Competition and Exaggeration in Experimental Cheap Talk Games *

William Minozzi† Jonathan Woon‡

May 2, 2018

Abstract

We experimentally investigate the effect of a second opinion on information transmission. Our design varies the number of senders as well as the alignment and magnitude of senders' biases in a sequential, cheap-talk, sender-receiver game. We find that decision makers do no better when a second opinion is available, irrespective of the alignment or competition between advisers, than when they receive a single opinion. Despite the fact that messaging behavior differs across experimental conditions, receivers successfully extract the same amount of information—an amount greater than what is theoretically predicted. These findings are consistent with senders using a simple strategy of naïve exaggeration, with receivers correctly recognizing this and adjusting their behavior accordingly.

Keywords: Strategic information transmission; Sender-receiver games; Multiple senders; Laboratory experiment

JEL Classifications: C72, C92, D82, D83

---

*We thank Dan Butler and Alistair Wilson for their valuable comments and suggestions on early versions of the manuscript. Previous versions of this paper were presented at the University of Warwick Political Economy Seminar, the London School of Economics Political Economy Seminar, the University of East Anglia Experimental Economics Seminar, Texas A & M's Political Economy Seminar, the Economic Science Association North American Conference, the Southern Political Science Association meeting, the Behavioral Models of Politics Conference, and the Midwest Political Science Association Meeting.

†Associate Professor, Department of Political Science, 2137 Derby Hall, Ohio State University, Columbus, OH 43210. minozzi.1@osu.edu

‡Corresponding author. Associate Professor, Department of Political Science, Department of Economics (secondary), and Pittsburgh Experimental Economics Laboratory, 4443 Wesley W. Posvar Hall, University of Pittsburgh, Pittsburgh, PA 15260. woon@pitt.edu

1 Introduction

Many decisions in business, politics, and society depend on competitive communication from the informed to the empowered. Conventional wisdom holds that consumers, patients, and investors should seek a "second opinion" when making important decisions. For similar reasons, managers and executives may seek the opinions of subordinates from different divisions whose agendas may differ from each other or from those of the entire organization. Litigation entails adversarial competition between advocates and witnesses seeking to persuade juries. Legislative committees call multiple experts and witnesses to testify in public hearings. Indeed, competitive communication pervades democratic politics, from the interactions between lobbyists and elected officials, to the relationships among political parties, candidates, and voters.

The conventional wisdom holds that competition supports more informed decisionmaking, and game theoretic analysis provides formal support for this idea. The seminal theory of Crawford and Sobel (1982) demonstrates that strategic incentives limit the information that can be transmitted from a single adviser to a decisionmaker. Building on this theory, Gilligan and Krehbiel (1989) and Krishna and Morgan (2001*a*,*b*) show that merely adding a second opinion can dramatically increase the amount of transmittable information in equilibrium, even up to full revelation under some conditions. In related strategic environments, the analyses of Battaglini (2002), Milgrom and Roberts (1986), and Minozzi (2011) also demonstrate that much more information can be transmitted with multiple advisers than with a single adviser.[1] Following the rigorous analysis offered by game theorists, decisionmakers ought to consult multiple experts with differing preferences to maximize the information content of their actions.

However, a growing body of experimental evidence complicates the equilibrium understanding of strategic information transmission, thus raising questions about the empirical value of a second opinion. That value depends critically on comparing the information that is transmittable

[1] In cheap talk settings, the analytic focus is on incentive compatibility in disclosure and the credibility of communication. Relatedly, Gentzkow and Kamenica (2017) consider the problem of information acquisition, assuming any information gathered is truthfully revealed. In their setting, they show that the effect of competition is generally ambiguous.

by a single expert and by multiple experts. Equilibrium analysis suggests the former is quite low. Yet experimental evidence demonstrates that senders tend to overcommunicate relative to equilibrium benchmarks (see, e.g., Blume et al., 1998; Cai and Wang, 2006). Given that single experts transmit so much more information than would be expected based on equilibrium analysis, the relative empirical value of a second opinion remains unclear.

In this paper, we measure the value of a second opinion by experimentally examining behavior in competitive and non-competitive strategic information transmission games. Our design entails three configurations of senders: one single sender, two senders with aligned preferences, and two senders with opposed preferences. When there are two senders, messages are sent sequentially, so that the more extreme sender selects a (public) message first, and the more moderate sender chooses their message second. The sequential nature of the game mirrors a consumer seeking a second opinion. In every configuration, the more moderate sender has precisely the same bias across each of the three settings, and the magnitude of the extremist's bias is the same in the conditions with aligned and opposed senders. Thus, in our experiment, we fix the preferences of players within each set of sessions, varying only the presence of a second, relatively extreme sender and the direction of that sender's bias. This setup effectively tests the theory put forward by Krishna and Morgan (2001*b*), who show that adding a second sender in a game with sequential messages can increase information transmission only when the senders' biases are opposed but not when they are aligned.

Surprisingly, we find that the value of a second opinion is negligible, regardless of whether the two senders are aligned or opposed. This finding matches equilibrium predictions in the former case but not the latter. The reason, however, is that senders consistently reveal more information than predicted by equilibrium theory across all three conditions. Sender behavior can be accurately characterized as a form of naïve exaggeration, in which messages are simply additive shifts away from the underlying state value. Interestingly, information transmission is similar across cases even though senders and receivers use different naïve exaggeration strategies in different conditions. Moderate senders temper their exaggeration when they respond to an aligned extremist compared to when they act alone, but exaggerate more in response to opposed extremists. Extremists generally exaggerate more (and thus send less information) than moderates. And receivers rely less on moderates in the presence of an aligned extremist, even less when the extremist is opposed, and yet consistently extract the same amount of information across conditions.

2 Related Literature

Theoretically, the direct antecedents to our experiment are the seminal cheap talk model of Crawford and Sobel (1982) and the extension to two senders by Krishna and Morgan (2001*b*). In these models, senders are perfectly informed about a one-dimensional state variable, and messages are costless and unverifiable. The focus of the theory is understanding how strategic incentives affect the amount of information that can be credibly communicated to the receiver when there is partial common interest between the players. Crawford and Sobel (1982) prove that only coarse (categorical) information can be transmitted in equilibrium and that the maximum amount of information is decreasing in the degree of preference divergence between the sender and receiver.

Krishna and Morgan (2001*b*) extend the model to the case of two senders who communicate sequentially to the receiver. The interaction between senders who have different preferences can increase information transmission only if their preferences diverge in different directions from the receiver. The equilibrium construction relies on the fact that the second sender's message can be used to discipline the first sender's message. The effect of adding the second sender can be dramatic: even when the preferences of the more moderate sender are far enough away to prevent any information transmission in a single sender environment, the receiver can learn the exact state for many possible states in the two sender environment (i.e., even when the only equilibrium is babbling in the former, there exists a semi-revealing equilibrium in the latter).

The literature includes a variety of other models in which more information is transmitted with two senders than with one. Gilligan and Krehbiel (1989) explore the analogous model with simultaneous messages and construct a semi-revealing equilibrium in which receivers learn the true state only when it is sufficiently extreme, and Krishna and Morgan (2001*a*) construct a fully revealing equilibrium in the same environment.[2] Minozzi (2011) shows that with additional uncertainty about senders' preferences, every state is revealed with positive probability but that a sender can confuse the receiver by jamming the opposing sender's truthful messages. In the multidimensional case, Battaglini (2002) shows that full revelation is generally possible even if preference divergence is large, as long as each sender shares a common interest with the receiver on a linearly independent dimension of the state space.

Information transmission also increases under different conditions of sender information and available messages. When senders have imperfect information about a binary state, Austen-Smith (1993) shows that multiple senders are informationally superior to one, and that sequential communication is more informative than simultaneous communication. In models where senders have partial or different information, Wolinsky (2002) shows that information can be elicited without the receiver's commitment if senders communicate with each other before communicating with the receiver; McGee and Yang (2013) show that this interaction exhibits strategic complementarity. When information is verifiable, Milgrom and Roberts (1986) demonstrate that competition compensates for a lack of strategic sophistication only if senders' interests are sufficiently opposed or if the receiver cares about senders' welfare. In an interest group environment, Battaglini and Benabou (2003) show that information decreases with the amount of active lobbying and expected preference conflict: lobbying is more frequent but less truthful when expected preference conflict is low, and less frequent but more truthful when expected conflict is high.

While the theoretical literature has looked at a variety of different environments, the experimental literature has examined only a few. Most studies focus on communication from a single sender in the basic one dimensional Crawford and Sobel (1982) framework. These studies find support for the comparative static prediction that information is decreasing in preference diver-

[2] Krehbiel (2001) argues that the construction relies on implausible beliefs, while Battaglini (2002) demonstrates that beliefs in fully revealing equilibria with a one dimensional state space are not robust to the possibility of small mistakes. Similarly, Baron and Meirowitz (2006) demonstrate the equivalence between equilibria in signaling games that rely on flexible specifications of beliefs and screening games with commitment.

gence, but also that senders overcommunicate, revealing more information than even the most informative equilibrium predicts (Blume et al., 1998, 2001; Dickhaut, McCabe and Mukherji, 1995). Cai and Wang (2006) explain overcommunication with a level-k theory of limited strategic sophistication, as do Kawagoe and Takizawa (2009) and Wang, Spezio and Camerer (2010).[3] In contrast to our environment, most one-dimensional, single-sender experiments also feature small state spaces, so it is not entirely clear whether overcommunication results from subjects engaging in level-k type thinking or from the application of a simple behavioral rule of naïve exaggeration.

The multi-sender, multi-dimensional framework of Battaglini (2002) has also received recent attention. Lai, Lim and Wang (2011) implement a discrete 2×2 multidimensional state space to explore the robustness properties of the fully revealing equilibrium (see Ambrus and Takahashi, 2008). They find that competition improves information transmission, as receivers identify the true state more often when there are two senders than one. By manipulating the size of the message space to influence out-of-equilibrium beliefs, they find faster convergence when the fully revealing equilibrium is robust. Vespa and Wilson (2016) implement a much richer environment, using circles to represent state spaces without boundaries. Their results suggest that while senders provide enough information for receivers to infer the state, receivers use a behavioral strategy that is responsive to the structure of senders' biases, in contrast to theoretical predictions. They also find that senders engage in greater exaggeration when biases are opposed than when they are aligned, but they do not find differences in overall payoffs between treatments.

In an environment with opposed senders, preference uncertainty, and a one-dimensional state space, Minozzi and Woon (2013, 2016) find a degree of information transmission similar to that predicted by the jamming equilibrium of Minozzi (2011).[4] Yet, their results also suggest that

[3] Gneezy (2005) uses a different cheap talk game to study deception. In his game, there is no hidden state variable. Instead, receivers are provided no information whatsoever about the possible payoffs and therefore face an ambiguous choice rather than a well-defined incomplete information sender-receiver game. He shows that lying is sensitive to distributional concerns, but Hurkens and Kartik (2009) also show that the data are consistent with a lying aversion explanation (i.e., there are some people who never lie). In a game where receivers can engage in costly punishment, Sánchez-Pagés and Vorsatz (2007, 2009) also argue that excessive truth-telling is associated with lying aversion.

[4] An experiment with a one-dimensional state space but without preference uncertainty is Battaglini et al. (2016), who experimentally test the predictions of Gilligan and Krehbiel (1989), comparing the effect of open versus closed rules in an environment with two competing committees offering simultaneous bills. They find support for the theory's qualitative predictions that information transmission is greater with the closed rule than the open rule and

senders exaggerate their messages in the direction of their biases in a way that is consistent with overcommunication in single sender environments. However, the amount of exaggeration is understated and exhibits a weak relationship with variation in the biases, which suggests behavior is more naïve than it is strategically sophisticated (i.e., than level-k reasoning). Receivers do well because averaging the senders' messages is optimal given such naïve exaggeration.

3 Theory and Hypotheses

Our experimental design compares behavior in the context of three different conditions defined by the set of senders and the configuration of their preferences: a *Single* sender, two *Aligned* senders, and two *Opposed* senders. In all three cases, play is sequential. Thus, in the conditions with two senders, the sender who moves second acts as the purveyor of a second opinion.

The Single sender condition replicates the seminal Crawford and Sobel (1982, hereafter CS) sender-receiver game. Here, the sender first observes the hidden state variable, which we refer to as the target t and which has a uniform distribution on X, the set of integers from -100 and 100. We refer to this sender as the "moderate" to facilitate comparison with the two-sender game. The sender then chooses a cheap talk message $m_M \in X$ to send. Finally, the receiver observes the messages and selects an action $a \in X$.

The Aligned and Opposed conditions are two sender games that implement the sequential setup of Krishna and Morgan (2001*b*, hereafter KM) by extending the Single sender environment to include a third player, who we refer to as the "extremist."[5] The extremist moves first, observing the target t and choosing a message $m_E \in X$. The moderate moves second, observing both t and m_E and choosing $m_M \in X$. The receiver observes both m_E and m_M and selects $a \in X$.

Payoffs are given by the linear loss between the action a and each player's "ideal action," the target that has been additively shifted. Each player has a different shift parameter $s_i \in \mathbb{R}$. The

that transmission is decreasing in the extremity of committees' preferences, but their data also suggest that subjects overcommunicate.

[5] Our usage of the term "extremist" differs somewhat from KM. We use the terms "moderate" and "extremist" to identify the relative preferences of the senders, whereas KM use the terms to refer to preferences in an absolute sense. Both of our senders are "moderates" in KM's usage.

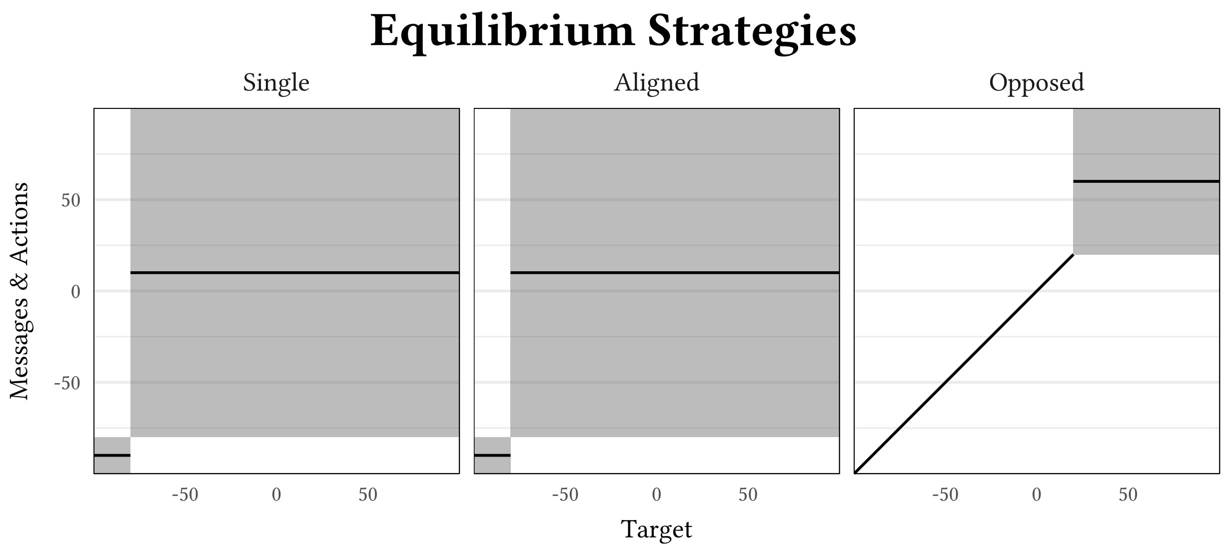

Figure 1: *Most Informative Equilibrium in Each Condition.* Each panel depicts the equilibrium predictions for the relationships between the target (i.e., the hidden state information), messages, and actions. Gray boxes represent partition message strategies, and black lines depict equilibrium actions.

receiver's shift is set to $s_R = 0$. To facilitate comparison, we set the moderate sender's shift to be $s_M > 0$ in all three conditions. Similarly, we fix the magnitude of the extremist sender's shifts at $|s_E|$ in each of the two-sender conditions. In the Aligned condition, both senders' shifts are on the same side of the receiver's shift $(s_E > s_M)$, and in the Opposed condition they are on opposite sides, with the extremist having shift $-s_E$. Formally, payoffs are given by $U(a; t, s) = -|t+s-a|$. The prior distribution of the target, values of the shift parameters, sequence of play, strategy sets, and all other aspects of the game are common knowledge.

