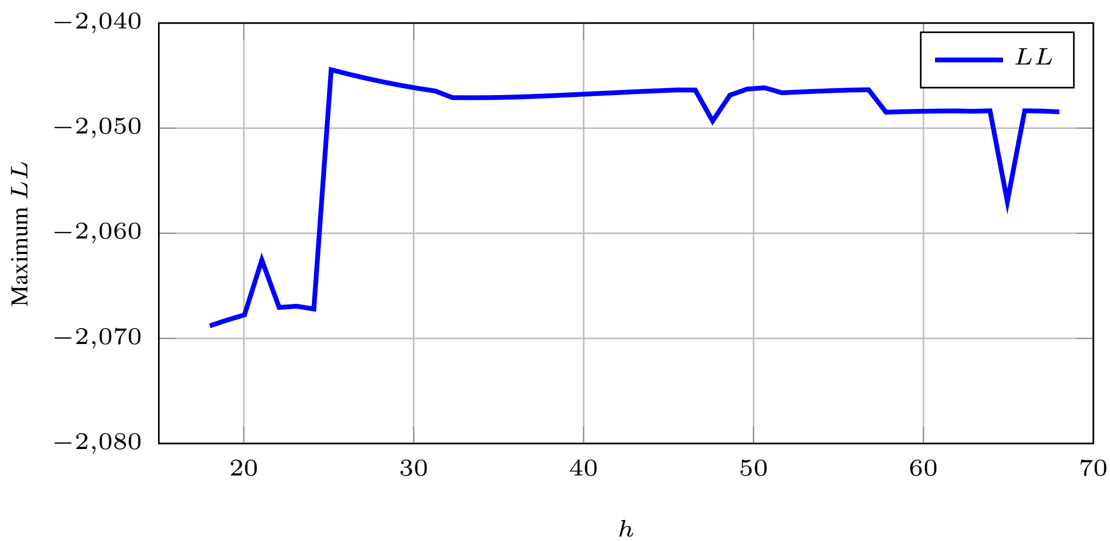

Figure 8: Maximum log-likelihood attained for varying levels of honesty.

Table 7: Parameter Estimates with Multiple Honesty Types or Variant Honesty Type with Multiple Honesty Types or Variant Honesty Type

	1	2	3	4
λ_0	0.0	0.0	0.0	0.0
	[0.0, 0.072]	[0.0, 0.057]	[0.0, 0.164]	[0.0, 0.079]
λ_1	0.117 ***	0.128 ***	0.114 ***	0.132 ***
	[0.026, 0.308]	[0.039, 0.552]	[0.035, 0.581]	[0.082, 0.488]
λ_2	0.437 ***	0.439 ***	0.463 ***	0.589 ***
	[0.22, 0.631]	[0.189, 0.651]	[0.132, 0.752]	[0.328, 0.773]
λ_3	0.446 ***	0.433 ***	0.423 ***	0.279 ***
	[0.272, 0.686]	[0.3, 0.675]	[0.166, 0.793]	[0.09, 0.499]
τ_0	0.547 ***	0.067 ***	0.763 ***	0.1 ***
	[0.035, 0.719]	[0.021, 0.141]	[0.601, 0.838]	[0.1, 0.464]
τ_{30}	0.374	0.331	–	–
	[0.0, 0.941]	[0.0, 0.38]		
τ_{20}	–	0.602 ***	–	0.9 ***
		[0.523, 0.9]		[0.9, 0.9]
τ_{60}	0.078	–	0.237 ***	–
	[0.0, 0.404]		[0.162, 0.399]	
n	1384.0	1384.0	1384.0	1384.0
$logl$	-2045.663	-2043.564	-2048.236	-2068.415
BIC	4149.187	4144.991	4147.1	4187.46

Maximum likelihood estimates of share of Lk players (λ_k) and share of players with utility from honesty h (τ_h). 95% BCa confidence intervals are reported beneath estimates. * denotes significance at 0.1 level, ** significance at 0.05 level, and *** significance at 0.01 level.

This problem only becomes worse when trying to estimate a "richer" model that contains two honest types. Because behavior is fairly similar for any type with $H < 57$, the parameters h_1 and h_2 are colinear and cannot be estimated precisely.[17] Therefore my preferred model contains one honest type, with H chosen to maximize the log-likelihood in Figure 8.

