

ACCEPTED VERSION

Sookie Xue Zhang, Ralph-Christopher Bayer

Delegation based on cheap talk

Theory and Decision: an international journal for multidisciplinary advances in decision sciences, 2022

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2021

*This is a post-peer-review, pre-copyedit version of an article published in **Theory and Decision: an international journal for multidisciplinary advances in decision sciences**, The final authenticated version is available online at: <http://doi.org/10.1007/s11238-022-09887-8>*

PERMISSIONS

<https://www.springer.com/gp/open-access/publication-policies/self-archiving-policy>

Self-archiving for articles in subscription-based journals

Springer journals' [policy on preprint sharing](#).

By signing the Copyright Transfer Statement you still retain substantial rights, such as self-archiving:

*Author(s) are permitted to self-archive a pre-print and an author's **accepted manuscript** version of their Article.*

.....

b. An Author's Accepted Manuscript (AAM) is the version accepted for publication in a journal following peer review but prior to copyediting and typesetting that can be made available under the following conditions:

(i) Author(s) retain the right to make an AAM of their Article available on their own personal, self-maintained website immediately on acceptance,

(ii) Author(s) retain the right to make an AAM of their Article available for public release on any of the following 12 months after first publication ("Embargo Period"): their employer's internal website; their institutional and/or funder repositories. AAMs may also be deposited in such repositories immediately on acceptance, provided that they are not made publicly available until after the Embargo Period.

An acknowledgement in the following form should be included, together with a link to the published version on the publisher's website: "This is a post-peer-review, pre-copyedit version of an article published in [insert journal title]. The final authenticated version is available online at: [http://dx.doi.org/\[insert DOI\]](http://dx.doi.org/[insert DOI])".

When publishing an article in a subscription journal, without open access, authors sign the Copyright Transfer Statement (CTS) which also details Springer's self-archiving policy.

See Springer Nature [terms of reuse](#) for archived author accepted manuscripts (AAMs) of subscription articles.

30 May 2022

<http://hdl.handle.net/2440/130549>

Delegation Based on Cheap Talk*

Sookie Xue Zhang[†], Ralph-Christopher Bayer[‡]

April 1, 2022

Abstract

We study a real-effort environment, where a delegator has to decide if and to whom to delegate a task. Applicants send cheap-talk messages about their past performance before the delegator decides. We experimentally test the theoretical prediction that information transmission does not occur in equilibrium. In our experiment, we vary the message space available to the applicants and compare the information transmitted as well as the level of efficiency achieved. Depending on the treatment, applicants can either submit a *Number* indicating past performance, an *Interval* in which past performance falls, or a free *Text* message. We observe that messages contain information in all treatments. Interestingly, information transmission occurs only in the treatments where messages are intervals or free text. Social welfare is higher if messages are intervals or free text than precise numbers. Gender and ethnicity stereotypes only influence delegation in the *Number* treatment, where no information transmission takes place.

Keywords: Delegation, Cheap Talk, Experiment, Efficient Allocation

JEL Classification: C90, D61, D82

*We appreciate the coordinating editor handling the paper and the anonymous reviewers for very helpful comments that led to a much improved paper. We gratefully acknowledge funding support from the Fundamental Research Funds for the Central Universities [Project No.2722019PY042].

[†]Corresponding author. Wenlan School of Business, Zhongnan University of Economics and Law, Wuhan 430064, China. E-mail address: sookie.xue.zhang@outlook.com.

[‡]School of Economics and Public Policy, University of Adelaide, SA 5005, Australia. E-mail address: ralph.bayer@adelaide.edu.au.

1 Introduction

Delegation is an important management tool to increase productivity by exploiting comparative advantage in skills (Holmström, 1977). Unfortunately, asymmetric information about individuals' productivity is a significant barrier to efficient delegation. Communication, as a potential remedy, plays an essential role in delegation processes. In theory, the effectiveness of cheap-talk communication crucially depends on the alignment of interests held by the parties involved (Crawford and Sobel, 1982). Information only flows in equilibrium and can only improve efficiency if the objectives of sender and receiver are not perfectly opposed. With perfectly opposed objectives, in equilibrium, delegators will disregard messages due to the inherent incentive of applicants to lie.

It is unclear if the incentive to lie prevents information transmission in reality. There is considerable evidence that humans often reveal private information through messages in individual decision making (e.g., Fischbacher and Föllmi-Heusi, 2013; Gneezy et al., 2018) and in games (e.g., Gneezy, 2005; Hurkens and Kartik, 2009; Sutter, 2009) even if they are incentivized not to do so. A recent meta-analysis of 90 experimental studies on patterns of reporting behaviour shows that subjects forgo, on average, about three-quarters of the potential gains from lying (Abeler et al., 2019). With this evidence in mind, we ask if cheap-talk communication can make delegation more efficient even if the agents have opposing interests. In order to answer this question, we analyze behaviour in an experimental cheap-talk delegation game with real effort, where only equilibria without information transmission exist. We are particularly interested in the effect of different message spaces on information transmission and efficiency. We allow for gender and ethnicity stereotypes to impact delegation decisions by using avatars that indicate our participants' gender and ethnic backgrounds.

We employ a one-principal, two-agents setting that captures the prevalent competitiveness among candidates in real-world delegation processes. Delegation processes we have in mind are outsourcing, delegating a task to consultants, or hiring an individual for a specific task. Our setting is less characteristic of delegation within a firm where there is more verifiable information and delegation happens repeatedly. Our delegation game is played once and learning is not possible. In our experiments, a delegator has to decide if she wants to delegate the task of adding numbers that pays a piece rate to one of two applicants or to perform the task herself. She makes the decision after receiving cheap-talk messages from applicants about their abilities. Applicants are paid a bonus if given the task and therefore have an incentive to send messages that achieve delegation. Delegators profit from delegating if the chosen applicant's performance in the addition task exceeds their own to an extent such that improved performance at least covers the delegation bonus. Our design eliminates the effects

of moral hazard, wealth effects and the delegators' uncertainty about their own performance. Standard theory in the arising game predicts that messages do not contain information and thus cannot help improve efficiency.

Our treatments are designed to investigate if information transmission occurs contrary to the theoretical prediction and if the message space available to the applicants impacts the amount of information transmitted. In the baseline treatment, senders submit a natural number to claim how many correct sums they calculated when they previously completed the task. This treatment provides us with a benchmark to quantify the size and frequency of misreports and the magnitude of efficiency gains due to delegation. We observe a large portion of senders who truthfully report (almost 50%) their past performance, while those who lied only exaggerated modestly. The modest and systematic lying implies that messages contain information, which can increase social welfare if extracted and acted upon by delegators. Unfortunately, delegators are not able to use the information contained in the messages. Delegators act as if the messages did not contain information. The delegation option is still welfare increasing, as delegators, who know that they are bad at calculating sums, are more likely to delegate. We observe a significant efficiency gain of 13.6 percent compared to the welfare level that would have resulted without delegation. Applicants with avatars indicating an Asian ethnicity or a different gender than the delegator are more likely to gain delegation. In the absence of information extraction, stereotypes are behaviourally relevant.

In the *Interval* treatment, we introduce coarseness by partitioning the message space into intervals of a fixed length. We find that the amount of information contained in messages is not statistically different from the benchmark treatment. However, compared to the *Number* treatment, delegators make better use of the information by conditioning their delegation decisions on the messages in a profitable way. As a result, we observe an 11.8 percent efficiency improvement over the baseline treatment.

In the *Text* treatment, we enrich the message space by allowing for natural language in messages. In contrast to theory, we observe that free-text messages affect both messaging and delegation behaviour. After analyzing the content of messages, we conclude that senders use different messaging modes depending on their actual past performance. Delegators can extract additional information based the mode of messaging. This leads to an efficiency improvement of 15.8 percent compared to the baseline treatment.¹ Stereotypes do not influence delegation decisions in both the *Text* or *Interval* treatment.

Our paper is related to the growing experimental literature that studies cheap-talk games with asymmetric information. One research program investigates information transmission with imperfectly aligned incentives. Papers typically focus on the class of simplified one-

¹This improvement is not significantly different from the improvement in the *Interval* treatment.

sender, one-receiver cheap-talk games (e.g., Dickhaut et al., 1995; Cai and Wang, 2006; Sánchez-Pagés and Vorsatz, 2007; Wang et al., 2010; Peeters et al., 2013; De Haan et al., 2015). The main finding in this literature is the persistence of non-equilibrium behaviour. Compared to equilibrium, the senders' messages contain too much information, and receivers rely too much on the messages being truthful. Some recent papers (e.g., Goeree and Zhang, 2014; Vespa and Wilson, 2016; Minozzi, 2018; Bayindir et al., 2020; Lohse and McDonald, 2021) investigate how adding an additional sender impacts information transmission. The results heavily depend on the specific setting.