In any informative equilibrium of the Single sender game, CS show there will be a partition of the state space, and the sender will only be able to convey coarse information about which subset in the partition contains the target. The most informative equilibrium predictions for each condition are shown in Figure 1, with the Single sender game in the leftmost panel.[6] The breakpoints between partitions are the points at which the sender is known to be indifferent

[6]See the Appendix for precise statements of the equilibria shown in Figure 1. We focus here on describing the key properties of the CS and KM equilibria rather than any specific construction given the multiplicity of equilibria with the same properties. For example, there is an uninformative equilibrium in which the sender randomizes as well as one in which the sender chooses a single precise message for all possible states. Nevertheless, the relationship between the unobserved states and the receiver's action is the same in all uninformative equilibria.

between the actions that the receiver would take if he knew the target were a member of either of the two subintervals that meet at that point. Importantly, because we do not vary the preferences of the moderate sender in our experiments, we focus on the single most informative equilibrium across all instances of our single sender experiment.[7]

For the two-sender games, we follow KM's equilibrium analysis of perfect Bayesian equilibria. Turning to the Aligned condition (which they label "Experts with Like Biases"), KM show that the addition of an aligned extremist cannot alter the information transmitted in equilibrium.[8] In fact, all players would be worse off in any equilibrium in which the receiver listens to both senders. The reason for this is the interaction between the senders in conjunction with the equilibrium requirement that one of the senders be indifferent at any breakpoint of the partition equilibrium.[9] If the receiver listens to the second, more extreme sender, the breakpoints become distorted in such a way that the largest partition becomes even larger, resulting in the loss of information. The best that the receiver can do is therefore to ignore the more extreme sender.

A very different outcome results in the Opposed condition (KM's "Experts with Opposing Biases"). Here, the interaction between the two opposed senders has a disciplining effect. KM prove that as long as the moderate is not too extreme, there is always an informationally superior equilibrium compared to when the moderate is the only sender.[10] The most informative equilibrium is a semi-revealing equilibrium in which every state is revealed except for states where the moderate sender would prefer the most extreme outcome.[11] In this equilibrium, the extremist reveals the true state and the moderate essentially confirms that the message is accurate. If the extremist attempts to lie, then the moderate disagrees and makes a recommendation that would lead to a worse outcome for the extremist, thereby punishing the lie. Information is therefore revealed when the senders' messages agree; when messages disagree, the receiver pays more attention to

[7] See Chen, Kartik and Sobel (2008) for a refinement that justifies informativeness as a selection criterion.

[8] KM's Proposition 2 states "When experts have like biases, there is *never* a monotonic PBE with both experts that is informationally superior to the most informative PBE with a single expert." Furthermore, they note that they "know of no instance where admitting a nonmonotonic equilibrium reverses the welfare comparisons" (759).

[9] See KM's Examples 1 and 2.

[10] Specifically, KM's Proposition 3 states "When experts have opposing biases and at least one is a moderate, there is *always* a PBE with both experts that is informationally superior to the most informative PBE with a single expert."

[11] See KM's Figure 2 and the corresponding equilibrium construction.

the message that is less self-serving.

Equilibrium analysis therefore generates a set of research hypotheses regarding the degree of information transmission within and across our experimental conditions. Since sender-receiver games have a multiplicity of equilibria, we focus on overall information transmission as the link between theory and experiment rather than the behavioral strategies of senders or receivers. Following CS, equilibrium analysis predicts that the advice of a single expert will be of limited value. Following KM, equilibrium analysis further predicts that a second opinion will be useful only if the senders want to pull the decision maker in opposite directions. We summarize our specific predictions in the following three hypotheses.

- H1 In the Single sender condition, the information transmitted will be limited and follow a partition structure.
- H2 Information transmitted in the Aligned condition will not be greater than in the Single sender condition.
- H3 Information transmitted in the Opposed condition will be greater than in either the Single sender or the Aligned conditions.

4 Experimental Design and Procedures

Our experimental sessions were conducted using z-Tree (Fischbacher, 2007) at the Pittsburgh Experimental Economics Laboratory. Subjects were recruited through the lab's centralized database, and most subjects were undergraduates at the University of Pittsburgh. No subjects were recruited from the authors' classes, and each subject participated in only one session.

Each session corresponded to one of the experimental conditions (Single, Aligned, or Opposed) and began with subjects giving informed consent and being seated at separate PC workstations. Written instructions were then distributed, presented on computer screens, and read out loud. Subjects then took a computer-based quiz about the instructions, followed with immediate and private feedback and explanations of the correct answers. Consistent with lab policy, no deception or false feedback was used in the experiment.

After the instructions and quiz, subjects were randomly assigned to a role that remained fixed throughout the session. In the Single condition, those roles were referred to as R (receiver) and S (moderate sender). In the Aligned and Opposed conditions, the roles were R (receiver), S1 (extremist sender), and S2 (moderate sender). In the experiment, we avoided the use of contextual labels and referred only to the abbreviations R, S, S1, and S2.

In each session, subjects played 24 rounds of the relevant game. At the beginning of a round, subjects were randomly matched into groups of either two or three, depending on condition, with one subject in each role in each group. Groups were selected with replacement so that it was possible to be matched with the same group in different rounds. To avoid reputation effects, subjects where not shown identifying information about the other subjects in their group.

The same sequence of randomly selected targets t was used for every session, regardless of condition.[12] This element of our design facilitates comparison of information transmission across sessions and conditions, which is our primary analytic focus. In all conditions, we referred to the receiver's ideal action t as simply "R's target." The moderate's shift is $s_M = 40$, which generates an ideal action for the moderate of $t + s_M$, which we referred to as "S's target" in the single sender condition and "S2's target" in the two sender conditions. We chose the value of s_M to ensure that the most informative equilibrium in the Single and Aligned conditions is a two-partition equilibrium. In the Aligned and Opposed conditions, the magnitude of the extremists' shift is $s_E = 80$, so "S1's target" is $t+s_E$ in the Aligned condition and $t-s_E$ in the Opposed condition. The divergence between the receiver's and extremist's target is sufficiently great that the extremist would never reveal any information to the receiver in a single sender game.

The experimental interface presented information both textually and graphically (see Figure 2). This interface intuitively conveys the spatial nature of the payoff functions and includes a payoff calculator, thus allowing subjects to focus cognitive resources on thinking strategically rather than on computing payoffs. Our interface gives senders the option to choose a point message (the default) or, following Cai and Wang (2006), to randomize over a range of messages. If senders

[12] The realized target values were 36, 84, -78, 42, 53, -95, 91, -46, 99, -42, -75, -51, 19, -22, 42, 79, 94, -32, -4, 84, 8, -29, -89, and -25.

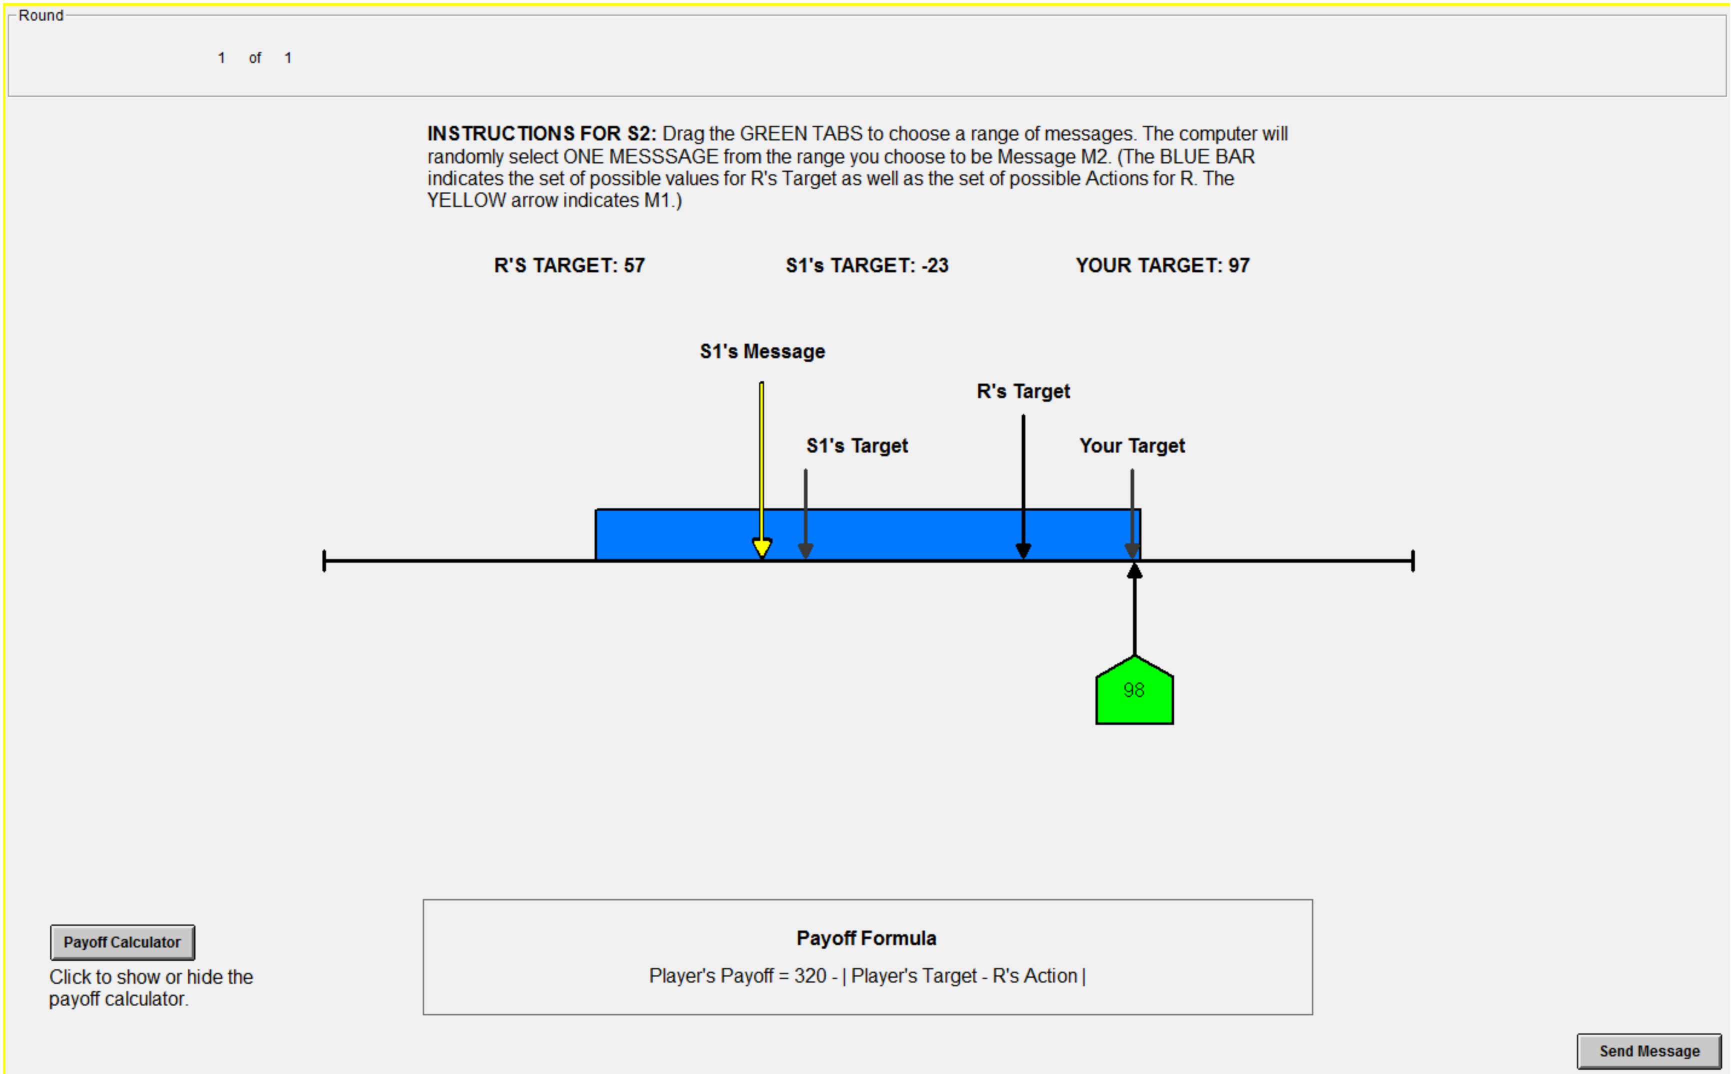

Figure 2: *Experimental Interface.* A snapshot of the experimental interface depicting the choice for the moderate sender after the extremist sender has already selected a message.

chose a range, they indicated the minimum and maximum message, and then the computer would select a single message from the range with uniform probability. The use of ranges and explicit randomization was intended to encourage the use of partition strategies and to discourage the kind of overcommunication found in previous experiments that results from overprecision.

At the end of every round, subjects were shown the results from that round for their group, including all messages, the action, every player's target, and every player's payoff. Each subject was also shown the results from her groups in previous rounds, but not the results for groups they did not belong to. Payoffs for each round were denominated in "points," with 320 points being the maximum possible points a player could earn in a round (if the receiver's action matched their own target exactly).

At the end of the experimental session, one random round of play was selected to calculate payoffs. Points from that round were converted to cash at the rate of $\$1$ for every 20 points, and subjects also received a $\$7$ show-up fee.