Conclusions about honesty: Table 7 presents several variant assumptions about honesty. Column (1) is a model with three honesty types $H = \{0, 30, 60\}$ and column (2) is another model with three honesty types, $H = \{0, 30, 20\}$. Finally model (3) restricts honest players to $h = 60$ and model (4) restricts honest players to $h = 20$. A likelihood ratio test does not reject my preferred model in favor of either of the nested models.[18] Comparisons of the Bayesian Information Criteria of the non-nested models strongly supports my preferred model over (4) and moderately supports it over (3).

Despite the underlying identification problems, I can conclude from my data that a substantial fraction of subjects have moderate preferences for truth-telling and believe that other subjects do as well. While I am cannot estimate the exact distribution of honesty in my subjects, I am able to show both that is not extremely weak (all subjects having $h < 25$) and that few subjects seem to have strong levels of honesty ($h \geq 57$, the cutoff at which behavior in the precise condition changes).

5.3 Relative Importance of Honesty and Bounded Rationality

The parameter estimates in my models can be used to ask a series of counterfactual questions that allow the relative importance of honesty and bounded rationality for each message space to be

[17]Throwing much more data at the model could solve this problem, but it would require hundreds of subjects. A different experiment would better address this question, for instance by choosing payoffs and signals that make distinguishing between $H = 20$ and $H = 40$ senders easier.

[18]Model (1) in Table 7: $D = 0.963, \chi^2(1) = 0.488$; model (2): $D = 5.162, \chi^2(1) = 0.108$.

gauged.

Table 8 reports the predicted mean distance between messages and states and between actions and messages for several assumptions about λ, τ, and β.[19] The first row reports on the value of each statistic in the data. The expected distance between m and s is 20% less and the expected distance between m and a is 25% less under the vague treatment than under the precise.

The remaining rows reports on the predicted values of these statistics for different parameterizations of my model. The second row shows that my preferred model predicts the informativeness of messages in each treatment well, but less so how receiver choices, predicting only a 6% reduction in the distance between m and a.

The parameter set (2) measures how important honesty is for understanding communication by looking at the predictions of my preferred model if τ_0 is constrained to be 1. Without honest types, less communication is predicted than occurs and the vague treatment effect is predicted to be smaller than it actually is.

The row for parameter set (3) reports a similar exercise for bounded rationality by counterfactually increasing strategic sophistication. λ is set such that all players are the most sophisticated type possible ($\lambda_3 = 1$). Again less communication is predicted to occur than actually occurs, and the predicted effect of the vague treatment is smaller than the true effect. The error in the predicted sizes of the treatment effect are smaller for the increased sophistication model than for the no honesty model however, suggesting that honesty is the more important contributor for this comparative static.[20]

[19] E.g., given observations of actions chosen and messages $(a_1, m_1), \ldots, (a_N, m_N)$, the mean distance between messages and actions is

$$\mathrm{E}[m - a] = \frac{1}{N} \sum_{a \in A} \hat{\Pr}(a \mid m) |a_i - m_i|$$

where $\hat{\Pr}(a \mid m)$ is the predicted conditional probability in equation (9). $\mathrm{E}[m - a]$ is a measure of how trusting action choices are and $\mathrm{E}[m - s]$ measures how informative messages are.

[20] Parameter set (3) predicts a 16% reduction in $\mathrm{E}[m - s]$ and a 9% reduction in $\mathrm{E}[m - a]$, compared to a predicted 9% reduction in $\mathrm{E}[m - s]$ and an 8% *increase* in $\mathrm{E}[m - a]$ for parameter set (2).

Table 8: Mean Deviation of Messages from Truth and Actions from Claimed State in Data, Preferred Model, and Counterfactual Predictions. No honesty assumption reports predicted statistics for $\tau_{30} = 0$ but λ and β from preferred model. Increased sophistication assumption reports predicted statistics for $\lambda_0 = \lambda_1 = \lambda_2 = 0$ and $\lambda_3 = 1$ but τ and β from preferred model.