Several recent studies explore the role of communication in other applied settings than in the delegation process. Lundquist et al. (2009) investigate how promises affect buyer and seller interactions in a bargaining game. They find that freely formulated messages lead to the fewest lies and the most efficient outcomes. Charness and Dufwenberg (2011) study if communication can alleviate hidden information problems embedded in principal-agent relationships and find that free-form communication is effective in promoting efficient contracting in some cases but not in others. Serra-Garcia et al. (2011) study information transmission in a public-good game with different levels of precision in language and find that leaders frequently use vagueness to hide inconvenient truths.²

The remainder of the paper is organized as follows. The next Section explains the logic why standard theory predicts no information transition taking place. Section 3 explains the experimental design and lays out the hypotheses. The results and a brief conclusion follow in Sections 4 and 5.

2 No information transmission in delegation with cheap talk

We begin by laying out the logic of why no efficiency-enhancing information transition occurs in the equilibrium of our cheap-talk game. Suppose that there are a delegator and two potential applicants. All three privately know their productivity for a specific task since they have performed it before. The delegator needs the task performed and can either do it herself or delegate the task to one of the applicants. The higher the productivity of the person performing the task, the higher is the revenue for the delegator. Applicants prefer to be chosen for the task as they then receive a bonus from the delegator. The delegator

²Our study also contributes to the recent experimental research on preferences for maintaining control in delegation decisions by showing that in our setting, delegation rates are comparatively high. See Owens et al. (2014), Hausfeld et al. (2020), Bartling et al. (2014), and Fehr et al. (2013) for papers demonstrating that the preference for control often leads to very low delegation rates.

asks the applicants to send a costless, non-verifiable message. Upon receiving the messages, the delegator decides to whom to delegate or to perform the task herself. In this setting, efficiency-enhancing information transition is not possible, as low productivity senders have an incentive to imitate messages of high-productivity senders, which then implies that a rational delegator should disregard the messages. This result is independent of the message space applicants can use when sending their messages.

More formally, the argument goes as follows. Suppose there were two messages a and b , where the delegator is more likely to delegate to a person sending message a . Then nobody would ever want to send a message b . This implies that in equilibrium no two messages will be sent with positive probability that lead to different likelihoods of gaining delegation. In order to make sure that the delegator finds it optimal to make the same delegation decision regardless of the message sent, she must assign the same expected underlying productivity to all messages, which implies that messages cannot be informative. The interested reader can find a formal treatment of the problem in the Appendix A.1.

3 Experimental Design

We present the design of an experiment by which we can test both the theoretical predictions and a variety of alternative hypotheses. We require a simple environment that allows us to investigate if and which kind of messages improve social welfare in a delegation environment. We decided to use a real effort task that requires non-negligible mental effort and inert ability instead of a task with stated effort and induced effort cost. This choice was guided by the aim to capture the salient factors of the delegation of tasks in real-world environments. The real-effort task chosen is the repeated addition of five two-digit numbers (see Niederle and Vesterlund, 2007), which has been shown to allow for plenty of variation in individual performance. The use of a real-effort task has advantages and disadvantages. By using actual ability in a real-effort task instead of inducing types, we surrender control over the distribution of types and prior beliefs. In exchange, we gain salience, as the messages subjects send and the beliefs they hold are concerned with the task they have experienced first-hand.

The participants complete the seven-minute addition task twice. After choosing an avatar that best represents gender and ethnicity and answering some priming questions (details on this follow below), participants are asked to complete the task for the first time. Completing the task allows participants to learn their abilities and form beliefs about their abilities relative to others. Participants are paid a piece rate for the number of sums solved in seven minutes. At this stage, participants do not know what is to follow.

Once the addition task has been completed for the first time (*Task 1*), new instructions explaining the delegation stage are distributed. Participants are randomly matched into groups of three. One participant of the group is randomly assigned the role of a delegator. The other two will play the role of applicants. The applicants are asked to send a message to the delegator, which signals their ability in the addition task. Depending on the treatment, the message can either be a number (sums solved in *Task 1*), one of a range of predetermined intervals (of sums solved in *Task 1*) or free text up to 400 characters. The delegator sees the avatars and messages of the applicants and then decides to delegate the addition task to one of the remaining two players or not. Delegation implies that the delegator’s performance in the **already completed** task is being replaced by the applicant’s performance in the **future** task.

Once the delegation decision has been made, all participants perform the addition task a second time (*Task 2*). The payment scheme is identical to that used the first time round. The number of solved sums is multiplied by a piece rate. All three participants, regardless of their role, receive the payment for their performance in *Task 2*. The payoffs for *Task 1* depend on the delegator’s decision. In case no delegation has taken place, the payoff for *Task 1* is the piece-rate payment for the participants’ own performance. If delegation has taken place, the chosen applicant receives a bonus paid by the delegator in addition to the own piece-rate pay. In return, the delegator’s performance in *Task 1* is replaced by that of the chosen applicant in *Task 2*. Hence, a delegator increases overall efficiency if she delegates to an applicant who solves more sums in *Task 2* than she did in *Task 1*. She improves her own payoff if the higher performance of the chosen applicant at least covers the bonus she has to pay. Table 3.1 summarizes the payoffs in case delegation takes place.

Table 3.1: Payoffs if the delegator delegates to applicant A

	Payoff for <i>Task 1</i>	Payoff for <i>Task 2</i>
Delegator <i>D</i>	rate* <i>A</i> ’s perf in <i>Task 2</i> - bonus	rate* <i>D</i> ’s perf in <i>Task 2</i>
Applicant <i>A</i>	rate* <i>A</i> ’s perf in <i>Task 1</i> + bonus	rate* <i>A</i> ’s perf in <i>Task 2</i>
Applicant <i>B</i>	rate* <i>B</i> ’s perf in <i>Task 1</i>	rate* <i>B</i> ’s perf in <i>Task 2</i>

Our design has a variety of features. Paying a piece rate in *Task 2* even if an applicant achieves delegation prevents moral hazard. Potential income effects are minimized by randomly choosing one task for payment. Moreover, we remove the potentially confounding factor of delegators’ uncertainty about their own ability by having delegation replace their past instead of their future performance.

In addition to the direct financial incentives, we introduce a social dimension by using a set of priming questions and avatars to make ethnic and gender cues salient. The priming questions sensitize subjects’ perception of gender differences through activating natural social identities.³ After priming, eight avatars representing males and females of four ethnic backgrounds are shown. Participants are asked to choose one that best represents them.⁴ Groups of three are formed randomly with the restriction that they are not single-sex. The delegator is randomly chosen from the gender that occurs twice in the group. We obtain groups with either a female or male delegator who can always choose between a male and a female applicant. Salient gender and ethnicity cues from the avatars allow for gender and ethnicity stereotypes to potentially influence delegation decisions. Table 3.2 summarizes the design and the rationale for its components.

3.1 Treatments

Our study consists of three treatments. In the baseline treatment, players send a precise message of the form “The number of correct answers I had in the addition task was x .” We call this baseline the *Number* treatment. The message space in the *Number* treatment is a copy of the type space and consists of non-negative integers up to a maximum. In this treatment, the message space has exactly the same cardinality as the type space, which is the minimum size for full information transmission to be technically possible.

In our second treatment, we partition the message space into intervals to study situations where it is impossible to communicate precisely. In the *Interval* treatment, players send a message of the form “The number of correct answers I had in the addition task was in the range from a to b .” The intervals senders can choose are predetermined. There are different options about how to design the intervals. The intervals could be fixed relative to the natural numbers (i.e. everybody faces the same intervals, say 0 to 3, 4 to 7, etc.) or fixed relative to the performance of the delegator. Fixing the interval to natural numbers has the disadvantage that the inference problem the delegators face differs according to their past performance. Hence, we decided to fix the interval relative to the delegator’s past performance.

The delegators’ performance is located on the upper end of an interval, and the length of

³This technique has been commonly used in social psychology (eg: Bargh and Pietromonaco 1982; Aronson et al. 1998). The questions used for priming were: Who are your role models? Please list some of your favourite food. What kind of movie do you prefer to watch for a casual weekend? Please describe your ideal way of relaxation. What is the major responsibility in your current life? What is your childhood dream? Please name some achievements that you are proud of.