5 Results

We report results from 18 experimental sessions (six sessions for each of the Single, Aligned, and Opposed conditions) with 240 subjects (10 to 18 per session). In total, there were 91 moderate senders, 58 extremist senders, and 91 receivers. Table 1 presents details by session. Each session lasted under 2 hours. Payments ranged between $\$13.25$ and $\$23.00$ (including show-up fee), with an average payment of $\$20.21$. Our analysis is divided into three parts. First, we test our main hypotheses regarding the overall amount of information transmission (Section 5.1). Next, we characterize the behavior of subjects (Section 5.2). Finally, we describe how receivers respond to the senders' messages (Section 5.3).

Table 1: Session Details

Condition	Session	# Subjects	# Groups	# Obs.
Single	3	10	5	120
	4	10	5	120
	9	10	5	120
	10	12	6	144
	13	12	6	144
	18	12	6	144
		66	33	792
Aligned	1	15	5	120
	6	12	4	96
	8	15	5	120
	12	15	5	120
	15	12	4	96
	16	18	6	144
		87	29	696
Opposed	2	15	5	120
	5	12	4	96
	7	15	5	120
	11	18	6	144
	14	15	5	120
	17	12	4	96
		87	29	696
Total		240	91	2184

An observation (last column) is a group in a single period.

5.1 Information Transmission

Equilibrium analysis generated three primary hypotheses concerning the amount of information transmission. First, information transmission should be minimal in the Single sender condition (H1). When a single sender's shift is $s_M = 40$, the most informative equilibrium involves a partition into two regions of the state space, with the breakpoint at $t = -80$. Receivers should only be able to tell which region the target is in, but otherwise should not be able to guess its value.[13] Second, the additional message provided by the extremist sender in the Aligned condition should not reveal any more information. So, receivers in the Aligned condition should not be able to infer any more than they do in the Single condition (H2). Third, information transmission

[13] Even with two partitions, our parameterization with a shift of 40 implies that there should be very little information transmission. This is because the lower region of the partition is extremely small compared to the upper region, and there is babbling within each region.

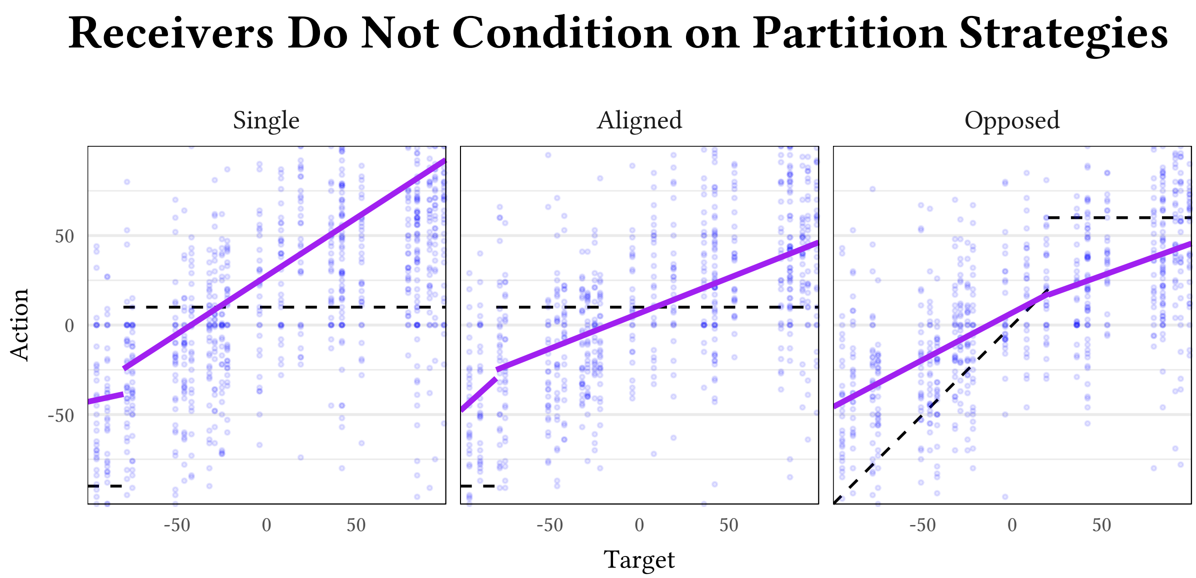

Figure 3: *Targets and Actions.* Solid lines display regression predictions based on models from Table 2 over equilibrium partitions, dashed lines indicates equilibrium predictions in the same partitions, and dots depict observed target-action pairs. More information is transmitted in each condition than is predicted by equilibrium theory.

should be greater with the addition of the extremist in the Opposed condition (H3), as competition disciplines sender behavior. Here, the semi-revealing equilibrium specifies a discontinuity at $t = 20$, with full-revelation below the breakpoint (where the receiver learns the target exactly) and babbling above it (where the receiver only infers the target is in the upper region).

Information transmission fundamentally refers to the relationship between observed *Actions* and *Targets*. The equilibrium strategies make specific predictions about these relationships. Therefore, we build models based on the piecewise structure of these equilibria. Specifically, we estimate treatment-specific regressions of *Action* on *Target*, interacted with the appropriate indicator for whether the *Target* was located in the *Low Region*: $t \in [-100, -80]$ for the Single and Aligned conditions and $t \in [-100, 20]$ for the Opposed condition. Figure 3 displays the equilibrium predictions (dashed lines) and the fitted regression lines (solid lines) based on these models, along with the observed *Target-Action* pairs. Table 2 reports the corresponding regression estimates.

Table 2: Regressions of Action on Target and Partition

	Single	Aligned	Opposed
Target	0.43^*	0.40^*	0.36^*
	(0.02)	(0.02)	(0.08)
Low Region	-31.60	36.09	-2.75
	(87.35)	(82.65)	(6.69)
Target $\times$ Low Region	-0.22	0.51	0.17
	(0.95)	(0.90)	(0.10)
Intercept	10.11^*	6.61^*	9.66
	(3.38)	(3.09)	(6.85)
n Observations	792	696	696
n Subjects	33	29	29
n Sessions	6	6	6
Error terms		Group SD	
Subject	12.9	9.4	3.8
Session	5.0	5.3	6.8
Residual	36.4	35.1	33.1

The table presents the results of separate mixed effects models of *Actions* for each treatment condition. For the Single and Aligned models, the *High Partition* begins at *Target* $= -80$; for Opposed, it begins at *Target* $= 20$. The model also includes random intercepts for Session and Subject to control for the panel structure of the data. $^*p < 0.05$.

In the Single and Aligned conditions, *Actions* are much more tightly related to *Targets* than predicted by equilibrium analysis. In the equilibria of the Single and Aligned conditions, receivers are predicted to choose *Action* $= -90$ in the low region and 10 in the high. The regression lines in the left and center panels of Figure 3 should therefore be flat. Similarly, the slope on *Target* should be zero, as should its sum with the slope on the interaction of *Target* and *Low Region*. But this is not what we observe. Instead, all of the regression lines in the left and middle panels have positive slopes. The estimated coefficients on *Target* in the Single and Aligned models from Table 2 bear this out, as they are both around 0.4. The estimated slopes on the interaction term alter these somewhat, though imprecisely, as few *Target* values were in the lower region.

Turning to the Opposed condition, we again see little evidence of equilibrium relationships between *Actions* and *Targets*. Indeed, the evidence suggests consistency in information transmission across conditions, rather than the sharp increase that is predicted. In the semi-revealing equilibrium, *Action* should equal *Target* in the lower region, and it should be a constant (60) in the upper. Therefore, the slope on *Target* should be 0, as it captures the relationship in the upper region, and that on the interaction term should be 1. That is, the regression line in the lower region should be diagonal, while that in the upper region should be flat. However, both regression lines have positive slope, and in the lower region the slope is smaller than predicted in equilibrium. In fact, there is striking similarity across all three regression models, in contrast to the differences predicted by equilibrium analysis.

Theory also predicts substantial discontinuities in the relationships between actions and targets in each condition, yet we see no evidence of this phenomenon. Figure 3 reveals negligible distances between regression lines at each of the breakpoints. Formally, in the Single and Aligned conditions, there should be a discontinuity at $t = -80$. The difference between the estimated actions just above and below this discontinuity should be $10 - (-90) = 100$. In the regression, this quantity is estimated by -1 times the sum of the slope on *Low Region* and $t = -80$ times the slope on the interaction term. Our estimate of this value is 15, with 95% interval $[-9, 39]$, in the Single Condition, and 6 $[-18, 29]$ in the Aligned condition. Neither is near 100, and we cannot reject the null of continuity in the piecewise regression. There should also be a discontinuity in the Opposed condition, at $t = 20$, where the *Action* should jump by $60 - 20 = 40$. This quantity is estimated by by -1 times the sum of the slope on *Low Region* and $t = 20$ times the slope on the interaction term. Instead, our estimate is actually negative, -1 $[-12, 10]$, which is closer to 20 than the estimates for the other two conditions, yet again we cannot reject the null of continuity and the estimate is still far from equilibrium prediction.

Despite the lack of support for equilibrium predictions, this analysis yields limited support for one of our hypotheses about information transmission. To be sure, there is substantially greater information transmission than predicted in both the Single condition, evidence against H1. However, the results for the Single and Aligned conditions are quite similar, indicating that the addition of the aligned extremist did not alter the relationship between *Actions* and *Targets*, thus supporting H2. Yet the broad consistency across all conditions means that the addition of the Opposed Extremist did not increase information transmission, in contrast to H3. Thus, it would be misleading to focus solely on H2 as support for the equilibrium predictions, as information transmission was similar across all three conditions.

In more substantive terms, we can measure the value of second opinions directly, based on how much more money receivers earned in those two conditions. Specifically, we regress the receiver's payoffs—converting points earned to US Dollars using the rate from the experiment sessions—on indicators for the two sender treatments, leaving the Single condition as the excluded category. The models reported in Table 3 also include session and subject random effects to account for the panel structure of our experimental data. Pooling all rounds together, receivers' payoffs ranged from $6.90 to $16, with a mean of about $14, excluding the show-up fee.[14] In keeping with H2, the overall effect of adding a second, aligned opinion is negligible, a statistically insignificant increase of $0.05. But in contrast to H3, Opposed second opinions offer only slightly more, a similarly insignificant $0.19, about one percent of the average payment. Moreover, the results do not change much when we allow for learning. If we include an indicator for the *Second Half*, i.e., the last 12 rounds, and interactions with the treatment indicators (column 2), we see very little change in the value of second opinions.

Our analysis of information transmission has yielded two principal findings. First, senders transmit far more information than equilibrium predicts in the Single sender condition, refuting H1. Thus, we find evidence of overcommunication, consistent with previous work. Second, adding a second sender has no effect, regardless of the configuration of the senders' preferences. While the results are consistent with H2, we find no evidence to support H3, contrary to the theoretical expectations provided by KM's analysis. Receivers are no better and no worse off from consulting a second opinion than they are consulting a single opinion.

These findings are closely related to each other. The addition of the opposed sender does not add value in part because senders in the Single and Aligned conditions already transmit so much information. For example, in the upper pooling regions, the correlation between *Action* and *Target* should be zero in the Single and Aligned conditions. Instead, it is 0.54 in the Single and

[14] Actual payoffs depended on a randomly selected round.

Table 3: Regression of Payoffs on Treatment

	[1]	[2]
Aligned	0.05	−0.05
	(0.23)	(0.24)
Opposed	0.18	0.29
	(0.23)	(0.24)
Aligned × Second Half		0.16
		(0.15)
Opposed × Second Half		−0.23
		(0.15)
Second Half		0.26*
		(0.10)
Intercept	13.96*	13.82*
	(0.16)	(0.17)
n Observations	2184	
n Subjects	91	
n Sessions	18	
Error terms	Group SD	
Subject	0.3	0.3
Session	0.4	0.4
Residual	1.5	1.5

The table presents the results of a mixed effects model of Payoffs on Treatment conditions, with the Single condition as the excluded category. Points were calculated as in the experiment ($320 - |\text{Target} - \text{Action}|$), and then converted to US Dollars at the rate used in the experimental sessions. The model also includes random intercepts for Session and Subject to control for the panel structure of the data.
$^*p < 0.05$.

0.53 in the Aligned. Yet the addition of the opposed sender also falls somewhat short of delivering full revelation in the lower region. There, the correlation between *Action* and *Target* should be one. Instead, it is 0.48, quite close to that in the other two conditions' pooling regions. In each case, the deviations from equilibrium predictions about information transmission work to erase differences across conditions.

5.2 Sender Behavior

To help understand why second opinions from opposed experts failed to improve information transmission in our experiment, we now focus on describing observed behavior rather than testing theoretically derived hypotheses. We begin with the senders. A few details of the experiment are important to keep in mind throughout the analysis. First, senders in the experiment could choose either to send a single message, or to identify an interval from which a single message would be randomly drawn. About 29% of selected messages were singletons; the remainder were intervals.[15] Therefore, we first analyze messages using a statistical model that accounts for the complete message strategies (i.e., minima, maxima, and midpoints of intervals), the boundaries of the message space, and the panel structure of the data. We also explore whether senders used the interval technology differentially across treatments by modeling the *Width of the Message Intervals*. Second, the two-sender conditions (Aligned and Opposed) both had a sequential structure, in keeping with KM and the general theme of seeking second opinions. Extremists sent their messages first, and moderates followed. Therefore, we report models that regress moderates' messages on targets and the messages from extremists. Third, senders played many periods of the game, and they may have learned over the course of the session to conceal more information, or to draw receivers closer to their desired actions, and so we fit models that allow for such differences. The results of these investigations suggest that senders use a strategy of naïve exaggeration. Therefore, our final analysis of senders is to estimate parsimonious, one-parameter models of messages for each sender.

Figure 4 plots the observed midpoints of the message intervals of senders in the experiment by treatment condition and sender bias, along with the results from a model of senders' message strategies. Focusing first on the raw data, many of the midpoints appear clustered above the truth-telling 45 degree line. The cluster is shifted out for extremists in the Aligned condition, in keeping with their larger shift parameter, and appropriately reflected (below the line) for extremists in the Opposed condition. Thus, there is a striking pattern of exaggeration where the midpoints appear to be chosen in a direction consistent with the sender's bias. For senders with positive shifts (all but the extremist in the Opposed condition), this strategy of exaggeration leads many message

[15] Senders sent single messages more often the more they played, although the changes were small. A paired-samples t-test indicates a statistically significant difference between the fraction of selected messages that were singletons in the first halves ($M = 0.25, SD = 0.11$) and the second halves ($M = 0.33, SD = 0.11$), ($t(17) = -6.49, p < .001$). There were no significant differences by sender type or condition.