Parameter Set	*Precise Treatment*		*Vague Treatment*	
	$\mathrm{E}[m - s]$	$\mathrm{E}[m - a]$	$\mathrm{E}[m - s]$	$\mathrm{E}[m - a]$
Actual data	1.52	1.30	1.22	0.98
(1) Predicted	1.47	1.20	1.24	1.13
(2) No honesty	1.55	1.20	1.39	1.29
(3) Increased sophistication	1.58	1.28	1.33	1.21

6 Conclusion

Both preferences for truth-telling and bounded rationality have been advanced to explain anomalous overcommunication in strategic information transmission experiments. My analysis shows that bounded rationality is the major cause of overcommunication if the message space is extremely restricted relative to spoken language by only allowing precise messages. However, when the message space is enriched slightly by adding explicitly imprecise messages, preferences for truth-telling interact with bounded rationality to cause even greater overcommunication. Since natural language is much richer and capable of nuance than the more expressive message space I consider, these findings suggest that the external validity of strategic information transmission experiments can be improved by allowing a wider variety of messages.

References

Agranov, Marina and Andrew Schotter, "Ignorance Is Bliss: An Experimental Study of the Use of Ambiguity and Vagueness in the Coordination Games with Asymmetric Payoffs," 2012, *4* (2), 77–103.

Blume, Andreas and Oliver Board, "Intentional Vagueness," *Erkenntnis*, 2014, *79*, 855–899.

_ , Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle, "Evolution of Communication with Partial Common Interest," *Games and Economic Behavior*, 2001, *37* (1), 79–120.

_ , Oliver Board, and Kohei Kawamura, "Noisy talk," 2007, *2*, 395–440.

Cai, Hongbin and Joseph Tao-Yi Wang, "Overcommunication in Strategic Information Transmission Games," *Games and Economic Behavior*, 2006, *56* (1), 7–36.

Chen, Ying, "Perturbed Communication Games with Honest Senders and Naive Receivers," *Journal of Economic Theory*, 2011, *146*, 401–424.

Crawford, Vincent P., "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentations of Intentions," *American Economic Review*, March 2003, *93*, 133–149.

_ and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), 1431–52.

_ , Miguel Costa-Gomez, and Nagore Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 2013, *51* (1), 5–62.

DiCiccio, Thomas J. and Bradley Efron, "Bootstrap Confidence Intervals," *Statistical Science*, 1996, *11*, 189–228.

Duffy, John, Tyson Hartwig, and John Smith, "Costly and discrete communication: an experimental investigation," *Theory and Decision*, 2014, *76* (3), 395–417.

Efron, Bradley, "Bootstrap Methods: Another Look at the Jackknife," *Annals of Statistics*, 1979, *7*, 1–26.

Gneezy, Uri, "Deception: The Role of Consequences," *American Economic Review*, 2005, *95* (1), 384–94.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, lies, and costly talk," *Journal of Economic Theory*, 2007, *134* (1), 93–116.

Myerson, Roger B., *Game Theory: Analysis of Conflict*, Cambridge: Harvard University Press, 1997.

Nagel, Rosemarie, "Unravelling in Guessing Games: An Experimental Study," *American Economic Review*, 12 1995, *85* (5), 1313–1326.

Sanchez-Pages, Santiago and Marc Vorsatz, "An experimental study of truth-telling in a sender-receiver game," *Games and Economic Behavior*, October 2007, *61* (1), 86–112.

_ **and** _ , "Enjoy the silence: an experiment on truth-telling," *Experimental Economics*, June 2009, *12* (2), 220–241.

Serra-Garcia, Marta, Eric van Damme, and Jan Potters, "Hiding an inconvenient truth: Lies and vagueness," *Games and Economic Behavior*, 2011, *73*, 244–261.

Stahl, Dale O. and Paul W. Wilson, "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, December 1994, *25* (3), 309–327.

Sutter, Mattias, "Deception Through Telling the Truth?! Experimental Evidence from Individuals and Teams," *Economic Journal*, 2009, *119* (1), 47–60.

Wang, Joseph Tao-Yi, Michael Spezio, and Colin F. Camerer, "Pinocchio's Pupil: Using Eye-tracking and Pupil Dilation to Understand Truth-telling and Deception in Games," *American Economic Review*, 2010, *100* (3), 984–1007.

A Instructions

Communication Experiment Instructions

Thank you for participating in the experiment today! As a courtesy to me and the other participants, I ask you to observe a few ground rules:

- Focus your attention on the experiment rather than reading, texting, or other activities.
- Do not talk while the experiment is ongoing.
- Do not use other programs on the computers.
- Do not look at other people's computer screens.