⁴Subjects do not know how they will be used in order to prevent strategic choices as observed in Charness et al. (2020).

Table 3.2: The stages of the experiment

Action	Actor	Payoff consequence	Rationale
1. Priming questions	Everybody	none	Activate attitudes toward gender and ethnicity
2. Choice of avatar	Everybody	none	Make gender and ethnicity of participant identifiable to others
3. Summation <i>Task 1</i>	Everybody	piece rate	Measure ability and allow for belief formation
4. Sending message	Applicants	free	Treatment variation: different message spaces
5. Delegation	Delegators	Delegator's Task 1 performance is replaced by chosen applicant's Task 2 performance; Applicant receives bonus	Replacing Task 1 performance controls Delegator's beliefs about own productivity
6. Summation <i>Task 2</i>	Everybody	piece rate	Piece rate for everybody prevents moral hazard

the interval equals the number of questions (i.e. four) required to pay for the bonus.⁵ This design implies that delegation is profitable whenever the actual performance of the chosen applicant lies in an interval at least two above that of the delegator. This is the same for all delegators. An advantage of this design is that the location of the delegator's performance within an interval can be eliminated as a cause for behavioural differences across delegators. In the *Interval* treatment, the message space is of lower cardinality than the type space, and full information transmission is technically impossible.

In our final treatment, we go the other way. Here, our message space exceeds the size of the type space, as we allow for free-form text messages. In the *Text* treatment, applicants can write a text message of up to 400 characters. Using natural language allows for more than just sending a message that communicates a number or a range of numbers. In standard game

⁵Suppose the performance of a delegator is 16, then the messages the applicants can send consists a set of consecutive intervals $\{[1, 4], [5, 8], [9, 12], [13, 16], [17, 20], \dots\}$.

theory, message spaces that are at least as large as the type space are sufficient for maximum information transmission if preferences are aligned. However, previous experimental studies (Charness and Dufwenberg, 2006; Chen and Houser, 2017) find that free-form communication in written messages can improve cooperation by fostering trust in games with hidden action. The *Text* treatment allows us to test if the theoretically superfluous size of the message space helps information transmission.⁶

3.2 Hypotheses

Theory predicts no information transmission in all treatments, which is at the heart of our Hypotheses.

Hypothesis 1: Senders' messages are not correlated with actual performance in any treatment.

Hypothesis 2: Delegators in all treatments disregard messages and base their delegation decisions only on their prior beliefs.

Hypothesis 3: Overall efficiency does not differ across treatments.

Hypothesis 4: Avatars representing gender and ethnicity have no impact on delegation decisions.

In light of the recent experimental literature on lie aversion (see Abeler et al., 2019, for a metastudy) or bounded rationality in cheap-talk games (e.g., Kawagoe and Takizawa, 2009), we expect that messages will contain information. Messages contain information whenever they are correlated with actual past performance. Then a delegator who knows or guesses the link between messages and past performance can update her beliefs after receiving the messages. Following the literature on lying aversion (Gneezy, 2005; Gneezy et al., 2018), we expect systematic lies of limited size, which should result in messages containing information about performance.

Hypothesis A1: Messages contain information about the senders' productivity. The degree of information differs across treatments.

We do not have a clear prior on whether the coarser message space in the *Interval* treatment leads to more or less information in messages than the *Number* treatment. We conjecture that more information will be contained in messages in the *Text* treatment. The

⁶Note that we only vary the messages participants can send across treatments. Participants, who make themselves available for delegation, have to send a message in all treatments.

opportunity to send free text messages allows choosing a style of the message and not only the message itself. If applicants with different levels of past performance use different communication styles, then not only the message content but also the style contains information. Earlier studies show that humans avoid lying by using imprecise messages (Serra-Garcia et al., 2011; Wood, 2016), are able to express their social motivations (Cason and Mui, 2015), and make promises in free-form communication (Charness and Dufwenberg, 2006; Chen and Houser, 2017; Vanberg, 2008). Hence, we expect messages to contain the most information in the *Text* treatment.

Hypothesis A2: Delegators make use of at least some information contained in the messages.

Messages containing information is not sufficient for information transmission. The delegators must also be able to extract the information. The degree of information extraction will depend on the delegator trusting in a relation between messages and performances to exist and on how accurate her guess about the nature of the relation is. If lies are modest, then information extraction depends largely on how much trust delegators put in the applicants' truthfulness.

Past research suggests that we should expect more trust in treatments with restricted message spaces (Cai and Wang, 2006; Wang et al., 2010). Bounded rationality might play an important role here. Wood (2016) finds that the increased complexity of a vague message space makes subjects less sophisticated in their reasoning and leads them to trust more in messages compared to a message space that only allows precise messages. Therefore, we expect more information transmission in the *Interval* treatment than in the *Number* treatment. We further conjecture that the uncertainty of the location of the actual performance in a messaged interval might be more salient than the uncertainty about its truthfulness. Evidence suggests that people do not take strategic uncertainty sufficiently into account when it is confounded with risk (Huberman and Rubinstein, 2001).

Finally, we conjecture that subjects will make good use of messages in the *Text* treatment by picking up cues contained in the chosen format of messages. Research on deception shows that people are able to read cues and spot a lie in informal written communication (Chen and Houser, 2017) and face-to-face communication (Belot et al., 2012; Konrad et al., 2014). In sum, we conjecture the ranking of the degree of information transmission across treatments to follow the order $Number < Interval < Text$. As efficiency depends predominantly on information extraction if the information content of the message is comparable, we formulate the following alternative hypothesis on efficiency:

Hypothesis A3: Social welfare is greatest in the *Text* treatment followed by the *Interval* treatment and the *Number* treatment.

Research in psychology has shown that there are relatively stable ethnic and gender stereotypes (e.g., Garg et al., 2018), such as “Asians are good at maths and women are not.” Additionally, there is an underrepresentation of females (e.g., Ortiz-Ospina, 2018) and ethnic minorities in leadership positions (Eagly and Chin, 2010). It is possible that these two facts are causally linked. Hence, we hypothesize that stereotypes activated by avatars and priming questions impact delegation decisions and negatively impact efficiency in our environment.

Hypothesis A4: The likelihood of achieving delegation is influenced by the avatar of the applicant.

4 Results

Treatments were programmed using Z-tree (Fischbacher, 2007) and carried out at AdLab, the Adelaide Laboratory for Experimental Economics. Participants were recruited with the help of the online system ORSEE (Greiner, 2015). Overall, 342 subjects, which were predominantly university students, participated in the experiments and earned on average 15.4 Australian dollars.

The bonus for gaining delegation was 4 Australian dollars.⁷ Subjects were also asked two belief-eliciting questions after they completed *Task1*. Firstly, we asked participants for their beliefs about the percentage of people who solved more correct sums in the session. This elicitation was incentivized by 5 Australian dollars paid to the person with the guess closest to the truth. The second question elicited the participants’ beliefs about how well they will be able to perform the same task the second time.