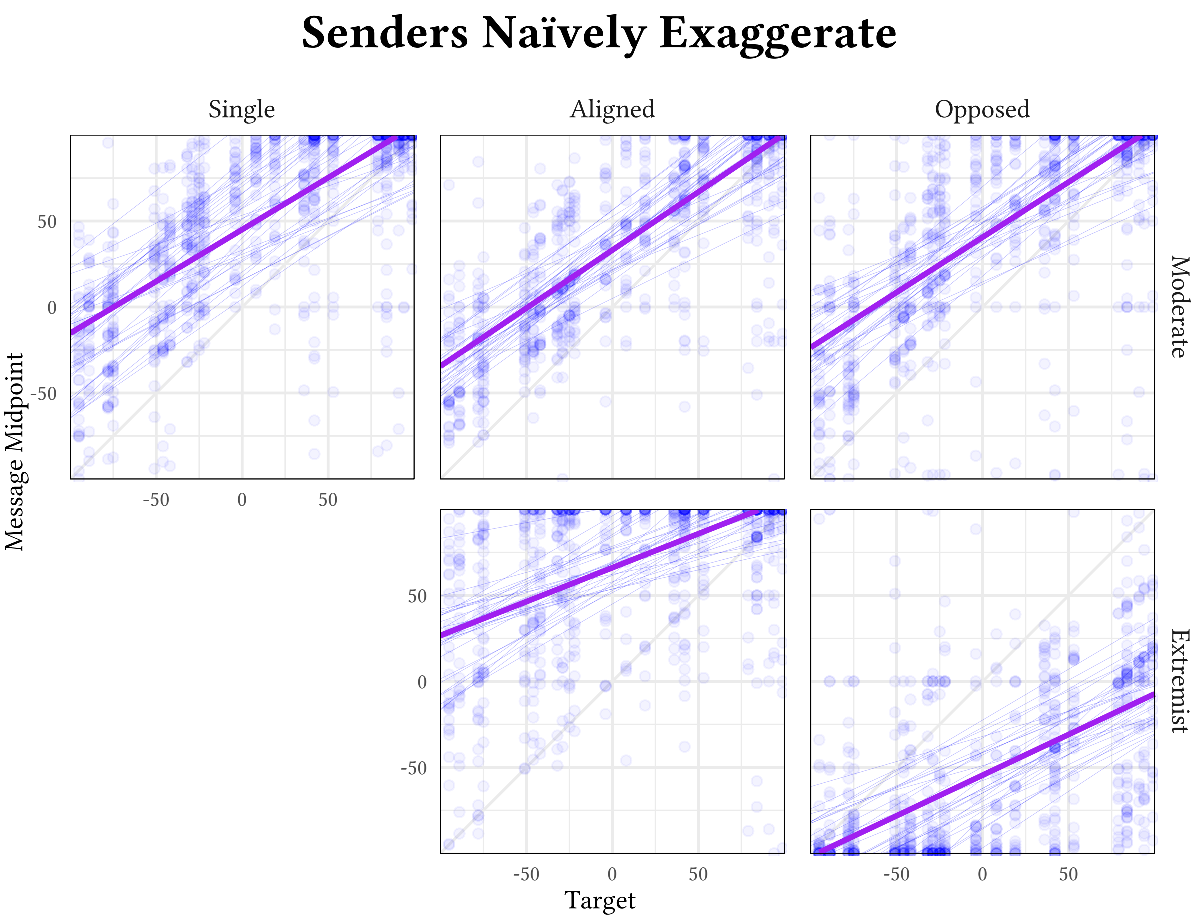

Figure 4: *Observed Message Strategies.* Points indicate the midpoints of the message ranges selected by senders, and lines depict the estimated midpoints based on a simultaneous equations, censored, mixed effects model. Regardless of treatment or bias, senders engage in naïve exaggeration, bounded by the available message space.

midpoints to run into the upper boundary of the message space, and for senders with negative shifts (Opposed extremists), it leads their message midpoints into the lower boundary.

We investigated senders' message strategies more formally with a statistical model, the results of which reinforce the visual analysis. Specifically, we estimated a simultaneous equations, censored, mixed effects model that was sufficiently complex to capture senders' selection of message ranges and the ceiling/floor effects created by the boundaries.[16] The estimates from this model

[16] The simultaneous equations are necessary given that senders choose two values (the minimum and maximum of the range), the censoring accounts for the boundedness of the message space, and the mixed effects account for the panel structure, including subject and session level heterogeneity. See the Appendix for a full description of and

are depicted as lines plotted over the message midpoints. The thickest line shows the grand mean across subjects and sessions, revealing that senders use a strategy in which the messages are generally inflated functions of the underlying target. The thin lines represent sender-level estimated strategies, suggesting that while there is a non-trivial amount of heterogeneity, this heterogeneity mostly takes the form of intercept shifts whereby different senders inflate their message by different amounts (rather than variation in the slopes of the individual message functions).

Thus, the model estimates depicted in Figure 4 suggest that senders used a strategy of naïve exaggeration when selecting messages. Formally, naïve exaggeration is specified by a simple rule. If the target is t, send the message $t + E$, if possible. If $t + E$ falls outside the bounds of the message space, send the message at the closer boundary. The censored regression lines in each of the panels of Figure 4 display exactly this strategy, even though the model is sufficiently rich to allow a variety of alternative strategies.

Of course, aspects of the experiment offered possibilities for variation in behavior across conditions. For example, senders could use the interval technology to select their messages, and they may have done so differently based on the presence and preferences of a second sender. Partition pooling strategies, in particular, would be made considerably easier by this feature of the experimental interface.[17] Since the equilibria in the Single and Aligned condition include larger pooling intervals than those in the Opposed condition, message intervals for both moderates and extremists might be more precise (narrower) in the Opposed condition than in the other two. Therefore, we examine the extent to which senders use intervals and how their message ranges vary with the configuration of preferences.

Table 4 presents regression models of the *Width of the Message Intervals* on a set of indicators for the type of sender (*Moderate* or *Extremist*) in each treatment condition. In general, the average *Width of the Message Intervals* is modest in the Single sender condition, covering less than 10% of the message space. The *Width* does not vary much at all with the condition or type of sender.

code for the model.

[17]Nevertheless, the default message is a point message, and we explicitly indicate that such point messages are permissible in the instructions.

Table 4: Regression of Width of the Message Interval on Treatment/Type

	[1]
Moderate/Aligned	7.87
	(6.20)
Extremist/Aligned	1.34
	(6.15)
Moderate/Opposed	2.60
	(6.22)
Extremist/Opposed	11.82
	(6.31)
Intercept	18.23*
	(4.44)
n Observations	3576
n Subjects	149
n Sessions	18
Error terms	Group SD
Subject	10.1
Session	9.6
Residual	30.6

The table presents the results of a mixed effects model of the Width of the Message Interval (Maximum Message - Minimium Message) on the interaction of Treatment conditions and Sender types, with the Moderate in the Single condition as the excluded category. The model also includes random intercepts for Session and Subject to control for the panel structure of the data. $^*p < 0.05$.

While the coefficients for the indicator variables are all positive, none are statistically significant at conventional levels.[18]

In addition to the interval technology, some senders could have taken advantage of the sequential nature of the experiment. Moderate senders, specifically, may have conditioned their strategies on the extremist's realized message. Equilibrium analysis offers some expectations for the relationship between these messages. In any two-partition equilibrium of the Aligned condition, the moderate's message will be uncorrelated with both the *Target* and the extremist's message within each region. However, in the semi-revealing equilibrium of the Opposed condi-

[18] We also estimated a similar model with interaction terms of the four covariates from Table 4 with the indicator for whether a period was in the *Second Half*. No interaction terms were significant.

Table 5: Regression of Moderate Message on Extremist Message

	Aligned	Opposed	Joint
Extremist Message	0.06^*	-0.09^*	0.06
	(0.03)	(0.04)	(0.03)
Target	0.62^*	0.63^*	0.62^*
	(0.02)	(0.03)	(0.03)
Opposed			5.78
			(6.69)
Extremist Message × Opposed			-0.15^*
			(0.05)
Target × Opposed			0.00
			(0.04)
Intercept	26.69^*	32.37^*	26.62^*
	(5.68)	(4.82)	(4.65)
n Observations	696	696	1392
n Subjects	29	29	58
n Sessions	6	6	12
Error terms		Group SD	
Subject	2.3	5.8	3.1
Session	13.1	9.2	9.5
Residual	34.2	43.9	39.4

The table presents the results of mixed effects models of the Moderate Sender's message on the Extremist Sender's and the Target. The first two models separate the Aligned and Opposed conditions, and the third is a joint model with interaction terms to allow comparisons of the two. The models also include random intercepts for Session and Subject to control for the panel structure of the data. $^*p < 0.05$.

tion, the requirement that the moderate "confirm" the extremist's truthful message implies that the moderate's message will be highly correlated with the extremist's message.

We test for this effect by regressing the *Moderate Message* on its preceding *Extremist Message*, controlling for *Target*. The results appear in Table 5. In the Aligned condition, *Moderate Message* is positively and significantly correlated with *Extremist Message*. This suggests that moderates may have attempted to collude with extremists due to the alignment of their incentives. In the Opposed condition, the coefficient is negative and statistically significant. The latter finding suggests that, rather than confirming extremists' messages, moderates choose their messages to counteract them. The greater the exaggeration by extremists (the farther to the left the message is from the target), moderates respond with greater exaggeration of their own (messages farther to the right). However, the coefficients on the extremist's message are relatively small in magni-

Table 6: Message Models over Time

	Single	Aligned Extremist	Opposed Extremist	Aligned Moderate	Opposed Moderate
Target	0.61*	0.31*	0.49*	0.62*	0.67*
	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)
Second Half	17.80*	13.39*	−13.37*	2.43	17.89*
	(2.41)	(3.15)	(3.08)	(4.26)	(5.04)
Target × Second Half	−0.12*	−0.09	−0.17*	−0.01	−0.18*
	(0.04)	(0.05)	(0.05)	(0.04)	(0.06)
Extremist Message				0.05	−0.06
				(0.04)	(0.06)
Extremist Message × Second Half				0.05	0.01
				(0.06)	(0.07)
Intercept	33.98*	54.62*	−46.37*	24.57*	25.48*
	(6.97)	(4.47)	(8.25)	(6.33)	(4.97)
n Observations	792	696	696	696	696
n Subjects	33	29	29	29	29
n Sessions	6	6	6	6	6
Error terms		Group SD			
Subject	2.0	6.0	3.1	2.6	5.7
Session	16.6	9.6	19.5	12.8	8.5
Residual	33.2	42.1	40.1	32.3	42.0

The table presents the results of mixed effects models of the the midpoint of senders' message intervals on *Target*, an indicator for the *Second Half* of periods, and their interaction. Separate models are presented by sender type. The models also include random intercepts for Session and Subject to control for the panel structure of the data. $^*p < 0.05$.

tude, especially compared to the coefficients on target, which suggests that while moderates are somewhat responsive to extremists, they largely ignore them.

Finally, senders may also have altered their strategies as they played more, learning as they gained experience. To analyze changes in message strategies, we estimated regression models of the midpoints of the senders' message intervals. We regressed these quantities on the *Target*, an indicator for whether a period was in the *Second Half*, and their interaction.[19] For moderate senders, we also included the *Extremist Message* and its interaction with *Second Half*, to account for the sequential nature of the game. The results appear in Table 6.

Across the board, the slopes on *Target* are positive and significant, and the intercepts are

[19] We also estimated alternative models using the minimum and maximum of the ranges as the dependent variables. We see similar results regardless of which dependent variable we use, so we report the midpoint models for ease of presentation.

large and signed similarly to senders' biases, in keeping with naïve exaggeration. Moreover, the coefficients on *Second Half* are positive and significant for the single sender, the aligned extremist, and the opposed moderate; and negative and significant for the Opposed extremist. All four appear to have learned to exaggerate more as play went on. A notable exception is the Aligned moderate, for whom we see no such effect. Finally, there appears to be no substantial relationship between time and the *Extremist Message* in the regression of moderates' messages in the two-sender conditions.

Given the welter of evidence suggesting that senders naïvely exaggerate, our final step is to estimate such exaggeration explicitly. Therefore, we return to our formal characterization of naïve exaggeration, and consider a rule in which a message (midpoint) is the target t plus some constant E. If $t + E$ exceeds the message ceiling of 100, assume the message sent is 100, and if it is below the floor at -100, let the message be -100. To estimate E at the level of the individual sender, we selected the value of E that minimized subject-specific root mean squared error. Ultimately, we have one estimated exaggeration parameter per subject.

The resulting distributions of estimated exaggeration parameters are depicted in Figure 5. In the Single condition, the distribution is bimodal. One mode occurs at about 40, which is this sender's shift parameter, and the second, larger mode occurs at 80, or twice the shift. This distribution suggests a kind of level-k thinking.[20] That is, some senders appear to simply add their shift to the target. Others—perhaps reasoning that receivers will catch on to the simple one-shift addition—instead add twice their shift to the target. In a crude level-k model focusing on the interaction between a single sender and the receiver, the magnitude of a level-k sender's exaggeration is indeed k multiplied by their shift.[21]

Remarkably, the higher mode vanishes completely in the Aligned condition. Focusing on the distribution of estimated exaggeration parameters for moderate senders, we can reject the

[20] See, for example, Camerer, Ho and Chong (2004), Crawford (2003), Nagel (1995), and Stahl and Wilson (1995).

[21] To see this, suppose level-0 senders tell the truth, $m_0(t) = t$, and level-0 receivers believe them, $a_0(m) = m$. Level-1 senders would exploit level-0 receivers' credulity and exaggerate enough so that the resulting action is equal to their own targets, so $m_1(t) = t + s$. But then level-1 receivers correctly invert this, so $a_1(m) = m - s$. Level-2 senders then exaggerate enough so that level-1 receivers choose $t + s$, so $m_2(t) = t + 2s$, and so on.

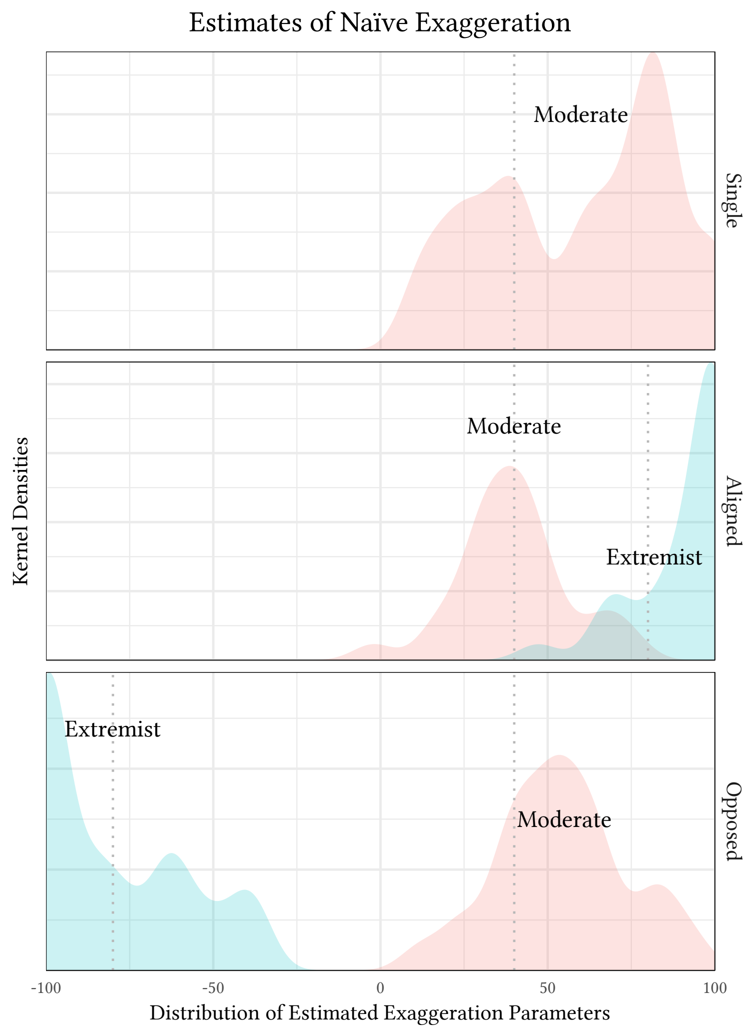

Figure 5: *Estimates of Naïve Exaggeration.* The figures display kernel densities of estimated exaggeration parameters for each type of sender in each condition.

equality of these distributions between the Single and Aligned conditions (Kolmogorov-Smirnov $p = 0.001$). The second mode is also substantially reduced in the Opposed condition. Comparing the moderate distributions in the Aligned and Opposed conditions indicates a modest shift to the right (Kolmogorov-Smirnov $p = 0.004$). Thus, it appears that moderate senders are responsive to the presence of a second opinion, which generally tempers the exaggeration of the moderate senders, regardless of the second sender's direction of bias. Interestingly, the distributions for the Single and Opposed moderate are closer to each other (Kolmogorov-Smirnov $p = 0.111$), although the latter remains slightly left-shifted. Thus, introducing competition, rather than alignment, between the senders has a small but noticeable effect of increasing exaggeration by moderates, but not enough to match the level of exaggeration in the Single condition. While senders behave in ways that are for the most part naïve, they are nevertheless somewhat responsive to—that is, not completely ignorant about—changes in the strategic environment.