Questions are welcome at any time. If you have a question, raise your hand and an experimenter will come to help.

Timing

Here is what will happen in this session:

1. You read the consent form, ask any questions you have, and return a signed copy to me.
2. You read the instructions, ask questions, and answer comprehension questions on next page. I check the comprehension questions when you have finished them.
3. We run the first experiment (described below).
4. We run the second experiment (similar to first, will be described when we get to this stage).
5. We run the third experiment (similar to first, will be described when we get to this stage).
6. You fill out payment / demographic information on your computer.
7. Your computer displays your earnings, you fill out a receipt, and I pay you as you leave.

Experiment 1 instructions

There are 16 rounds in this experiment. Every round you will be matched with a random participant, who will be a different person every round. One person in each pair will be an "A" and the other person will be a "B". You will alternate between being an A and a B.

Each round there will a number **s** generated for each pair. **s** is equally likely to be 1, 2, 3, 4, 5. Each A learns what the pair's **s** is, but B is not shown **s**.

The A member then chooses a message to send to B. The following messages can be chosen:

- "s is 1."
- "s is 2."
- "s is 3."
- "s is 4."
- "s is 5."

After A has finished choosing the message, B gets to read it. B then chooses a number **w**. **w** can be 1, 2, 3, 4, or 5.

A and B both earn payoffs each round based on **s** and **w**. For B, the highest payoff is when she chooses the same **w** as **s**. B's are best off choosing the average number they think **s** is. For example, if as a B you think **s** is equally likely to 1, 2, or 3, you are best off choosing w equal to 2. For A, the highest payoff is when B chooses **w** equal to **s** + 2. The tables below show A and B's exact payoffs.

B payoffs

S	w is 1	w is 2	w is 3	w is 4	w is 5
s is 1	100	75	43	7	-32
s is 2	75	100	75	43	7
s is 3	43	75	100	75	43
s is 4	7	43	75	100	75
s is 5	-32	7	43	75	100

A payoffs

S	w is 1	w is 2	w is 3	w is 4	w is 5
s is 1	43	75	100	75	43
s is 2	7	43	75	100	75
s is 3	-32	7	43	75	100
s is 4	-72	-32	7	43	75
s is 5	-115	-72	-32	7	43

These payoffs are in experimental units. In addition, you receive 200 units at the beginning of each experiment. At the end of the session, units will be converted to dollars: each unit is worth 1 cent.

Examples

s=2 and **w**=4. Then B receives 43 and A receives 100.

s=5 and **w**=3. Then B receives 43 and A loses 32.

Please complete these questions once you have finished reading the earlier pages. When you have finished them, raise your hand. These questions are to make sure that everyone understands how the first experiment works. They do not affect your earnings.

Comprehension questions

If **s** is 4, then

B earns the most if she chooses **w** equal to ________

A earns the most if B chooses **w** equal to ________

If **s** is 2 and B chooses **w** to be 5, then

B receives ________

A receives ________

(More difficult questions)

If the highest-paying **w** for A is 2, what is the highest-paying **w** for B? ________

If the highest-paying **w** for B is 4, what is the highest-paying **w** for A? ________

Experiment 2 instructions

There are 12 rounds instead of 16 in this experiment.

This experiment has the same basic structure as the previous one. A learns **s** and sends a message to B, who then decides on **w**. Earnings each round are calculated in exactly the same way and you will receive 200 units at the beginning.

The only difference is that now A's will choose from a different fixed set of messages. Now the message possibilities are:

- "s is 1."
- "s is 1, 2, or 3."
- "s is 2."
- "s is 2, 3, or 4."
- "s is 3."
- "s is 3, 4, or 5."
- "s is 4."
- "s is 5."

Experiment 3 instructions

There are 12 rounds in this experiment.

This experiment has the same basic structure as the previous one. A learns **s** and sends a message to B, who then decides on **w**. Earnings each round are calculated in exactly the same way and you will receive 200 units at the beginning.

The only difference is that now A's will sometimes not be told **s** exactly. Half of the time, A's will be told **s** precisely, but the other half of the time A's will only be told one of the following

- s is 1, 2, or 3.
- s is 2, 3, or 4.
- s is 3, 4, or 5.