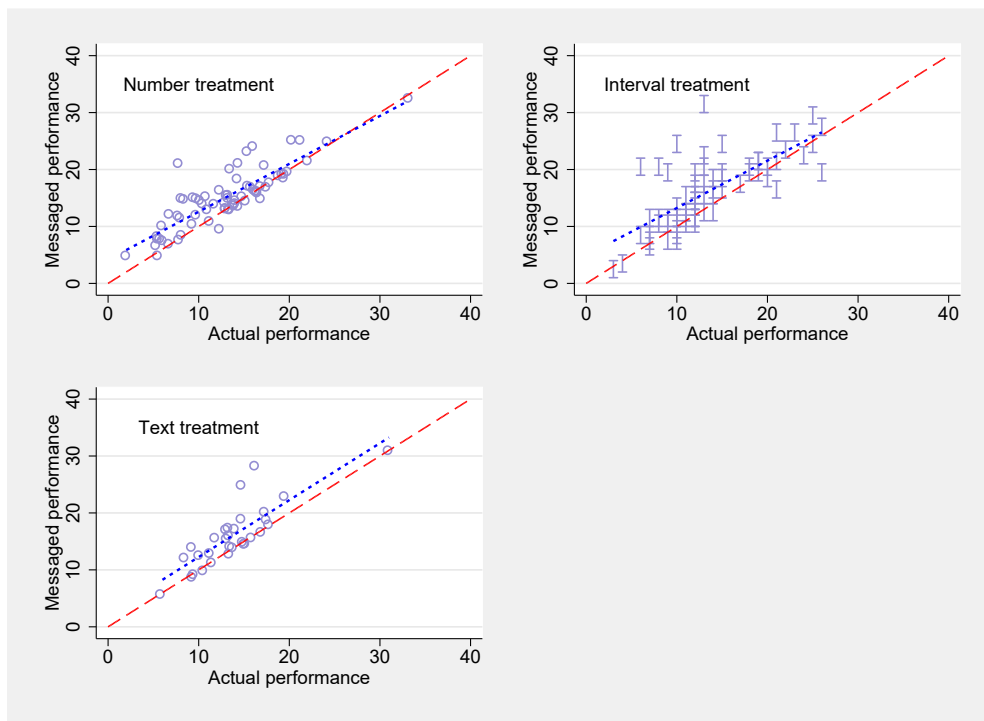
We ran five sessions for each treatment with between 15 and 27 subjects per session. Overall, we had 117 individuals in the *Number* treatment, 117 in the *Interval* treatment, and 105 in the *Text* treatment. The majority of participants (324 out of 339) chose avatars that coincided with their real gender. On average, participants calculated 13 correct sums in *Task 1* and 16 in *Task 2* and correctly anticipated their performance increase. We also observed some aggregate overconfidence. On average, participants believed that only 38 percent of the population outperformed them. There is no statistical difference in performances and beliefs between male and female participants (two-sided rank-sum test, $p > 0.1$).

⁷We used a pilot to estimate the distribution of the number of sums our population can solve and then set the bonus such that delegation is profitable for the delegator in about 50 percent of groups.

4.1 Information content of messages

There are 71, 70, and 65 available sender-subjects in the *Number*, *Interval*, and *Text* treatments, respectively.⁸ Figure 4.1 plots the applicants' messaged performance against their actual performance. Note that numerical information in messages are represented as points in the *Number* treatment, intervals in the *Interval* treatment, and a selection of points in the *Text* treatment (only subjects who mentioned a number are represented). Hypothesis 1, drawn from standard game theory, implies that messages and actual performance should not be correlated. Rejecting Hypothesis 1, we observe a very regular pattern. The majority of the data points lie on or slightly above the 45-degree line, indicating truthful reports or moderate but systematic lies. In all three treatments, there is considerable information contained in the messages.

Figure 4.1: Messaged performance vs. actual performance



Note: Sample size: 71 in the *Number* treatment, 70 in the *Interval* treatment and 30 in the *Text* treatment.

The intuition gained from a visual inspection is confirmed by regression models that con-

⁸20 subjects out of 226 in total (7 in the *Number* treatment; 8 in the *Interval* treatment; and 5 in the *Text* treatment) did not enter the cheap-talk stage, as they preferred not to be considered for delegation. When asked for a reason, some low-performance applicants stated to prefer staying silent rather than lying. High skill participants who did not enter stated that the bonus did not compensate them enough for handing a very high payoff to somebody undeserving.

trol for subjects’ prior beliefs about their relative performance and individual characteristics. Table A.1 in the Appendix presents estimation results from OLS regressions in the *Number* and *Text* treatments and an Interval regression in the *Interval* treatment. Beyond the insights from plotting, we learn that the coefficients on the prior belief variable for all treatments are close to zero and not significant. This suggests that subjects did not try to improve their delegation chances by increasing the size of their lies whenever they believed to be relatively noncompetitive.⁹

Next, we compare the quality of the information in messages for the two structured message spaces, i.e., *Number* vs *Interval*. The difficulty of comparison stems from the different measurement units between the two treatments. While the units in the treatments differ, the messages in both treatments can be ranked. Hence, we calculate the Spearman rank correlation coefficient. We obtain a correlation of 0.818 in the *Number* and 0.755 in the *Interval* treatment.¹⁰ The correlation seems to be slightly larger in the *Number* treatment, but the difference is not statistically significant ($p > 0.1$, 1000 bootstraps). For robustness, we map numbers into corresponding intervals and calculate the size of a lie as the difference between the messaged performance band and the actual performance band. The two distributions for the two treatments are not statistically different from each other (Kolmogorov-Smirnov test, $p > 0.1$). Figure A.1 in the Appendix shows histograms of the size of lies for the two treatments.

Lastly, we analyze the information contained in the kind of message an applicant sends in the *Text* treatment. We observe the following regularities: 31 subjects mentioned their precise performance in *Task 1*; 25 subjects expressed their competence by mentioning their mathematical ability or their math-related background; 24 subjects promised either better future performance, more effort or more earnings to the delegator; 20 subjects tried to reduce the social distance by using words such as “we”, “trust”, “believe”, “help” and their synonyms. There are 11 messages that do not belong to any of the categories above and only contain babble that is unrelated to the situation. Accordingly, we use five dummies, i.e. *Number*, *Ability*, *Promise*, *Trust*, and *Babble*, to record if a particular way of communication was observed in a text message. Note that these categories are not exclusive. Many subjects used multiple kinds of messages in their communication. The only exclusive category is babbling. Table 4.1 reports the average performance conditional on a particular messaging type used. For later use, we also report the frequency of being chosen for delegation in each case.

It is noticeable that subjects who were babbling performed significantly worse than those

⁹The propensity to lie is greater, the lower the performance is, and there are no treatment differences in the propensity of lying. See Table A.2 in the Appendix for the Probit regression that shows that.

¹⁰The Pearson correlation coefficient on the numerical data in the *Number* treatment is quite similar (0.858).

Table 4.1: Average performance conditional on verbal cues

Cues	Average Performance	Chosen Frequency	Number of Observations
Babble	9.545	0.000	11
No babble	13.463	0.241	54
Number	13.774	0.355	31
Not numeric	11.912	0.059	34
Promise	14.167	0.250	24
No promise	12.000	0.171	41
Trust	14.100	0.150	20
No trust	12.222	0.222	45
Ability	11.840	0.280	25
Not ability	13.400	0.150	40

who did not (rank-sum test, $p = 0.02$, two-sided). We also observe that subjects who mentioned their past performance performed significantly better than those who did not (rank-sum test, $p = 0.03$, two-sided). Subjects who made a promise had performed marginally better than those who did not (rank-sum test, $p = 0.07$, two-sided). However, verbal cues such as signalling trust or commenting on general ability were not informative with respect to average actual performance (rank-sum test, $p > 0.1$, two-sided).

Result 1 Messages contain a substantial amount of information in all three treatments.

Result 2 The amount of information contained in the messages is similar in the *Number* and *Interval* treatments.

Result 3 Free-form text messages contain some additional information as low-performance senders tend to either exaggerate less or avoid outright lying by babbling. Observing if somebody sent a message mentioning past performance or a promise contains some further information.

4.2 Delegation decisions

Delegators who knew their past performance were asked to estimate the percentage of participants that outperformed them and received at most two cheap-talk messages. We define a message as potentially “profitable” if the messaged performance is strictly greater than the delegator’s own performance plus the bonus. The variable “Message” counts the

number of profitable messages. Table 4.2 reports the estimated coefficients and average marginal effects from probit models that predict delegation.

Table 4.2: Probit regressions on delegation decisions

Treatments Dependent Var.	<i>Number</i>		<i>Interval</i>		<i>Text</i>	
	Coefficient	A.M.E.	Coefficient	A.M.E.	Coefficient	A.M.E.
Prior Belief	0.022* (0.012)	0.006** (0.003)	0.006 (0.008)	0.002 (0.002)	0.015 (0.014)	0.003 (0.003)
One Message (vs. None)	0.015 (0.611)	0.004 (0.157)	0.952* (0.562)	0.260* (0.150)	2.334*** (0.678)	0.677*** (0.133)
Two Messages (vs. None)	0.780 (0.584)	0.241 (0.188)	1.493** (0.635)	0.465** (0.192)	–	–
Controls	YES		YES		YES	

Note: Coefficient with standard error in parentheses.

***, **, * denote significance at $p=0.01$, 0.05 and 0.10 respectively.

[1] Dependent variable is a dummy variable, i.e.1=delegation and 0=no delegation. Constant term is included.

[2] Sample size: 39 in the *Number* treatment, 39 in the *Interval* treatment, and 34 in the *Text* treatment.

[3] Controls for individual characteristics include gender, and math training.