We speculate that these changes in the moderate senders' behavior may be because they believe that their messages will be perceived as more trustworthy when compared to the extremists' messages. Indeed, the estimated exaggeration parameters for extremists are essentially at the ceiling and floor of the possible values, and we cannot reject the equality of the distributions of their magnitudes (i.e., comparing the distribution of aligned extremists' exaggeration with the distribution of the reflection about the origin of the Opposed extremists').

5.3 Receiver Behavior

Based on our analysis of information transmission in Section 5.1, receivers did not limit their actions to one or two values, as they would have if they believed senders played partition or babbling strategies. The substantial correlation between targets and actions suggests that receivers learn much more than they would if senders strategically limited the information they convey. Indeed, in Section 5.2, we showed that senders engage in a strategically naïve fashion. Consistent with level-k models of limited strategic sophistication, senders play simple "addition" strategies. The question we address in this section is: How do receivers respond to the senders' messages? That is, to what extent can we say that receivers best respond to senders' observed behavior?

Given that senders engage primarily in exaggeration, a sophisticated receiver should be able to invert the messages to back out relatively accurate guesses about the true underlying state. Furthermore, they should be able to recognize that exaggeration by extremists is greater than exaggeration by moderates. These considerations generate two expectations for what we should observe for receiver behavior in the experiment. First, since senders engage in "addition" strategies, we should therefore observe receivers engage in "subtraction." Second, we should observe that receivers' pay more attention to the messages of moderates than those of extremists.

Figure 6 plots receivers' actions against messages, separately for each condition and type of sender, along with smoothed LOESS regression predictions. The visual evidence is consistent with a strategy of inverting messages by subtraction. When receivers respond to senders with positive shifts (all but the opposed extremists), we find that a substantial proportion of actions fall below the 45 degree line. Similarly, when they respond to opposed extremists, we see that actions fall primarily above the 45 degree line, consistent with subtracting negative exaggeration. Furthermore, we see evidence that subtraction is the greatest at the boundaries in the direction of the senders' shifts, whereas messages at the boundaries opposite of those shifts appear to be taken at face value. And for messages that are unexpected, at least based on the naïve exaggeration rule, the "subtraction" rule breaks down and loses predictive accuracy. In those cases, receivers tend to choose actions closer to 0, the expected value of the target based on prior information alone. These patterns are consistent with strategic sophistication on the part of receivers in that they subtract the most when senders are most likely to have exaggerated and subtract the least (or perhaps none at all) when messages are much less likely to have been generated by exaggeration.

Next, we model receivers' actions as a function of messages and experimental condition. Specifically, Table 7 presents results of regression models of *Action* as a function of the *Moderate Message* and *Extremist Message*. We include estimates for separate models for moderates' messages (first column) and extremists' messages (second column) and joint model specifications that include both messages and appropriate interactions (last three columns). The moderate-only and extremist-only models allow us to consider each sender's message in isolation, thereby facilitating their interpretation. The joint model, of course, provides a more accurate picture of how much receivers rely on each sender's message, by effectively accounting for the correlation between messages. Regardless, there are few important inferential differences between these models. Finally, we also account for the possibility of learning by estimating the joint model separately for the first and second half of the periods. Each model includes random effects (intercepts) for sessions and subjects to account for the panel structure of the data.

Consistent with the visual interpretation of Figure 6, the coefficient on moderates' messages

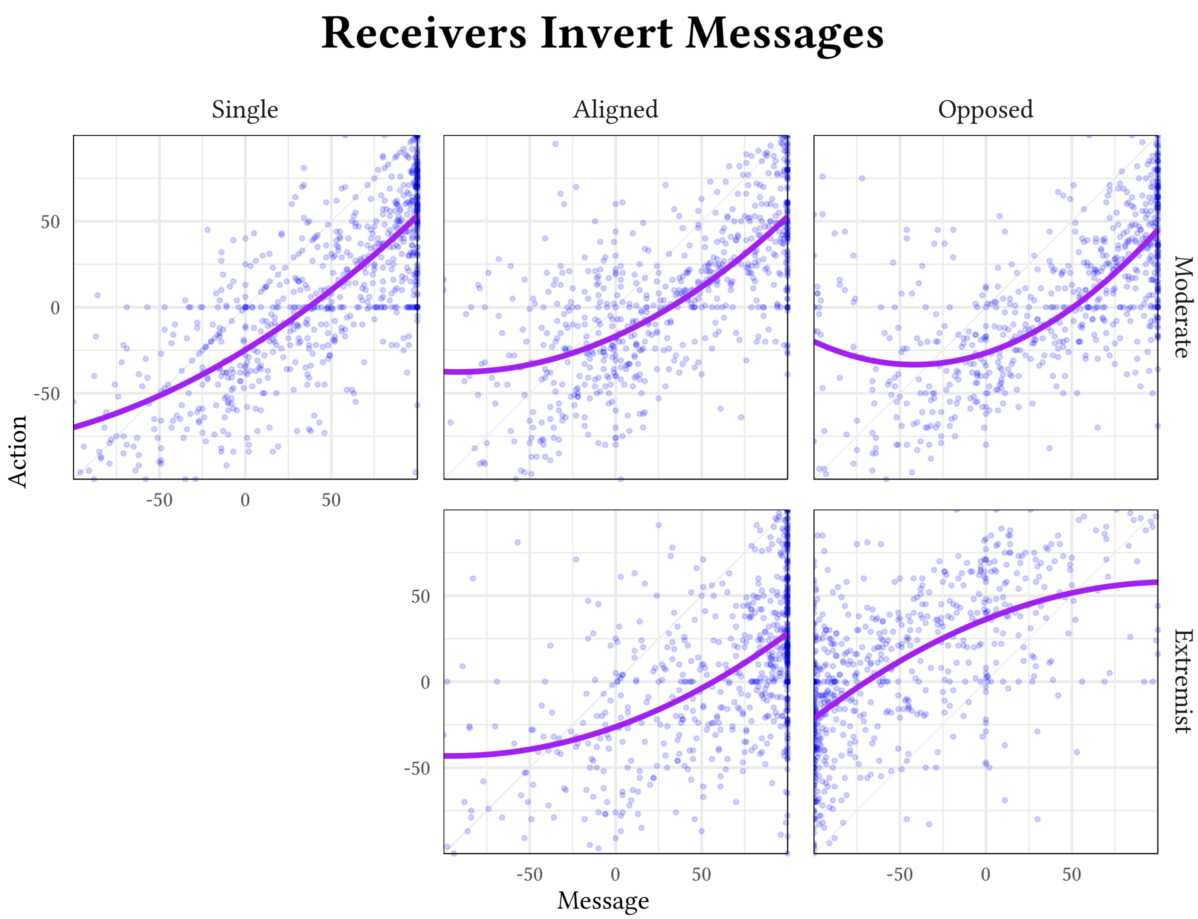

Figure 6: *Actions are correlated with messages.* Points depict observed message-action pairs, and lines illustrate predictions from local regression (LOESS) models. There are substantial correlations between actions and messages throughout each condition and with both types of sender.

in the single sender condition is relatively large, but less than 1. This magnitude is consistent with subtracting greater amounts at the upper boundary than in the interior of the message space. Indeed, the intercept is negative and statistically significant, providing strong evidence for subtraction, and its magnitude indicates that receivers subtract approximately 22 units when the moderate's message is $m_S = 0$. Contrast this amount of subtraction with the amount at the boundary, which we can estimate from the predicted action $\hat{a} = 52$ when $m_S = 100$, implying that receivers subtract 48, an amount that is more than twice what they subtract at $m_S = 0$ and is slightly larger than the moderate sender's bias.

Table 7: Regressions of Actions

	Moderate	Extremist	Joint	First Half	Second Half
Moderate Message	0.74^*		0.74^*	0.70^*	0.82^*
	(0.02)		(0.02)	(0.03)	(0.03)
Extremist Message		0.48^*	0.30^*	0.26^*	0.35^*
		(0.03)	(0.02)	(0.03)	(0.03)
Moderate Message × Aligned	-0.15^*		-0.24^*	-0.20^*	-0.32^*
	(0.03)		(0.03)	(0.04)	(0.04)
Moderate Message × Opposed	-0.25^*		-0.35^*	-0.28^*	-0.47^*
	(0.03)		(0.03)	(0.04)	(0.04)
Extremist Message × Opposed		0.02	0.08^*	0.11^*	0.04
		(0.04)	(0.03)	(0.05)	(0.04)
Aligned	10.75^*		-5.06	-3.16	-5.94
	(5.02)		(4.53)	(4.56)	(5.62)
Opposed	8.56	54.60^*	32.54^*	26.81^*	43.24^*
	(5.05)	(4.24)	(4.44)	(4.67)	(5.67)
Intercept	-22.45^*	-22.51^*	-22.73^*	-19.23^*	-29.18^*
	(3.51)	(3.13)	(2.94)	(2.97)	(3.82)
n Observations	2184	1392	2184	1092	1092
n Subjects	91	58	91	91	91
n Sessions	18	12	18	18	18
Error terms			Group SD		
Subject	11.1	7.2	11.2	8.4	14.1
Session	6.3	3.7	4.6	4.7	5.0
Residual	30.3	36.4	27.4	29.1	25.0

The table presents the mixed effects models of *Action* on the Moderate and Extremist Senders' Messages, separately in the first two columns, and joint, interacted with treatment conditions, in the third column. The model also includes random intercepts for Session, and Subject to control for the panel structure of the data. $^*p < 0.05$.

The addition of a second message changes receivers' behavior in two ways. First, receivers are less responsive to the moderate's message in the presence of a second opinion, based on the negative and statistically significant coefficients for the interaction terms in the moderate message-only. Furthermore, the positive coefficients for the Aligned and Opposed condition indicators mean that receivers subtract less from a moderate message of $m_M = 0$ when there are two messages than when there is a single message. To understand the change in the receivers' subtraction strategy, consider the change in magnitude of the slope in conjunction with the change in intercept. In the Aligned condition, the estimated magnitude of subtraction at the boundary $m_M = 100$ remains similar (a predicted action of $\hat{a} = 47$ implies subtracting 53) even though the amount of subtraction for messages in the middle of the space ($m_M = 0$) is smaller.

In the Opposed condition, the estimated amount of subtraction is even larger at the boundary (predicted action $\hat{a} = 35$ implies subtracting 65) while similar in the middle of the space as in the Aligned. These results suggest that in the presence of a second aligned sender, receivers rely less on the moderate's message.

Turning to the relationship between *Action* and the *Extremist Message*, we find that receivers respond to extremists, but rely on them less than they rely on the moderate's message. The evidence for this claim is the smaller magnitude of the coefficient for the extremist's message (which indicates the slope in the Aligned condition). Receivers also subtract greater amounts from the extremist's message than the moderate's message, particularly at the boundaries. In the Aligned condition, the model suggests that given an extremist's message of $m_E = 100$, the predicted action would be $\hat{a} = 25$, with the receiver therefore subtracting 75 from the message. In the Opposed condition, a message at the lower boundary $m_E = -100$ leads to a predicted action of $\hat{a} = -18$, with the receiver therefore subtracting -82. When either sender's messages at the boundaries are considered alone, receivers subtract an amount roughly equal to the sender's shift. The results thus suggest that receivers take second opinions from extremists seriously but discount them relative to opinions from moderates.[22] Receivers thus appear to be sufficiently sophisticated to take into account the difference in the senders' biases when interpreting their messages.

The last three columns of Table 7 present the model of the receiver's strategy that takes both senders' messages into account, first combining all 24 periods and then dividing the periods into the first and second halves. We examine two quantities of interest computed from the estimated coefficients on the messages and their interactions. The first quantity of interest is the sum of the message coefficients, which tells us the overall credence that receivers find in the set of observed messages. If the coefficients sum to 1, it would indicate that the receiver completely inverts the message by subtraction only, whereas coefficients that sum to less than 1 indicate that receivers

[22] This result is not surprising, given that the model of actions as a function of the extremist's message treats the moderate's message as if it were not observed when in fact they were. Although the extremist model estimates therefore suffer from omitted variable bias and are misleading, we find them to be analytically useful.

shade their actions toward the middle of the space, subtracting less in the interior than at the boundary. We saw the latter in response to moderates in the Single sender condition (coefficient of 0.74). In the two-sender conditions, the combined weights are slightly higher but roughly similar in magnitude: 0.80 in the Aligned condition, and 0.77 in the Opposed condition. Comparing these quantities for the first and second halves of the experimental periods, we see a modest increase in all three conditions: from 0.70 to 0.83 in the Single condition, from 0.77 to 0.86 in the Aligned, and from 0.80 to 0.90 in the Opposed. Receivers appear to rely more heavily on the senders' messages later in the experiment. Nevertheless, the similarity of these quantities across conditions helps to explain why there is no effect of adding a second sender or of the configuration of preferences on overall information transmission. Given that senders of all types naïvely exaggerate, senders in all conditions extract the same amount of information from the observed messages.

The second quantity of interest is the relative weights receivers place on the different senders' messages (computed as a proportion of the sum). In the Single sender condition, the weight on the moderate is necessarily 100%. In the Aligned condition, the weights are approximately 63% on the moderate's message and 37% on the extremist's message. These weights are not far from what one might expect from inverting the ratio of the senders' shift parameters in a manner consistent with receivers' correctly believing that senders' naïvely exaggerate. Interestingly, receivers rely much less on moderates in the Opposed condition, as the weights are nearly equal, with 51% of the weight on the moderate's message and 49% on the extremist's message. Whereas receivers take a weighted average in the Aligned condition, their behavior changes to taking a simple average in the Opposed condition.