These reports are accurate and when A is told that s is in one of these sets of numbers, each number is equally likely. If A learns that **s** is 2, 3, or 4, for example, that means there is a 1-in-3 chance that **s** is 1, a 1-in-3 chance that **s** is 3, and a 1-in-3 chance that **s** is 4.

The messages A can choose from are the exactly the same as in experiment two:

- "s is 1."
- "s is 1, 2, or 3."
- "s is 2."
- "s is 2, 3, or 4."
- "s is 3."
- "s is 3, 4, or 5."
- "s is 4."
- "s is 5."

Communication Experiment Instructions

Thank you for participating in the experiment today! As a courtesy to me and the other participants, I ask you to observe a few ground rules:

- Focus your attention on the experiment rather than reading, texting, or other activities.
- Do not talk while the experiment is ongoing.
- Do not use other programs on the computers.
- Do not look at other people's computer screens.

Questions are welcome at any time. If you have a question, raise your hand and an experimenter will come to help.

Timing

Here is what will happen in this session:

1. You read the consent form, ask any questions you have, and return a signed copy to me.
2. You read the instructions, ask questions, and answer comprehension questions on next page. I check the comprehension questions when you have finished them.
3. We run the first experiment (described below).
4. We run the second experiment (similar to first, will be described when we get to this stage).
5. We run the third experiment (similar to first, will be described when we get to this stage).
6. You fill out payment / demographic information on your computer.
7. Your computer displays your earnings, you fill out a receipt, and I pay you as you leave.

Experiment 1 instructions

There are 16 rounds in this experiment. Every round you will be matched with a random participant, who will be a different person every round. One person in each pair will be an "A" and the other person will be a "B". You will alternate between being an A and a B.

Each round there will a number **s** generated for each pair. **s** is equally likely to be 1, 2, 3, 4, 5. Each A learns what the pair's **s** is, but B is not shown **s**.

The A member then chooses a message to send to B. The following messages can be chosen:

- "s is 1."
- "s is 2."
- "s is 3."
- "s is 4."
- "s is 5."

After A has finished choosing the message, B gets to read it. B then chooses a number **w**. **w** can be 1, 2, 3, 4, or 5.

A and B both earn payoffs each round based on **s** and **w**. For B, the highest payoff is when she chooses the same **w** as **s**. B's are best off choosing the average number they think **s** is. For example, if as a B you think **s** is equally likely to 1, 2, or 3, you are best off choosing w equal to 2. For A, the highest payoff is when B chooses **w** equal to **s** + 2. The tables below show A and B's exact payoffs.

B payoffs

S	w is 1	w is 2	w is 3	w is 4	w is 5
s is 1	100	75	43	7	-32
s is 2	75	100	75	43	7
s is 3	43	75	100	75	43
s is 4	7	43	75	100	75
s is 5	-32	7	43	75	100

A payoffs

S	w is 1	w is 2	w is 3	w is 4	w is 5
s is 1	43	75	100	75	43
s is 2	7	43	75	100	75
s is 3	-32	7	43	75	100
s is 4	-72	-32	7	43	75
s is 5	-115	-72	-32	7	43

These payoffs are in experimental units. In addition, you receive 200 units at the beginning of each experiment. At the end of the session, units will be converted to dollars: each unit is worth 1 cent.

Examples

s=2 and **w**=4. Then B receives 43 and A receives 100.

s=5 and **w**=3. Then B receives 43 and A loses 32.

Please complete these questions once you have finished reading the earlier pages. When you have finished them, raise your hand. These questions are to make sure that everyone understands how the first experiment works. They do not affect your earnings.

Comprehension questions

If **s** is 4, then

B earns the most if she chooses **w** equal to ________

A earns the most if B chooses **w** equal to ________

If **s** is 2 and B chooses **w** to be 5, then

B receives ________

A receives ________

(More difficult questions)

If the highest-paying **w** for A is 2, what is the highest-paying **w** for B? ________

If the highest-paying **w** for B is 4, what is the highest-paying **w** for A? ________

Experiment 2 instructions

There are 12 rounds instead of 16 in this experiment.

This experiment has the same basic structure as the previous one. A learns **s** and sends a message to B, who then decides on **w**. Earnings each round are calculated in exactly the same way and you will receive 200 units at the beginning.

The only difference is that now A's will choose from a different fixed set of messages. Now the message possibilities are:

- "s is 1."
- "s is 1, 2, or 3."
- "s is 2."
- "s is 2, 3, or 4."
- "s is 3."
- "s is 3, 4, or 5."
- "s is 4."
- "s is 5."