The regressions show that messages and beliefs significantly impact delegation behaviour in some treatments but not in others.¹¹ In line with Hypothesis 2, subjects ignore messages and resort to their initial belief about their relative performance in the *Number* treatment. A ten percentage-point increase in the subjective belief regarding the percentage of people who can calculate more sums increases the delegation likelihood by about 6%. The fact that the number of potentially profitable messages is irrelevant once we control for beliefs implies that delegators are not able to extract information from the messages in the *Number* treatment. In contrast, subjects made good use of messages in the remaining two treatments, just as our alternative Hypothesis A2 conjectured. The average frequency of delegation significantly increases by 47% if subjects receive two profitable messages in the *Interval* treatment and by 68% if subjects receive one profitable message in the *Text* treatment.¹²

Delegators seem to rely on the messages even a bit too much. Despite the fact that the messages contain information, there is still some noise. Therefore, rational updating

¹¹Note that the effect of profitable messages observed in the regression can be a confound of the causal effect of observing profitable messages and unobserved attributes of people who receive the profitable messages. Ideally, one would isolate the causal effect by correcting for selection. The performance of the delegator could be used as an instrument for selection. Our sample is too small to precisely estimate such a selection model. Similarly, just controlling for e.g. past performance is problematic due to collinearity. Figure A.2 in the Appendix, where we plot delegation fractions conditional on profitable messages separately for high and low performing delegators, provides some support for a causal effect to exist.

¹²Note that in the *Text* treatment, the impact of two potentially profitable messages was not estimable as this only occurred once and therefore perfectly determined the delegation decision.

implies that the prior belief should still have an impact on the decision. Our regressions show that it does not. Table 4.1 where we summarized the information contained in the communication style shows that delegators could extract some of it. The third column contains the delegation fractions depending on a message containing a particular mode of communication or not. Delegators never delegated to people who babbled and more often delegated to applicants who mentioned a precise past performance than those who did not. On average, this delegation behaviour is consistent with delegators extracting and responding to the information contained in the applicants' message styles.

Result 4 Information contained in messages is not transmitted in the *Number* treatment. Information is transmitted in the *Interval* treatment and even more so in the *Text* treatment.

Result 5 Subjects in the *Interval* treatment are more likely to delegate if they receive two potentially profitable messages.

Result 6 Subjects in the *Text* treatment never delegate if they receive babbling messages and typically delegate if they receive at least one profitable message.

4.3 Stereotypes

The eight avatars that we used span two dimensions: gender and ethnicity. There are four male and four female avatars representing four ethnicities, each. As a base, there is an ethnicity-neutral avatar, which could be White, Hispanic, Mediterranean, Middle Eastern, or Indian. About 47% chose the neutral avatar. Then there are three strongly ethnic avatars for each gender, which contain stereotypical elements for Asian (25%), White (22%) and Black (6%). To investigate the impact of stereotypes, we run Probit regressions where gaining delegation is the dependent variable. The independent variables of interest are gender and ethnicity dummies. We include a dummy that indicates if the applicant in question has sent the highest unique profitable message. Adding this dummy implies that we assume that a person who is not guided by stereotypes does not delegate to anybody else than to the person with a message that is higher than that of the other sender and profitable if true. Our results are robust to the use of other variables to control for messages.

Table 4.3 reports average marginal effects for two estimated specifications for each treatment. We find that stereotypes do not play a large role overall. Messaging has a much larger impact on all treatments. Sending the uniquely highest profitable message increases the probability to obtain delegation by around 0.21 in *the Number* and *Interval* treatments, and up to 0.76 in the *Text* treatment. Stereotypes for gender or ethnicity only play a role

Table 4.3: Probit regressions of the likelihood of being delegated to

Treatments	<i>Number</i>		<i>Interval</i>		<i>Text</i>	
Highest message	0.212**	0.227**	0.217**	0.197**	0.764***	0.794***
	(0.107)	(0.097)	(0.104)	(0.097)	(0.122)	(0.105)
Female	-0.086		-0.04		0.125*	
	(0.075)		(0.096)		(0.073)	
Delegator same gender		-0.164**		-0.124		-0.013
		(0.066)		(0.078)		(0.058)
Ethnicity dummies (non-specific is base)						
Asian	0.191*	0.249*	0.003	0.026	0.054	-0.002
	(0.114)	(0.139)	(0.100)	(0.084)	(0.085)	(0.075)
White	0.104	0.113	-0.063	-0.086	0.070	0.046
	(0.091)	(0.078)	(0.092)	(0.081)	(0.084)	(0.084)
Black	0.035	0.062	-	-	0.001	-0.022
	(0.095)	(0.083)	-	-	(0.049)	(0.048)

Note: Average Marginal Effects with standard error in parentheses.

***, **, * denote significance at $p=0.01$, 0.05 and 0.10 respectively.

[1] Dependent variable is a dummy variable, i.e. 1=delegated to and 0=not delegated to. Constant is included.

[2] Sample size: 78 in the *Number* treatment, 76 in the *Interval* treatment; and 70 in the *Text* treatment.

in the *Number* treatment. This is consistent with the observation that delegators in the *Number* treatment were not able to make good use of the messages as they were too afraid of facing lies. This leaves room for using other available information for the delegation decision. In the *Number* treatment, participants with avatars indicating an Asian ethnicity have an increased probability of being delegated to. While we don't find outright gender discrimination, indirectly gender still plays a role in that treatment. Everything else equal, delegators are significantly more likely to delegate to the opposite gender. While this seems counter-intuitive, there is a simple explanation. The participants who delegate are those who believe that their ability is low compared to the other participants. If these participants wrongly extrapolate from their own gender-performance relation to others, then they should delegate more often to the opposite gender. This effect is mainly driven by women delegating to men, as there is not a single case where a woman delegates to a woman. This is consistent with the findings in Bordalo et al. (2019) who find strong evidence that stereotypes play a role in females' beliefs only.¹³ We also find that females secure delegation more often than males in the *Text* treatment. This is not a stereotype effect and can be explained by females making more claims about their mathematical abilities, which yields higher delegation rates.

¹³We checked for similar effects of delegators delegating to their own or a different ethnicity but did not find any.

Result 7 Stereotypes only influence the delegation decision in the *Number* treatment, where delegators are not able to make good use of the information contained in the messages.