Based on these models, we conclude that receivers were both aware of senders' naïve exaggeration strategy, and that they worked to counteract it by developing a "subtraction" strategy of their own. To make this point explicit, we fit a rule-based model of receiver behavior similar to the naïve exaggeration model at the end of Section 5.2. First, we assume that receivers in the Single condition use simple rule in which they choose the action a equal to the moderate's message m_M

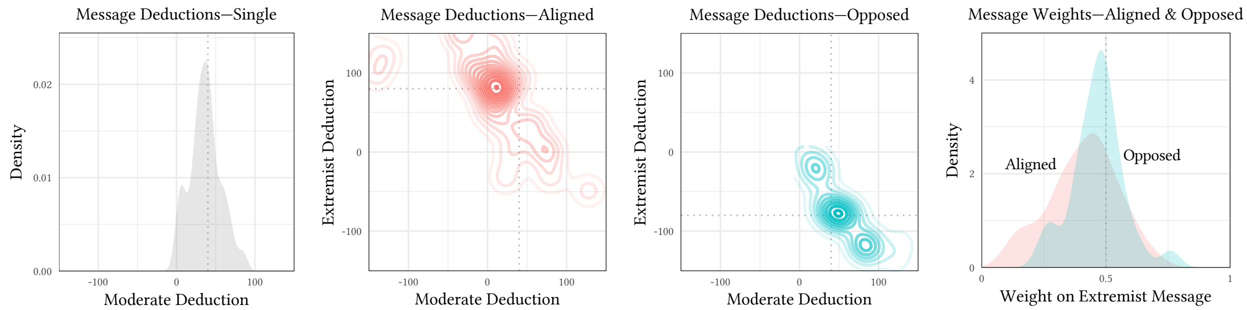

Figure 7: *Estimates of Message Deductions.* The leftmost panel shows the kernel density for receivers' deduction parameters in the Single condition. The middle two panels depict kernel densities for the estimated pairs of deduction parameters receivers used in the Aligned and Opposed conditions. The rightmost panel displays shows the kernel density of the estimated weights receivers placed on the message sent by extremists. Dotted gray lines depict exogenous shift parameters (left three plots) and the equal weight line (rightmost plot).

minus some deduction D. If $m_M - D$ exceeds the message ceiling of 100, assume the action chosen is 100, and if it is below the floor at -100, let the action be -100. For receivers in the Aligned and Opposed condition, we simply assume receivers follow that same rule for each sender, to yield potential actions $m_M - D_M$ and $m_E - D_E$. We then assume that receivers choose a weighted average of these two potential actions, appropriately censored at the boundaries. Thus we have a one-parameter model for receivers in the Single condition, and a three-parameter model for receivers in the Aligned and Opposed conditions. We again estimate parameters at the level of the individual receiver, selecting the values that minimize subject-specific root mean squared error.

Figure 7 displays the results for these simple behavioral models. The leftmost panel shows the kernel density for receivers' deduction parameters in the Single condition. Here, we see that there is again a mode at 40, the shift for the moderate sender. There are also smaller areas of mass near zero and 80, suggesting that some receivers either naïvely implemented the sender's message, or expected the preponderance of level-2 thinking we saw in the distribution of senders' exaggeration parameters.

The middle two panels of Figure 7 depict the kernel densities for the estimated pairs of deduction parameters receivers used in the Aligned and Opposed conditions. Densities are presented with contour plots, with darker areas indicating more mass. In both cases, the moderate message deduction is on the horizontal axis and that for the extremist message is on the vertical. The dotted gray lines show the shift parameters for each sender. Again, we see that the largest mode of each distribution is located very near the exogenous shift parameters. Finally, the last panel of Figure 7 shows the density of the estimated weights receivers placed on the message sent by the extremist. It appears again that receivers put somewhat more weight on messages from opposed extremists than on those from aligned extremists. However, there is substantial overlap in the distributions, meaning that even in the Aligned condition, receivers gave considerable weight to extremist messages.

We draw two important conclusions from our analysis of receiver behavior. First, receivers rely on the aggregate set of messages equally across treatments. Given senders' use of naïve exaggeration, this helps to explain why the overall levels of information transmission are similar across treatments. Second, receivers adjust their behavior across conditions in a way that appears to reflect an accurate understanding of how changes in the structure of the communicative environment induce changes in senders' exaggeration behavior. In contrast to senders' strategic naïveté, receivers display a remarkable degree of strategic sophistication. Taken together, these findings explain why second opinions have limited value.

6 Conclusion

In this paper, we experimentally investigated the effect of providing a second opinion in cheap talk games, comparing behavior and information transmission in environments with a single sender, two aligned senders, and two opposed senders. Our experimental setup closely mirrors the settings analyzed by Crawford and Sobel (1982) and Krishna and Morgan (2001*b*). Theoretically, equilibrium analysis leads to a series of stark predictions: adding an aligned sender should have no effect on behavioral or informational changes while adding an opposed sender should have a noticeable effect, yielding substantially more information revelation. Experimentally, we have shown evidence of a different set of phenomena. Comparing the Single sender condition with either of the Aligned or Opposed two-sender conditions, we see differences in behavior but with similar levels of information transmission. As senders modify their behavior due to changes in the strategic environment, receivers adjust their behavior accordingly. The results of our experiment thus complicate the theoretical understanding of competition, expertise, and credible communication.

Nevertheless, the differences between the theoretical model and these experimental results obscure an underlying similarity. In the theoretical model, the partition equilibria in the single sender and multiple aligned sender conditions results from the following chain of logic. If the receiver were to naïvely enact the message sent by the sender, then that sender would have an incentive to exaggerate. Sequential rationality then implies that receivers will not take the sender's message at face value, but instead invert the exaggerated message to infer the true state. Continuing this chain, the sender would exaggerate even more, and in the limit, no information transmission is possible (full separation cannot emerge). Recognizing the implications of exaggeration, but also that the sender and receiver have shared (i.e., partial common) interests, senders have incentives to share some information, but not too much. This latter insight generates an equilibrium with a partition structure in which senders must therefore refrain from saying too much and must be purposefully vague. A different pattern emerges in the opposed senders condition. Here, theoretically, the two opposing senders can coordinate their message strategies to successfully transmit the hidden target information for a large swath of the space. Competition should have a disciplining effect, which encourages the truth.

Experimentally, the evidence we have presented is consistent with the first part of this chain. Senders appear to behave as though they expect receivers to naïvely enact their messages, in a phenomenon that we have described as naïve exaggeration (Minozzi and Woon, 2016). This is true regardless of the alignment of the senders' preferences. In the Aligned condition, we observe senders inflate their messages in the same direction, and in the opposed condition, we observe them inflate their messages in opposite directions. The difference between theory and experiment can be explained by a failure of sequential rationality. That is, senders fail to realize that receivers expect this kind of naïve exaggeration. If they did, then they would be able to select a message strategy that would result in an increased payoff. Thus, observed behavior does not come close to approaching babbling. As a consequence, senders do not need to use partition strategies, or to discipline each other when their interests are opposed. Indeed, the process of devising such strategies strikes us as a cognitively challenging task, requiring tremendous sophistication, and it is not clear to us when or whether partition strategies would emerge naturally. Thus, we suggest that future research on strategic communication should focus on the link between cognition and credibility.

References

Ambrus, Attila and Satoru Takahashi. 2008. "Multi-sender Cheap Talk with Restricted State Spaces." *Theoretical Economics* 3(1):1–27.

Austen-Smith, David. 1993. "Interested Experts and Policy Advice: Multiple Referrals under Open Rule." *Games and Economic Behavior* 5(1):3–43.

Baron, David P. and Adam Meirowitz. 2006. "Fully-Revealing Equilibria of Multiple-Sender Signaling and Screening Models." *Social Choice and Welfare* 26(3):455–470.

Battaglini, Marco. 2002. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica* 70(4):1379–1401.

Battaglini, Marco, Ernest K. Lai, Wooyoung Lim and Joseph Tao-Yi Wang and. 2016. "The Informational Theory of Legislative Committees: An Experimental Analysis." Unpublished manuscript.

Battaglini, Marco and Roland Benabou. 2003. "Trust, coordination, and the industrial organization of political activism." *Journal of the European Economic Association* 1(4):851–889.

Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim and Geoffrey B. Sprinkle. 1998. "Experimental Evidence on the Evolution of Meaning of Messages in Sender-Receiver Games." *American Economic Review* 88(5):1323–1340.

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

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

Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics* 119(3):861–898.

Carpenter, Bob, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A. Brubaker, Jiqiang Guo, Peter Li and Allen Riddell. 2016. "Stan: A probabilistic programming language." *Journal of Statistical Software* in press. http://mc-stan.org/.
URL: *http://mc-stan.org/*

Chen, Ying, Navin Kartik and Joel Sobel. 2008. "Selecting Cheap-Talk Equilibria." *Econometrica* 76(1):117–136.

Crawford, Vincent and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50(6):1431–1451.

Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review* 93(1):133–149.

Dickhaut, John W., Kevin A. McCabe and Arijit Mukherji. 1995. "An Experimental Study of Strategic Information Transmission." *Economic Theory* 6(3):389–403.

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

Gentzkow, Matthew and Emir Kamenica. 2017. "Competition in persuasion." *The Review of Economic Studies* 84(1):300–322.

Gilligan, Thomas W. and Keith Krehbiel. 1989. "Asymmetric Information and Legislative Rules with a Heterogeneous Committee." *American Journal of Political Science* 33(2):459–490.

Gneezy, Uri. 2005. "Deception: The Role of Consequences." *American Economic Review* 95(1):384–394.

Hurkens, Sjaak and Navin Kartik. 2009. "Would I Lie to You? On Social Preferences and Lying Aversion." *Experimental Economics* 12(2):180–192.

Kawagoe, Toshiji and Hirokazu Takizawa. 2009. "Equilibrium Refinement vs. Level-k Analysis: An Experimental Study of Cheap-talk Games with Private Information." *Games and Economic Behavior* 66(1):238–255.

Krehbiel, Keith. 2001. "Plausibility of Signals by a Heterogeneous Committee." *American Political Science Review* 95(2):453–457.

Krishna, Vijay and John Morgan. 2001*a*. "Asymmetric Information and Legislative Rules: Some Amendments." *American Political Science Review* 95(2):435–452.

Krishna, Vijay and John Morgan. 2001*b*. "A Model of Expertise." *Quarterly Journal of Economics* 116(2).

Lai, Ernest K., Wooyoung Lim and Joseph Tao-Yi Wang. 2011. "Experimental implementations and robustness of fully revealing equilibria in multidimensional cheap talk." Working paper.

McGee, Andrew and Huanxing Yang. 2013. "Cheap Talk with Two Senders and Complementary Information." *Games and Economic Behavior* 79:181–191.

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

Minozzi, William. 2011. "A Jamming Theory of Politics." *Journal of Politics* 73(2):301–315.

Minozzi, William and Jonathan Woon. 2013. "Lying aversion, lobbying, and context in a strategic communication experiment." *Journal of Theoretical Politics* 25(3):309–337.

Minozzi, William and Jonathan Woon. 2016. "Competition, Preference Uncertainty and Jamming: A Strategic Communication Experiment." *Games and Economic Behavior* 96:97–114.

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

Sánchez-Pagés, Santiago and Marc Vorsatz. 2007. "An Experimental Study of Truth-telling in a Sender-Receiver Game." *Games and Economic Behavior* 61(1):86–112.

Sánchez-Pagés, Santiago and Marc Vorsatz. 2009. "Enjoy the Silence: An Experiment on Truth-telling." *Experimental Economics* 12(2):220–241.

Stahl, Dale O. and Paul W. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10(1):218–254.

Vespa, Emanuel and Alistair Wilson. 2016. "Communication with Multiple Senders: An Experiment." *Quantitative Economics* 7:1–36.

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

Wolinsky, Asher. 2002. "Eliciting Information from Multiple Experts." *Games and Economic Behavior* 41(1):141–160.

Appendix

A1 Equilibrium Specifications

Here, we explicitly state the most informative equilibria of the games corresponding to each of the experimental conditions. We focus on equilibria in which senders who are indifferent between sending informative messages and babbling send informative messages. There are many equilibria that are informationally equivalent to the ones stated here that relax this requirement.

In the single sender condition, the sender uses the message strategy

$$m^{\text{single}}(t) \in \begin{cases} U[-100,-80] & \text{if } t \leq -80 \\ U[-80,100] & \text{otherwise,} \end{cases}$$

and the receiver uses the action strategy

$$a^{\text{single}}(m) = \begin{cases} -90 & \text{if } m \leq -80 \\ 10 & \text{otherwise,} \end{cases}$$

with the beliefs

$$q^{\text{single}}(t|m) = \begin{cases} U[-100,-80] & \text{if } m \leq -80 \\ U[-80,100] & \text{otherwise.} \end{cases}$$

In the Aligned condition, the extremist sender, who moves first and has shift $s_E = 80$, uses $m_E^{\text{aligned}}(t) = m^{\text{single}}(t)$, the message strategy from the single sender condition. The moderate sender, who moves second and has shift $s_M = 40$, uses a similar strategy that simply ignores the extremist's message, so that $m_M^{\text{aligned}}(t, m_E) = m^{\text{single}}(t)$ for all m_E. The actions and beliefs of the receiver in the Aligned condition similarly ignore the extremist: $a^{\text{aligned}}(m_E, m_M) = a^{\text{single}}(m_M)$ and $q^{\text{aligned}}(t|m_E, m_M) = q^{\text{single}}(t|m_M)$. Finally, in the opposed condition, the extremist sender, who moves first and has shift $s_E = -80$, uses

$$m_E^{\text{opposed}}(t) = \begin{cases} t & \text{if } t < 20 \\ U[20,100] & \text{otherwise.} \end{cases}$$

The moderate sender, who moves second and has shift $s_M = 40$, plays

$$m_M^{\text{opposed}}(t, m_E) = \begin{cases} t & \text{if } m_E = t < 20 \\ U[20,100] & \text{if } t \geq 20 \\ m_E & \text{if } t < m_E \ \& \ t < 20 \text{ (off-the-path)} \\ \max(m_E + 80, t + 40) & \text{if } m_E < t < 20 \text{ (off-the-path).} \end{cases}$$

Finally, the receiver uses

$$a^{\text{opposed}}(m_M, m_E) = \begin{cases} m_E & \text{if } m_E < 20 \,\&\, m_M = m_E \\ 60 & \text{if } m_E > 20 \text{ or } m_M > 20 \\ m_E & \text{if } m_E, m_M < 20 \,\&\, m_M < m_E + 80 \text{ (off-the-path)} \\ m_M & \text{if } m_E, m_M < 20 \,\&\, m_M \geq m_E + 80 \text{ (off-the-path)}, \end{cases}$$

and has the beliefs

$$q^{\text{opposed}}(t|m_M, m_E) = \begin{cases} \mathbb{I}[t = m_E] & \text{if } m_E < 20 \,\&\, m_M = m_E \\ 60 & \text{if } m_E > 20 \text{ or } m_M > 20 \\ \mathbb{I}[t = m_E] & \text{if } m_E, m_M < 20 \,\&\, m_M < m_E + 80 \text{ (off-the-path)} \\ \mathbb{I}[t = m_M] & \text{if } m_E, m_M < 20 \,\&\, m_M \geq m_E + 80 \text{ (off-the-path)}. \end{cases}$$