Experiment 3 instructions

There are 12 rounds in this experiment.

This experiment has the same basic structure as the previous one. A learns **s** and sends a message to B, who then decides on **w**. Earnings each round are calculated in exactly the same way and you will receive 200 units at the beginning.

The only difference is that now A's will sometimes not be told **s** exactly. Half of the time, A's will be told **s** precisely, but the other half of the time A's will only be told one of the following

- s is 1, 2, or 3.
- s is 2, 3, or 4.
- s is 3, 4, or 5.

These reports are accurate and when A is told that s is in one of these sets of numbers, each number is equally likely. If A learns that **s** is 2, 3, or 4, for example, that means there is a 1-in-3 chance that **s** is 1, a 1-in-3 chance that **s** is 3, and a 1-in-3 chance that **s** is 4.

The messages A can choose from are the exactly the same as in experiment two:

- "s is 1."
- "s is 1, 2, or 3."
- "s is 2."
- "s is 2, 3, or 4."
- "s is 3."
- "s is 3, 4, or 5."
- "s is 4."
- "s is 5."

B Data

Three sessions were conducted in November 2011 with a total of 42 participants. Subjects were paid a $6 dollar show-up fee plus $0.01 per experimental unit earned. Average earnings were $24.12, with a minimum of $16 and a maximum of $30.

	Signal								
	1	**2**	**3**	**4**	**5**	**1, 2, or 3**	**2, 3, or 4**	**3, 4, or 5**	**N**
	No.	No.	No.	No.	No.	No.	No.	No.	No.
Message (noise condition)									
1	2	0	0	0	0	2	0	0	4
2	0	0	1	0	1	0	1	1	4
3	5	3	2	2	1	0	2	3	18
4	2	4	1	5	0	1	0	6	19
5	4	9	9	11	12	4	9	18	76
1, 2, or 3	1	3	1	1	0	6	2	2	16
2, 3, or 4	7	0	1	2	2	5	5	2	24
3, 4, or 5	5	3	8	4	5	13	9	12	59
N	26	22	23	25	21	31	28	44	220
Message (vague condition)									
1	2	0	1	1	2	–	–	–	6
2	3	5	1	1	0	–	–	–	10
3	10	2	2	0	2	–	–	–	16
4	0	2	4	3	2	–	–	–	11
5	8	4	12	17	27	–	–	–	68
1, 2, or 3	5	4	2	1	1	–	–	–	13
2, 3, or 4	8	2	3	4	1	–	–	–	18
3, 4, or 5	3	17	13	16	13	–	–	–	62
N	39	36	38	43	48	–	–	–	204
Message (precise condition)									
1	9	5	0	0	2	–	–	–	16
2	8	12	4	4	2	–	–	–	30
3	16	5	7	2	7	–	–	–	37
4	10	15	13	5	4	–	–	–	47
5	17	23	35	31	32	–	–	–	138
N	60	60	59	42	47	–	–	–	268

Source: novembexp.dta

Table 9: Sender Choices

	Action					
	1	**2**	**3**	**4**	**5**	**N**
	No.	No.	No.	No.	No.	No.
Message (noise condition)						
1	3	0	1	0	0	4
2	1	2	0	1	0	4
3	3	6	6	3	0	18
4	0	8	6	3	2	19
5	2	6	20	27	21	76
1, 2, or 3	3	8	5	0	0	16
2, 3, or 4	2	5	16	1	0	24
3, 4, or 5	0	13	29	15	2	59
N	14	48	83	50	25	220
Message (vague condition)						
1	4	1	0	1	0	6
2	1	3	4	2	0	10
3	6	3	5	2	0	16
4	0	4	2	4	1	11
5	3	2	21	21	21	68
1, 2, or 3	6	2	4	1	0	13
2, 3, or 4	2	1	14	1	0	18
3, 4, or 5	2	17	20	19	4	62
N	24	33	70	51	26	204
Message (precise condition)						
1	5	7	4	0	0	16
2	5	16	4	5	0	30
3	15	7	7	6	2	37
4	0	20	14	10	3	47
5	8	4	64	39	23	138
N	33	54	93	60	28	268

Source: novembexp.dta

Table 10: Receiver Choices