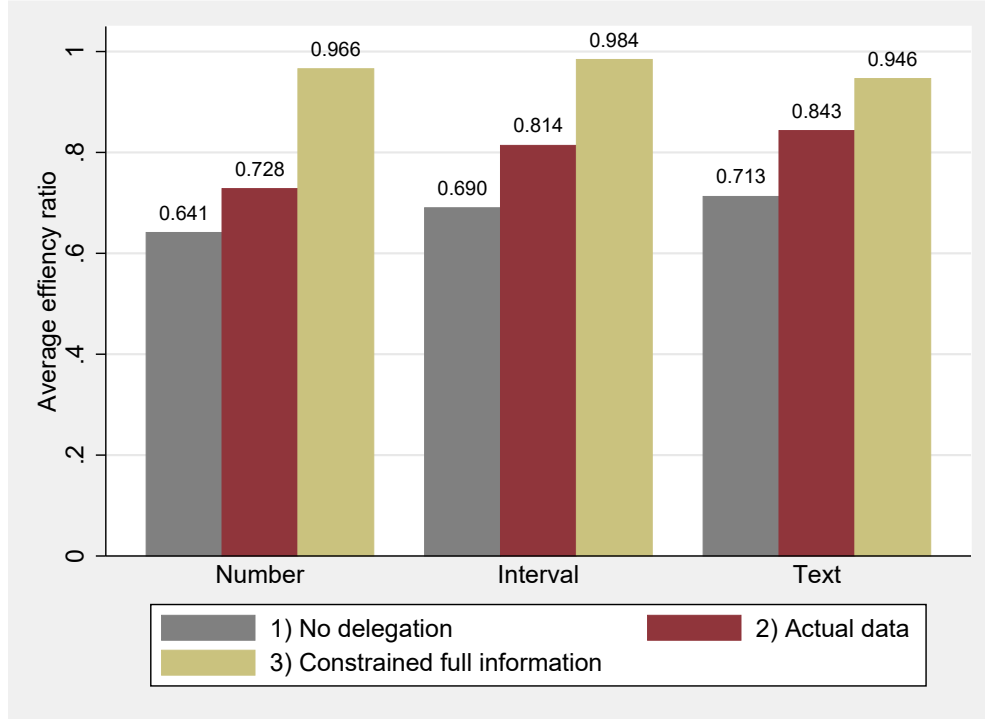
4.4 Social welfare

It remains to be checked if the varying degree of information transmission translates to efficiency differences. The average observed delegation frequency is 0.282 in both the *Number* and *Interval* treatments and 0.371 in the *Text* treatment. Recall that we calibrated the value of the bonus to target approximately 50% of rational delegations if individual performances were known. Due to asymmetric information embedded in the process of delegation, the observed average delegation frequencies across treatments are significantly below 50% (two-sided binomial probability test, $p < 0.01$). In what follows, we construct an efficiency ratio to measure the relative degree of efficiency in the treatments. Delegation can improve social welfare only if there is one player, whose performance in *Task 2* is better than the delegator's performance in *Task 1*. If the delegator always allocates the task to the most capable person, delegation is fully efficient. We normalize the efficiency measure by dividing the performance implemented by the delegator by the best performance in the group. In this way, we can compare the average efficiency ratios across treatments. The ratio can be interpreted as the average fraction of the available surplus extracted. If the delegator decision resulted in the task being performed by the person solving the most sums, all surplus is extracted and the ratio is one. If the delegator implements the performance of somebody who only solves half as many sums as the best person in the group, the ratio becomes 0.5.

We also compare the actual average efficiency ratio to two counterfactual scenarios to assess the efficiency gain through delegation based on communication. In the no delegation benchmark, we calculate the average efficiency ratio that would have occurred if delegation had not been possible. Our second benchmark calculates the constrained optimal efficiency ratio. Note that a delegator has to pay a delegation bonus to the chosen applicant whenever she delegates. Hence, we exclude cases where delegation would improve social welfare (i.e. there is an applicant with higher performance than the delegator), but the difference is not high enough such that it pays for the fully informed delegator to delegate. The constrained optimality benchmark represents the efficiency ratio that would have resulted from rational delegation under perfect information.

Figure 4.2 shows the average efficiency ratios by treatment calculated for 1) the no delegation scenario, 2) the actual experimental data, and 3) the constrained full information scenario. As expected, the actual efficiency levels in both the *Interval* and *Text* treatments are higher than in the *Number* treatment. The observed difference is significant for the *Text*

Figure 4.2: Average efficiency ratio in four scenarios by treatments



Note: [1] Efficiency ratio is calculated at group level and averaged over all groups within each treatment.
 [2] Sample size is 39 for the *Number* treatment; 39 for the *Interval* treatment and 35 for the *Text* treatment.
 [3] We only consider 206 available players' performance in Task 2 in the efficiency calculation for the three cases.

treatment ($p < 0.024$, rank-sum, two-sided) and weakly significant for the *Interval* treatment ($p < 0.073$, rank-sum, two-sided). Compared to the no-delegation benchmark, all treatments achieve a significantly higher efficiency ratio ($p < 0.05$, sign-rank, all two-sided). Hence, we reject Hypothesis 3 and find support for alternative Hypothesis A3. However, we still observe a significant efficiency loss due to asymmetric information compared to the constrained full information scenario ($p < 0.01$, sign-rank, two-sided) in all treatments.

Result 8 Allowing for delegation and one-way communication improves welfare in all treatments. Free-form communication achieves significantly higher efficiency than precise communication. Coarse communication attains a weakly significantly higher efficiency than precise communication. With cheap talk alone the constrained full information efficiency level can not be attained.

5 Conclusion

Our paper experimentally investigated if and how one-way cheap-talk communication can facilitate efficient delegation. We were particularly interested in testing whether the available message space impacts information transmission, where in theory, misaligned preferences do not allow for information to be transmitted in equilibrium. We found that information transmission occurred and that the message space applicants were allowed to use affected both the amount of information contained in applications and the efficiency gains from the delegation. Overall, allowing for free-text messages yields the best result. This is in contrast to the game-theoretical prediction that in cheap-talk and signalling games, nothing can be gained from extending the size of the message space beyond that of the type space. Applicants convey information not only by what they say but also by how they say it, and delegators can extract at least some of this information.

While our study shows that cheap talk can increase efficiency in a delegation scenario, the extracted surplus is still well short of what could be achieved without information frictions. We conclude that for many delegation situations, the delegator is likely better off requiring the applicant to send a costly signal (such as a lengthy proposal, response to selection criteria or the completion of training or education). The resulting signalling game will probably improve information transmission. However, this comes at a cost, as the signalling wastes resources and dissuades good applicants from applying. We leave the analysis of the circumstances under which signalling is more efficient for future research.

A Appendix

A.1 Cheap-talk delegation game

We consider a sequential move game with one principal, two agents and incomplete information. There are one delegator and two senders, indexed by $i \in I = \{D, 1, 2\}$. Types $\theta_1, \theta_2, \theta_D \in \Theta$ are independently drawn from the same distribution f with support $[0, \theta^{max}]$ and mean $\bar{\theta}$. In this context, a type refers to individual productivity. After observing their own type θ_i , two senders simultaneously send a message $m_i \in M_i$ to the delegator. M_i denotes the message space for sender i , which can be of smaller, equal or greater cardinality than the sender's type space.

The delegator, after learning his own productivity θ_D and the two messages m_1, m_2 , decides if she wants to conduct a project herself or to delegate to one of the two senders, i.e. $d \in A = \{0, 1, 2\}$. If the principal decides to conduct her own project, i.e. $d = 0$, then

the delegator's own productivity θ_D determines the project outcome. If sender i is chosen for delegation, i.e. $d = i$, then sender i receives a bonus, $B > 0$, from the delegator and the outcome of the project is determined by the chosen player's productivity θ_i . To make the question interesting, we assume that the least productive delegator is at least indifferent to delegate under the prior distribution, i.e. $\bar{\theta} \geq B$. The payoff functions, once delegation has occurred and the types are revealed are as follows:

$$U_D = \begin{cases} \theta_d - B & \text{if } d \neq 0 \\ \theta_D & \text{if } d = 0 \end{cases},$$

for the delegator, and

$$U_i = \begin{cases} B & \text{if } d = i \\ 0 & \text{if } d \neq i \end{cases},$$

for the senders $i = 1, 2$.

Together, the requirements of consistency of beliefs and sequential rationality in a Perfect Bayesian Nash Equilibrium imply that all messages sent in equilibrium necessarily need to lead to the same probability p_i for a sender to be delegated to.

Proposition 1. *If two messages m'_i and m''_i are both sent in equilibrium, then $p_i(m'_i, \cdot) = p_i(m''_i, \cdot)$ follows.*

Proof. In what follows we will show that in a Perfect Bayesian Equilibrium information transmission that leads to message-dependent delegation decisions is impossible. In equilibrium, we require consistent beliefs. In our case this means that delegators form beliefs about the expected future project outcomes depending on to whom they delegate, i.e. $\mu_1(m_1)$ and $\mu_2(m_2)$. Note that rational beliefs about the productivity of one sender cannot depend on the message of the other sender, as the other sender has no other information than the prior. The beliefs have to be formed using Bayes' Rule and in equilibrium have to be compatible with equilibrium messaging behaviour. Sequential rationality for the delegator implies:

$$d^* = \begin{cases} i & \text{if } \mu_i > \mu_{-i} \wedge \mu_i > \theta_D + B \\ D & \text{if } \max\{\mu_1, \mu_2\} < \theta_D + B \end{cases}$$

Note that a strategy for sender i is a mapping from her type space Θ_i into her message space M_i . Now suppose that a sender is assessing her probability p_i of being delegated to. In equilibrium, a sender will maximize this probability assuming that the other sender and

the delegator play their equilibrium strategies. Moreover, the sender in equilibrium will take into account how the delegator updates her beliefs. Therefore in equilibrium for the senders, the following will hold

$$m_i^*(\theta_i) = \arg \max_{m_i} p_i \{m_i, m_{-i}^*(\theta_{-i}); \mu_i(m_i), \mu_{-i}(m_{-i}); F\} \forall \theta_i \in \Theta_i.$$

Observe that p_i does not depend on θ_i . This implies that all messages that are sent in equilibrium necessarily need to lead to the same delegation probability. Messages that lead to a lower probability cannot be sequentially rational and therefore cannot be part of an equilibrium. \square

This proposition either implies that all messages m_i sent in equilibrium induce the same posterior expected ability μ_i or that induced differences in beliefs do not lead to differences in delegation behaviour for any possible delegator type. Our assumption that the abilities of the delegator and the senders are independently drawn from the same distribution rules out the latter, as long as the least productive delegator is prepared to delegate if she holds the prior beliefs.