A2 Statistical Model for Figure 4

To model message strategies, we used a simultaneous, censored, mixed effects regression model written and estimated in the Stan modeling language (Carpenter et al., 2016). Code for the model appears below.

```
data {
  // observations are split into four datasets based on whether
  // (1) message midpoint hits the boundary and
  // (2) message range hits the maximum size

  int<lower = 1> N_id; // number of subjects
  int<lower = 1> N_period; // number of periods

  // observations in which midpoint does not hit boundary and range
  //  does not hit maximum value
  int<lower = 1> N_1;
  int<lower = 1, upper = N_id> id_1[N_1];
  int<lower = 1, upper = N_period> period_1[N_1];
  vector[N_1] Target_1;
  vector[N_1] Message_Midpoint_1;
  vector[N_1] log_Message_Range_Size_1;

  // observations in which midpoint does not hit boundary but range
  //  does hit maximum value
  int<lower = 0> N_2;
  int<lower = 1, upper = N_id> id_2[N_2];
  int<lower = 1, upper = N_period> period_2[N_2];
  vector[N_2] Target_2;
  vector[N_2] Message_Midpoint_2;
  vector[N_2] log_Max_Message_Range_Size_2;

  // observations in which midpoint hits upper boundary
  int<lower = 0> N_3;
  int<lower = 1, upper = N_id> id_3[N_3];
```

```
  int<lower = 1, upper = N_period> period_3[N_3];
  vector[N_3] Target_3;
  real U; // right censor point for Message_Midpoint

  // observations in which midpoint hits lower boundary
  int<lower = 0> N_4;
  int<lower = 1, upper = N_id> id_4[N_4];
  int<lower = 1, upper = N_period> period_4[N_4];
  vector[N_4] Target_4;
  real L; // left censor point for Message_Midpoint
}

parameters {
  real a0;              // grand mean for intercept in midpoint model
  vector[N_id] a_e_id; // (unscaled) subject-level random intercept
  real<lower = 0> a_s_id; // scale for subject-level random intercept
  real b0;              // grand mean for slope in midpoint model
  vector[N_id] b_e_id; // (unscaled) subject-level random slope
  real<lower = 0> b_s_id; // scale for subject-level random slope
  real<lower = 0> s_midpoint; // scale for error in midpoint model
  real g0;              // grand mean for intercept in range model
  vector[N_id] g_e_id; // (unscaled) subject-level random intercept
  real<lower = 0> g_s_id; // scale for subject-level random intercept
  real h0;              // grand mean for slope in range model
  vector[N_id] h_e_id; // (unscaled) subject-level random slope
  real<lower = 0> h_s_id; // scale for subject-level random slope
  real<lower = 0> s_logrange; // scale for error in range model
}

model {
  // local declarations
  // containers for subject-level coefficients
  vector[N_id] a_id;
  vector[N_id] b_id;
  vector[N_id] g_id;
  vector[N_id] h_id;
  // containers for midpoint model linear predictors
  vector[N_1] m_midpoint_1;
  vector[N_2] m_midpoint_2;
  vector[N_3] m_midpoint_3;
  vector[N_4] m_midpoint_4;
  // containers for range model linear predictors
  vector[N_1] m_logrange_1;
  vector[N_2] m_logrange_2;
  vector[N_3] m_logrange_3;
  vector[N_4] m_logrange_4;

  // local definitions
  a_id = a0 + a_s_id * (a_e_id - mean(a_e_id));
  b_id = b0 + b_s_id * (b_e_id - mean(b_e_id));
  g_id = g0 + g_s_id * (g_e_id - mean(g_e_id));
  h_id = h0 + h_s_id * (h_e_id - mean(h_e_id));
  for (n in 1:N_1) {
   m_midpoint_1[n] = a_id[id_1[n]] + b_id[id_1[n]] * Target_1[n];
```

```
    m_logrange_1[n] = g_id[id_1[n]] + h_id[id_1[n]] * Target_1[n];
  }
  for (n in 1:N_2) {
    m_midpoint_2[n] = a_id[id_2[n]] + b_id[id_2[n]] * Target_2[n];
    m_logrange_2[n] = g_id[id_2[n]] + h_id[id_2[n]] * Target_2[n];
  }
  for (n in 1:N_3) {
    m_midpoint_3[n] = a_id[id_3[n]] + b_id[id_3[n]] * Target_3[n];
    m_logrange_3[n] = g_id[id_3[n]] + h_id[id_3[n]] * Target_3[n];
  }
  for (n in 1:N_4) {
    m_midpoint_4[n] = a_id[id_4[n]] + b_id[id_4[n]] * Target_4[n];
    m_logrange_4[n] = g_id[id_4[n]] + h_id[id_4[n]] * Target_4[n];
  }

  // prior
  a0 ~ normal(0, 10);
  a_s_id ~ cauchy(0, 2.5);
  a_e_id ~ normal(0, 1);
  b0 ~ normal(0, 10);
  b_s_id ~ cauchy(0, 2.5);
  b_e_id ~ normal(0, 1);
  s_midpoint ~ cauchy(0, 2.5);
  g0 ~ normal(0, 10);
  g_s_id ~ cauchy(0, 2.5);
  g_e_id ~ normal(0, 1);
  h0 ~ normal(0, 10);
  h_s_id ~ cauchy(0, 2.5);
  h_e_id ~ normal(0, 1);
  s_logrange ~ cauchy(0, 2.5);

  // posterior
  Message_Midpoint_1 ~ normal(m_midpoint_1, s_midpoint);
  Message_Midpoint_2 ~ normal(m_midpoint_2, s_midpoint);
  log_Message_Range_Size_1 ~ cauchy(m_logrange_1, s_logrange);
  // increment log probability for censoring
  target +=
    cauchy_lccdf(log_Max_Message_Range_Size_2 | m_logrange_2, s_logrange) +
    normal_lccdf(U | m_midpoint_3, s_midpoint) +
    normal_lcdf(L | m_midpoint_4, s_midpoint);
}
```

Table A1: Model of Midpoint & Range

	Estimate	Std. Error	95% CI
Midpoint Equation			
Intercept	0.28	0.01	$[0.27, 0.30]$
Target	0.62	0.01	$[0.60, 0.64]$
Subject Intercept Scale	0.48	0.03	$[0.43, 0.54]$
Subject Slope Scale	0.23	0.02	$[0.19, 0.27]$
Error Scale	0.43	0.01	$[0.42, 0.44]$
log(Range) Equation			
Intercept	0.21	0.00	$[0.20, 0.22]$
Target	0.04	0.01	$[0.02, 0.05]$
Subject Intercept Scale	0.13	0.01	$[0.11, 0.15]$
Subject Slope Scale	0.09	0.01	$[0.07, 0.11]$
Error Scale	0.08	0.00	$[0.08, 0.09]$

$n = 2190$. This table presents the estimated model based on the above code, and depicted in Figure 4. The constant is suppressed, as are the constituent terms for Moderate and Extremist, as the use of all conditions as interaction terms permits direct estimation of quantites of interest. The model was fitted in Stan 2.10, with four chains, each with 1000 warmup iterations and 1000 sampling iterations, after which diagnostics indicated convergence (e.g., $\hat{R} \leq 1.1$ and $n_{\text{eff}} > 400$ for all parameters).

Table A2: Message Models over Time (Min and Max)

	Single		Aligned Extremist		Opposed Extremist		Aligned Moderate		Opposed Moderate	
	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max
Target	0.63*	0.59*	0.33*	0.30*	0.41*	0.57*	0.65*	0.59*	0.69*	0.65*
	(0.03)	(0.02)	(0.04)	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.05)	(0.04)
Second Half	20.39*	15.37*	14.14*	12.67*	−8.47*	−18.38*	4.84	0.36	19.62*	16.50*
	(2.86)	(2.21)	(3.87)	(2.85)	(2.87)	(3.79)	(5.25)	(4.06)	(5.83)	(4.93)
Target × Second Half	−0.13*	−0.11*	−0.08	−0.11*	−0.11*	−0.22*	−0.01	−0.01	−0.18*	−0.18*
	(0.05)	(0.04)	(0.07)	(0.05)	(0.05)	(0.07)	(0.05)	(0.04)	(0.07)	(0.06)
Extremist Message							0.04	0.07	−0.05	−0.06
							(0.05)	(0.04)	(0.07)	(0.05)
Extremist Message × Second Half							0.03	0.06	0.00	0.02
							(0.07)	(0.05)	(0.09)	(0.07)
Intercept	23.75*	44.66*	44.30*	65.29*	−63.64*	−29.21*	11.38	37.58*	14.31*	37.02*
	(9.44)	(5.02)	(7.14)	(3.54)	(5.82)	(11.01)	(8.48)	(3.52)	(5.41)	(6.87)
n Observations	792	792	696	696	696	696	696	696	696	696
n Subjects	33	33	29	29	29	29	29	29	29	29
n Sessions	6	6	6	6	6	6	6	6	6	6
Error terms					Group SD					
Subject	3.0	1.7	8.0	4.3	3.9	3.6	3.8	2.0	8.1	5.1
Session	23.1	11.6	15.7	6.4	13.3	27.3	19.3	6.3	8.1	14.5
Residual	39.8	30.8	51.1	37.2	37.3	50.6	39.4	29.6	49.4	40.1

The table presents the results of mixed effects models of the the minima and maxima of senders' message intervals on *Target*, an indicator for the *Second Half* of periods, and their interaction. Separate models are presented by sender type and outcome variable. The models also include random intercepts for Session and Subject to control for the panel structure of the data. $^*p < 0.05$.

Appendix: Instructions for Aligned Condition

Instructions

General Information

This is an experiment in communication. The University of Pittsburgh has provided funds for this research. You will be paid in cash for your participation, and the exact amount you receive will be determined during the experiment and will depend on your decisions and the decisions of others. You will be paid your earnings privately, meaning that no other participant will find out how much you earn. These earnings will be paid to you at the end of the experiment along with the $7 participation payment.

Each participant has a printed copy of these instructions, and you may refer to them at any time.

If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you. Please do not talk, exclaim, or try to communicate with other participants in any way. Also, please ensure that your phones are off and your personal belongings are put away for the duration of the experiment. Participants intentionally violating the rules will be asked to leave the experiment and may not be paid.

Roles, Rounds, and Matching

Each participant will be assigned to one of three roles: S1, S2, or R. Your role will be assigned before the first round and will remain fixed throughout the experiment.

In this experiment you will make decisions in a series of rounds, and there are a total of 24 rounds. **Each round is a separate decision task.**

Before every round, you will be randomly matched with two other participants. In every group of three participants there will be one player in each role (one S1 player, one S2 player, and one R player).

You will not know the identity of the other participants you are matched with in any round, and your earnings for each round depend only on your action in that round and the actions of the participants you are matched with in that round.

Appendix: Instructions for Aligned Condition

Targets

In every round there will be a set of targets:

Player R's Target will be a randomly selected number between -100 and 100. Each number is equally likely to be R's target. R's Target for one round does not affect the value that is randomly selected for any other round.

Player S1's Target will always be 80 more than Player R's Target. For example, if R's Target is 100, then S1's Target is 180; if R's Target is 0, then S1's Target is 80, etc.

Player S2's Target will always be 40 more than Player R's Target. For example, if R's Target is 100, then S2's Target is 140; if R's Target is 0, then S2's target is 40, etc.

Note that it is possible for S1's or S2's Targets to be outside the set of possible Actions.

Sequence

The sequence of actions in every round is as follows:

1. **Player S1** observes the set of targets and chooses a **Range of Messages**. The minimum and maximum of the range can each be any number from -100 to 100. The computer then randomly selects **Message M1** from the numbers in the range chosen by S1. Each message in the range is equally likely to be selected.

2. **Player S2** observes the set of targets and the Message M1 and also chooses a **Range of Messages**. The computer then randomly selects **Message M2** from the range of numbers chosen by S2, with each message in the range equally likely to be selected.

3. **Player R** observes both Messages M1 and M2 (but none of the Targets or the ranges chosen by S1 or S2) and chooses an **Action**.

4. The computer determines each player's **Payoff** as a function of his or her Target and the Action chosen by R.

Appendix: Instructions for Aligned Condition

Payoffs

In each round, each player's payoff depends on how close R's Action is to his or her own Target. Specifically, each player earns 320 points if R's Action equals his or her own Target and 1 point less for each unit of difference between R's Action and his or her Target. Mathematically, this is described by the following formula (where the straight lines indicate absolute value):

$$\text{Player's Payoff} = 320 - |\text{Player's Target} - \text{R's Action}|$$

Note that the messages M1 and M2 are not part of the payoff formula. To illustrate, consider a few examples.

Example 1: R's Target is 50, so S1's Target is 130 and S2's Target is 90. If R chooses the Action 50, R's payoff is 320 since the Action equals R's Target. The difference between R's Action and S1's Target is 80, so S1's payoff is 240. The difference between R's Action and S2's Target is 40, so S2's payoff is 280. If R instead chooses the Action 0, then R's payoff would be 270, S1's payoff would be 190, and S2's payoff would be 230.

Example 2: R's Target is -60, so S1's Target is 20 and S2's Target is -20. If R chooses the Action -90, then R's payoff is 290, S1's payoff is 210, and S2's payoff is 250. If R instead chooses the Action 40, then R's payoff would be 220, S1's payoff would be 300, and S2's payoff would be 260.

Of course these are only a few examples. During the experiment, the software will provide you with a "Payoff Calculator" that will compute each player's payoff for any combination of R's Target and Action.

Payment

At the end of the experiment, the computer will randomly select ONE round to count for payment. Your points from this round will be converted to cash at the rate of $1 for every 20 points. We will pay you this amount in addition to the $7 participation payment.

Appendix: Instructions for Aligned Condition

Sample Screens

We will now see what the screens look like for each type of player during the experiment.