Proposition 2. *If $\bar{\theta} \geq B$, then for any two messages m'_i and m''_i sent with positive probability in equilibrium induce assessments $\mu(m'_i) = \mu(m''_i) = \bar{\theta}$ holds.*

Proof. To see this, first assume that $p_i(m_i, \cdot)$ is positive but below unity for all equilibrium messages sent by i and $\mu_i(m''_i) > \mu_i(m'_i)$ for two messages sent with positive probability. The probability to be delegated to is given by

$$p_i = \text{prob} \{ \mu_i(m_i) > \theta_D + B \} \cdot \text{prob} \{ \mu_{-i}(m_i) > \mu_{-i} | \mu_i(m_i) > \theta_D + B \}.$$

Note that the first probability is strictly increasing in μ_i , while the second probability is positive and non-decreasing, which implies that $\mu_i(m''_i) > \mu_i(m'_i)$ and $p_i(m'_i, \cdot) = p_i(m''_i, \cdot)$ are not compatible for positive p_i . It remains to be shown that different believed abilities together with delegation probabilities of zero or one are not possible in equilibrium. For delegation probabilities of one, all types of delegators would need to find it optimal to delegate. This is not possible since delegation is a dominated strategy for a delegator of type θ^{max} . Now turn to the case where equilibrium messages might imply different believed productivities but no delegation occurs ($p_i(m_i, \cdot) = 0$) for any equilibrium message. Note that an equilibrium delegation probability of zero for all messages sent requires that the least productive delegator type does not find it worthwhile to delegate regardless of the message. This implies that $\mu_i(m_i) \leq B$ for all i and messages sent in equilibrium, which given our assumption on $\bar{\theta} \geq B$

implies $\mu_i(m_i) = B$ for all messages sent with positive probability in equilibrium and $\bar{\theta} = B$. This follows from the fact that the selection of a subset of a distribution with a conditional expected value below the distribution means necessarily implies that the non-selected part of the distribution has a conditional expectation above the distribution mean.

The analysis above has shown that in any equilibrium senders send messages that are not informative with respect to the expected productivity. In what follows, we refer to this result, when we speak of the impossibility of information transmission. As messages do not contain any information, senders in equilibrium will only use their own productivity and prior beliefs in order to decide on delegation. Whenever a delegator estimates her productivity below the mean productivity plus bonus, she will delegate. Delegating to either sender or any stochastic tie-breaking procedure is part of equilibrium as long as it does not condition on the messages. Since delegation in equilibrium only uses information that is available without messages, the equilibrium efficiency level is not influenced by the ability of senders sending messages or the type of messages they can send. \square

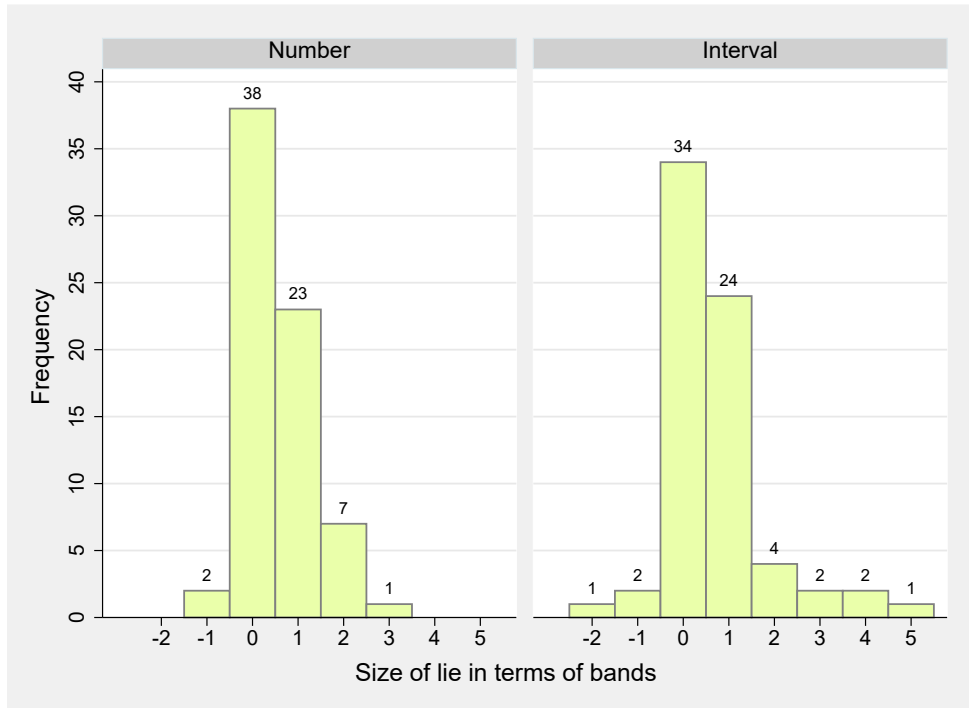
Corollary 1. *Allowing agents to send messages in equilibrium does not improve social welfare.*

This result immediately motivates our research. The natural question that arises is if there are behavioural mechanisms (such as lie-aversion) that allow humans to improve welfare over the non-message case. Moreover, one might ask if the kind of message senders can send has an impact on how much efficiency can be improved. Finally, we are interested in if stereotypes influence delegator decisions and efficiency.

A.2 Some additional analysis

A.2.1 Information contained in messages

Figure A.1: Size of a lie in terms of bands



Note: Sample size: 71 in the *Number* treatment, 70 in the *Interval* treatment.

Table A.1: Regressions of messages across treatments

	(1) OLS Regression <i>Number</i>	(2) Interval Regression <i>Interval</i>	(3) OLS Regression <i>Text</i>
Performance in <i>Task 1</i>	0.878*** (0.067)	0.766*** (0.104)	1.120*** (0.154)
Prior Belief	0.009 (0.014)	0.002 (0.023)	0.022 (0.035)
Female	-0.462 (0.713)	1.683 (1.088)	-0.489 (1.414)
Ethnicity dummies (non-specific is base)			
Black	0.001 (1.294)	3.100 (3.034)	-1.017 (2.384)
White	-0.739 (0.777)	-1.748 (1.482)	-1.367 (1.473)
Asian	-1.271 (0.968)	1.850 (1.242)	-2.470 (1.644)
Math Training	-0.795 (0.732)	1.380 (1.151)	-0.976 (1.237)
Constant Term	4.614*** (1.357)	3.520 (2.264)	1.815 (2.836)

Note: Coefficients with standard error in parentheses.

***, **, * denote significance at p=0.01, 0.05 and 0.10 respectively.

[1] Sample size: 71 in the *Number* treatment, 70 in the *Interval* treatment and 31 in the *Text* treatment.


Table A.2: Probit regression on messages being untruthful

	Coefficients	Std. Err.
Performance	-0.050**	(0.021)
Interval×Performance	0.003	(0.015)
Text×Performance	0.004	(0.019)
Constant	0.784***	(0.266)

Note: ***, **, * denote significance at $p=0.01$, 0.05 and 0.10 respectively.

A.2.2 Delegation, performance and profitable messages

Figure A.2: Delegation fractions by performance and messages



A.3 Sample instructions for the free-text treatment

Welcome

We are pleased to welcome you to the experiment today. During the experiment you will earn money depending on your actions. Communicating with other participants during the experiment is not allowed and will lead to exclusion from the experiment without payment.

The experiments today

In what follows we will ask you to complete a few tasks on the screen. The tasks are very simple and the instructions are given on the screen. You can make use of the provided pen and scratch paper during the tasks. Please raise your arm if you have a question and we will come to you and answer your question.

Don't turn over as yet!!!

A screen shot of the adding-up task is shown as below:

#	#	#	#	#	The Sum
					<input type="text"/>
Submit this sum as your final result					<input type="submit" value="Submit"/>

Your correct answers so far:

Your wrong answers so far:

Remaining Time [sec]:

Instructions (freeNoinfo)

So far you have completed one adding-up task for 7 minutes where you have earned one Dollar per correct answer. Your **performance** in the adding-up task refers to the number of correct answers given during a task.