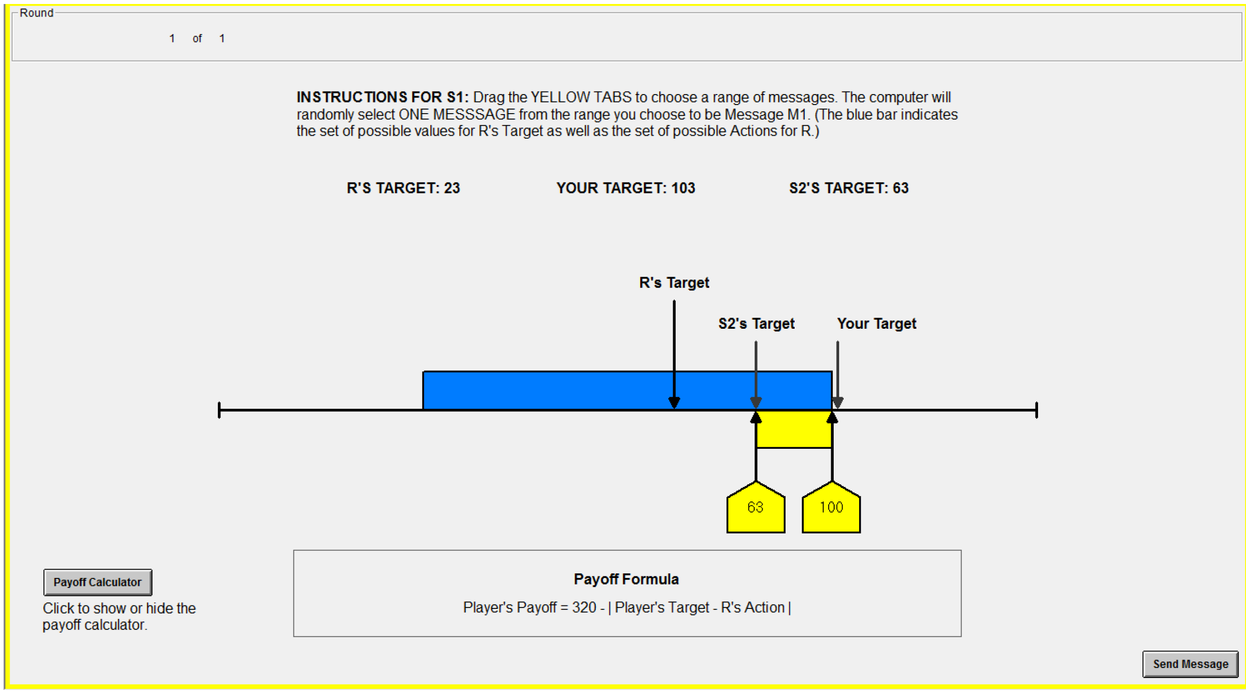

This is the screen that will be seen only by Player S1. There is a brief set of instructions at the top of the screen. Below it, the targets are presented both textually and graphically. To choose one end of the range of messages, Player S1 drags the yellow tab to any position along the blue bar. (The blue bar indicates the range of possible values for R's Target.) Go ahead and try dragging the yellow tab now. Player S1 will then see a second yellow tab, which can then be moved to any position along the blue bar to indicate the other end of the range. The yellow bar between the two tabs represents Player S1's range of messages, and Player S1 can adjust the yellow tabs until he or she is satisfied with the position and width of the range. Go ahead try adjusting both tabs.

Note that there is a button in the lower left marked "Payoff Calculator." You can click on this button to reveal two white tabs. You can use these tabs to calculate hypothetical payoffs for any possible values of R's Target and Action. If you move the white tabs to different positions, the text in the payoff calculator box changes to indicate what each player's payoff would be. Note also that you can hide the payoff calculator by clicking on the button again.

When Player S1 is finished choosing a range, he or she will click on the "Send Message" button in the lower right-hand corner of the screen. Please click on the "Send Message" button now to see Player S2's screen.

Appendix: Instructions for Aligned Condition

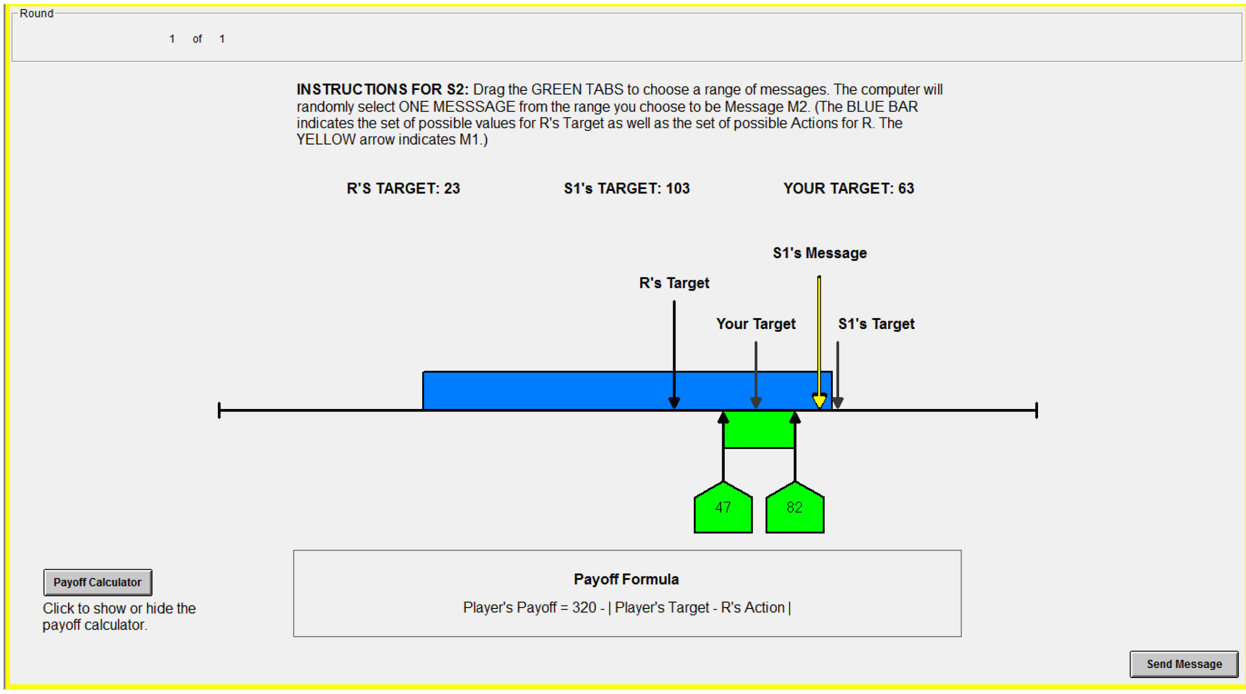

Player S2's screen is similar to Player S1's screen except that there is a yellow arrow indicating the position of Message M1 (which the computer randomly selected from the range chosen by Player S1). Player S2 then chooses a range of messages by dragging the green tabs. Go ahead and choose a range of messages for Player S2 and then click on the "Send Message" button.

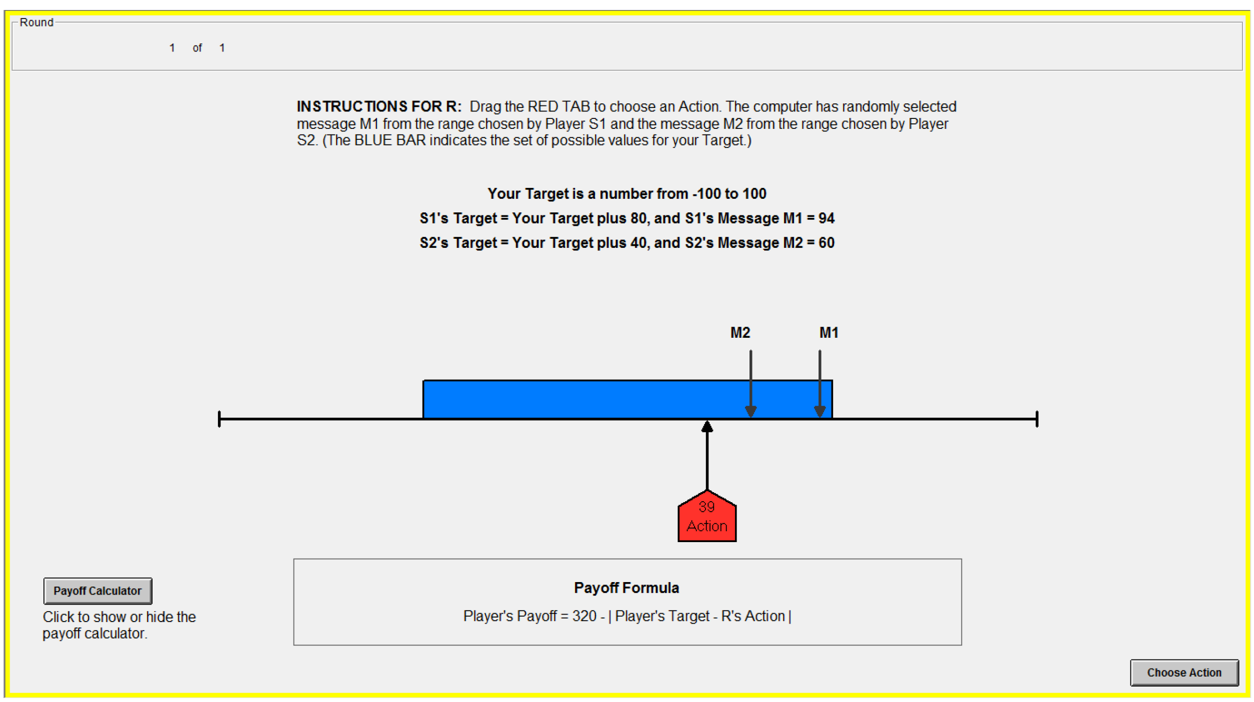

After Player S2 chooses his or her range, Player R will see this screen. Below the instructions, Player R will see the numerical values of Messages M1 and M2. Remember that the computer randomly selects these messages from the ranges chosen by Players S1 and S2. The messages are also indicated graphically in the middle of the screen, and there is a single red tab that Player R uses to choose an Action. As with the other tabs, it shows the numerical value of its location after it is moved. Note that Player R also has a payoff calculator with both calculator tabs set to a default value of 0. Go ahead and choose an Action and then click on the "Choose Action" button.

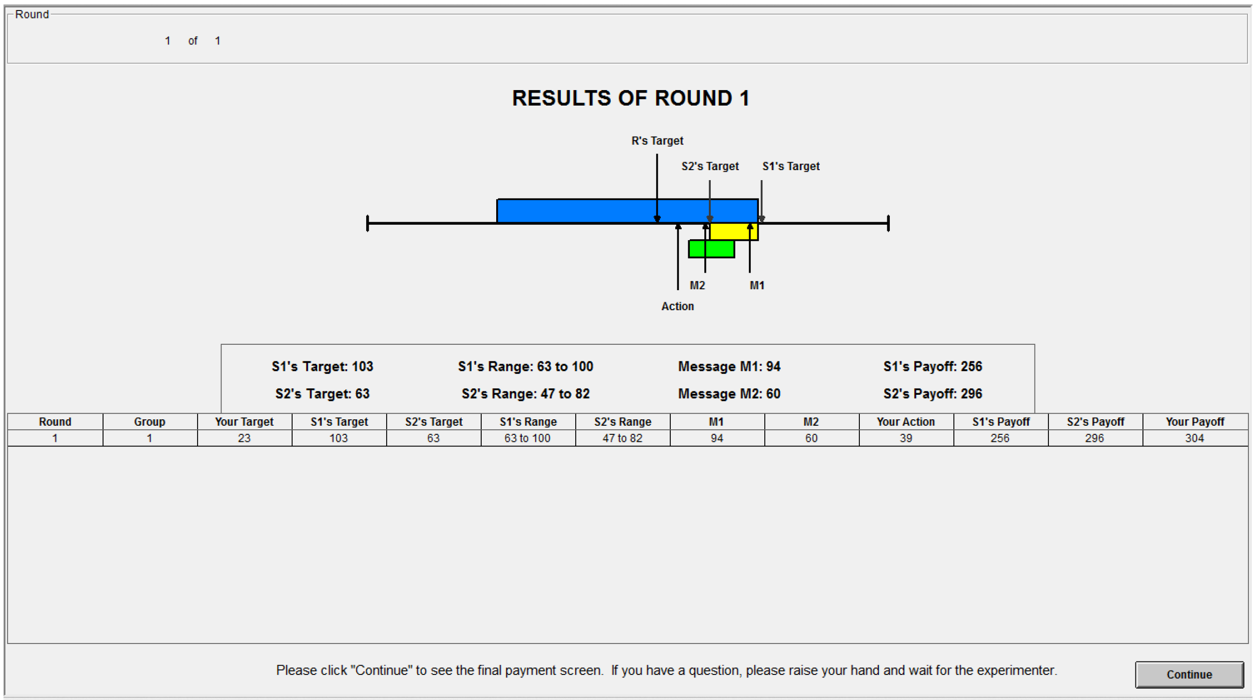

Round	Group	Your Target	S1's Target	S2's Target	S1's Range	S2's Range	M1	M2	Your Action	S1's Payoff	S2's Payoff	Your Payoff
1	1	23	103	63	63 to 100	47 to 82	94	60	39	256	296	304

At the end of every round, you will see this screen, which shows you the results from the round—including all of the targets, both ranges, both messages, the action, and the payoffs for each player in your group. At the bottom of the screen, it will show the results of every previous round that you played. Please click on the "Continue" button now.

Appendix: Instructions for Aligned Condition

SUMMARY

Targets

R's Target = Number between -100 and 100

S1's Target = R's Target + 80

S2's Target = R's Target + 40

Sequence

1. S1 sees all targets, then chooses a range of messages between -100 and 100; M1 is a random number from the range chosen by S1

2. S2 sees all targets and M1, then chooses a range of messages between -100 and 100; M2 is a random number from the range chosen by S2

3. R sees M1 and M2 and chooses an Action

Payoffs

Player's Payoff = 320 – |Player's Target – R's Action|

Payment

One round randomly selected for payment.

Appendix: Instructions for Aligned Condition

INSTRUCTION CHECK. To check your understanding of the decision tasks, please answer the questions below. When you are finished, feedback about the correct answers will be shown on the screen. Note that your quiz answers do not affect your earnings, but you must attempt to answer all of the questions before the computer will check your answers. During the quiz, you are free to refer to your printed instructions.

Once everyone has completed the instruction questions, we will begin the experiment. If you have any further questions at this time, please raise your hand and the experimenter will come to you.

1. In every round, will you be matched with same participants? [Yes, No]

2. Player R's target can be any number from: [0 to 10, 0 to 100, -100 to 100, -150 to 150]

3. If Player R's target is -40, then Player S1's target is: [-120, -80, 0, 40]

4. If Player R's target is 30, then Player S2's target is: [-50, -10, 70, 110]

5. Is -34 to 61 a valid range of messages for Player S1 or Player S2 to choose? [Yes, No]

6. Is -100 to 100 a valid range of messages? [Yes, No]

7. If Player S1 chooses the range 72 to 90, is it possible for the computer to randomly select 80 or 90 as the value for Message M1? [80 only, 90 only, Both, Neither]

8. If Player S2 chooses the range 18 to 18, is it possible for the computer to randomly select 18 or 20 as the value for Message M2? [18 only, 20 only, Both, Neither]

9. If Player R's Target is 10 and Player R chooses the action 50, how many points will <u>Player S1</u> receive? [200, 240, 260, 280]

10. If Player R's Target is -80 and Player R chooses the action 0, how many points will <u>Player S2</u> receive? [180, 240, 280, 320]

11. Suppose you are Player R, Message M1 is 0, Message M2 is 40, and you choose Action 30. If the Target turns out to be 65, how many points will you receive? [65, 255, 285, 320]