Overview

In this part of the experiment, you will be randomly grouped with two other participants before you play a **delegation game**. In the delegation game one person (the **Delegator**) decides if he¹ wants the performance of somebody else to count instead of his own performance. So a delegator can make a nice profit if he chooses a person who performs much better than himself. In the meantime the other group members try to convince the Delegator to choose them, since being chosen results in a handy bonus.

Delegation game

As part of the delegation game the adding-up task you have already performed will be repeated. We call the adding-up task you have done already **Task 1** and the repetition you will do **Task 2**. Everybody will be repeating the adding up task and will be paid again one Dollar per correct answer.

However, one participant of the group will be randomly chosen to become the **Delegator**. The **Delegator** can decide if he wants to transfer the responsibility for his payoff from **Task 1** to any of the two other players (this is called to delegate). In each group there are one Delegator and two non-delegator players.

If a **Delegator** decides to delegate, then his performance from the completed **Task 1** will be replaced by the performance of the chosen person in the upcoming adding-up **Task 2**. So the Delegator's earnings from **Task 1** will be replaced by one Dollar for each correct answer the chosen person achieves in **Task 2**. A **Delegator** will pay a bonus of four dollars to the person he delegates to. The chosen player will receive this bonus on top of his normal earnings from performing Task 2.

A **Delegator** who decides not to delegate will be paid according to his own performance in the already completed **Task 1**.

Before the Delegator has to decide if and to whom he is delegating, non-delegators will have to indicate if they are available to be chosen. Then the non-delegator can send a short paragraph up to 400 characters message to the Delegator. When the Delegator decides he will see the messages sent and the avatars (the little pictures you chose) but will have **no information** about the actual past performance of the non-delegator players.

¹ We adopt a male perspective instead of saying she or he to avoid the cumbersome expression.

After the delegation decision has been taken, it will be shown to all group members. Then all participants will do **Task 2** and again add up numbers for 7 minutes, for which the payment is one Dollar per correct answer.

Earnings summary

The computer will randomly decide if you will be paid your earnings from Task 1 **or** Task 2.

Earnings for Task 1

Below you can find the payoffs for Task 1 separately for the case where the Delegator chose to delegate and if not.

The three shaded cells show where payoffs differ depending on whether the **Delegator** delegates or not.

Earnings after delegation:

	<i>Delegator who delegated</i>	<i>Chosen player</i>	<i>Not chosen player</i>
<i>Earnings for Task 1</i>	1 Dollar * performance of the chosen player in Task 2	1 Dollar * own performance in Task 1	1 Dollar * own performance in Task 1
<i>Bonus</i>	- 4 Dollars bonus to the chosen player	+ 4 Dollars bonus for being chosen	--

Earnings without delegation:

	<i>Delegator who did not delegate</i>	<i>Other player</i>	<i>Other player</i>
<i>Earnings for task 1</i>	1 Dollar * own performance in Task 1	1 Dollar * own performance in Task 1	1 Dollar * own performance in Task 1
<i>Bonus</i>	--	--	--

Earnings for Task 2

Delegation does not influence the payoff from **Task 2**. Everybody will be paid 1 Dollar times the performance in **Task 2**.

Please turn over if you would like to see a summary of the sequence of decisions in the delegation following delegation game.

Sequence of decisions summary

Stage 1

Everyone is informed about their role either as a Delegator or a non-delegator player.

Stage 2

Non-delegator players decide whether they want to be considered by the Delegator as a person he can delegate to.

Stage 3

Non-delegators send messages, to the Delegator of their group.

Stage 4

Delegators receive messages but don't know if they are truthful. Then they see the avatars of the interested non-delegator players and make a delegation choice.

Stage 5

Everyone is informed the result of the delegation.

Stage 6

Everyone repeats the adding-up task.

Stage 7

Earnings are displayed.

Are there any questions before we begin?

References

- Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.
- Aronson, J., Quinn, D. M., and Spencer, S. J. (1998). Stereotype threat and the academic underperformance of minorities and women. (*Mimeo*).
- Bargh, J. A. and Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, 43(3):437.
- Bartling, B., Fehr, E., and Herz, H. (2014). The intrinsic value of decision rights. *Econometrica*, pages 2005–2039.
- Bayindir, E. E., Gurdal, M. Y., Ozdogan, A., and Saglam, I. (2020). Cheap talk games with two-senders and different modes of communication. *Games*, 11(2).
- Belot, M., Bhaskar, V., and Van De Ven, J. (2012). Can observers predict trustworthiness? *Review of Economics and Statistics*, 94(1):246–259.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2019). Beliefs about gender. *American Economic Review*, 109(3):739–773.
- Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.
- Cason, T. N. and Mui, V.-L. (2015). Rich communication, social motivations, and coordinated resistance against divide-and-conquer: A laboratory investigation. *European Journal of Political Economy*, 37:146–159.
- Charness, G., Cobo-Reyes, R., Meraglia, S., and Sánchez, Á. (2020). Anticipated discrimination, choices, and performance: Experimental evidence. *European Economic Review*, 127:103473.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6):1579–1601.
- Charness, G. and Dufwenberg, M. (2011). Participation. *American Economic Review*, 101(4):1211–37.
- Chen, J. and Houser, D. (2017). Promises and lies: can observers detect deception in written messages. *Experimental economics*, 20(2):396–419.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, pages 1431–1451.
- De Haan, T., Offerman, T., and Sloof, R. (2015). Money talks? an experimental investigation of cheap talk and burned money. *International Economic Review*, 56(4):1385–1426.
- Dickhaut, J. W., McCabe, K. A., and Mukherji, A. (1995). An experimental study of strategic information transmission. *Economic Theory*, 6(3):389–403.

- Eagly, A. H. and Chin, J. L. (2010). Diversity and leadership in a changing world. *American Psychologist*, 65(3):216–224.
- Fehr, E., Herz, H., and Wilkening, T. (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review*, 103(4):1325–1359.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547.
- Garg, N., Schiebinger, L., Jurafsky, D., and Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1):384–394.
- Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, 108(2):419–53.
- Goeree, J. K. and Zhang, J. (2014). Communication & competition. *Experimental Economics*, 17(3):421–438.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.
- Hausfeld, J., Fischbacher, U., and Knoch, D. (2020). The value of decision-making power in social decisions. *Journal of Economic Behavior & Organization*, 177:898–912.
- Holmström, B. R. (1977). *On incentives and control in organizations*. Doctoral Dissertation, Stanford University.
- Huberman, G. and Rubinstein, A. (2001). Correct belief, wrong action and a puzzling gender difference. *SSRN*.
- Hurkens, S. and Kartik, N. (2009). Would i lie to you? on social preferences and lying aversion. *Experimental Economics*, 12(2):180–192.
- Kawagoe, T. and Takizawa, H. (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.
- Konrad, K. A., Lohse, T., and Qari, S. (2014). Deception choice and self-selection—the importance of being earnest. *Journal of Economic Behavior & Organization*, 107:25–39.
- Lohse, J. and McDonald, R. (2021). Absolute groupishness and the demand for information.
- Lundquist, T., Ellingsen, T., Gribbe, E., and Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, 70(1):81–92.
- Minozzi, William & Woon, J. (2018). Competition and exaggeration in experimental cheap talk games. *mimeo*.

- Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? *Quarterly Journal of Economics*, 122(3).
- Ortiz-Ospina, E. (2018). Economic inequality by gender. *Our World in Data*. <https://ourworldindata.org/economic-inequality-by-gender>.
- Owens, D., Grossman, Z., and Fackler, R. (2014). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics*, 6(4):138–161.
- Peeters, R., Vorsatz, M., and Walzl, M. (2013). Truth, trust, and sanctions: On institutional selection in sender-receiver games. *The Scandinavian Journal of Economics*, 115(2):508–548.
- Sánchez-Pagés, S. and Vorsatz, M. (2007). An experimental study of truth-telling in a sender–receiver game. *Games and Economic Behavior*, 61(1):86–112.
- Serra-Garcia, M., Van Damme, E., and Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1):244–261.
- Sutter, M. (2009). Deception through telling the truth?! experimental evidence from individuals and teams. *Economic Journal*, 119(534):47–60.
- Vanberg, C. (2008). Why do people keep their promises? an experimental test of two explanations¹. *Econometrica*, 76(6):1467–1480.
- Vespa, E. and Wilson, A. J. (2016). Communication with multiple senders: An experiment. *Quantitative Economics*, 7(1):1–36.
- Wang, T. Y., Spezio, M., and Camerer, C. F. (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.
- Wood, D. H. (2016). Communication-enhancing vagueness. *Available at SSRN 2749363*.