

Research paper

Project selection with biased advice: An experiment on competitive cheap talk[☆]

Miguel A. Martínez-Carrasco ^a,^{1,*} Eric Schmidbauer ^b, John Hamman ^c^a School of Management, Universidad de los Andes, Colombia^b Department of Economics, University of Central Florida, Orlando, FL, United States of America^c Department of Economics, Florida State University, Tallahassee, FL, United States of America

ARTICLE INFO

JEL classification:

D82

G31

Keywords:

Cheap talk

Multiple senders

Project selection

ABSTRACT

When agents with private information compete for resources from an uninformed decision-maker and are biased towards their own favored projects (e.g., a CEO decides which division manager's project to fund), they have incentive to strategically communicate about their project's value. However, possible future interaction can mitigate this problem even without reputational concerns, since an agent who induces acceptance of a low-valued project today consumes firm resources that crowd out better opportunities that may arrive in the future. We study this organizational environment both theoretically and empirically using laboratory experiments. We hypothesize and find that truth telling is easier to support as low-quality projects lose value or become more likely to occur, but harder to support as agent competition grows. We see an interesting behavioral result in which beliefs influence responsiveness to parameter changes. Specifically, as agents grow more pessimistic about the likelihood of truthful reporting by their competitors, they respond more sharply to parameter changes, in line with the model's predictions.

1. Introduction

Information is often diffuse within organizations and those who hold it may misrepresent what they know to alter the behavior of decision makers. One such setting is the “project advocacy” model (Li et al., 2016 among others) where multiple agents are privately informed about their own project and promote it to a decision maker with a limited budget. It has been shown that such agents might not misrepresent even in the most pessimistic case where they (i) do not internalize any benefit to the organization as a whole, and (ii) have no regard for reputation (Schmidbauer, 2017). The basic idea is that an agent who currently has a poor quality project might truthfully reveal this if the continuation value in the game is high enough, which depends on how much more valuable and likely higher value projects are in the future.

Despite the importance of communication within organizations on performance, empirical evidence on experts' reporting behavior is difficult to collect by its very nature: an expert with an information advantage over the decision maker likely has this same advantage over the econometrician, making it hard to compare the true and reported states of the world. In this paper, we use lab experiments to explore the extent of truthful communication in a project selection context and how it varies with the

[☆] For financial support we thank the International Foundation for Research in Experimental Economics (IFREE) Small Grants program. For helpful comments we thank participants at several conference and seminar presentations. Brian Gifford, Rebecca Gilland and Nicolas Salas provided research assistance.

* Corresponding author.

E-mail addresses: ma.martinezc1@uniandes.edu.co (M.A. Martínez-Carrasco), eschmidb@ucf.edu (E. Schmidbauer), jhamman@fsu.edu (J. Hamman).

¹ Author order determined using the American Economic Association's Author Randomization Tool.

<https://doi.org/10.1016/j.jebo.2025.106936>

Received 16 February 2024; Received in revised form 11 December 2024; Accepted 7 February 2025

0167-2681/© 2025 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

number of agents, the value of high and low projects, and their likelihood. We do so to validate theoretical predictions as well as explore some of its counterintuitive comparative statics. As an example of the latter, the project advocacy model predicts an agent with a low project is more willing to reveal this (hoping for a better project in the future that is accepted then) the *less likely* a high value project is to occur, since this also means high projects are less likely now and so the chance of being preempted by a competing agent is lower.

An example of the project advocacy setting is that of division managers of a company who report to the CEO about the profitability of projects available within their unit. When a budget constraint does not allow for all potential projects to be adopted, the CEO must pick only the best to fund even though they likely are at an information disadvantage relative to the better-informed division managers. A similar dynamic is at play when multiple lobbyists provide information to a politician about the benefits from various interventions or spending policies, only some of which can be adopted. In each case the more informed agents may choose to exaggerate their claims with the hope of diverting resources to themselves to their own favored cause, regardless of whether it is the best overall project available to the principal.

We adapt and experimentally test a model of competition between agents over time and show that this competition can distort communication and leads to misallocated resources. However, we show that two factors serve to limit the distortion. First, possible future interaction can improve communication even without reputational concerns, since an agent (division manager, lobbyist) who induces acceptance of a low-valued project today consumes firm resources that crowd out even better opportunities that may arrive in the future. This effect disincentivizes egregious misrepresentations and its strength depends in part on the distribution of future projects' profitability across divisions. Second and relatedly, the concern for crowding out is only relevant if there is some budget remaining in the future. Since this is more likely when there are fewer competing agents, we show that communication is less distorted within firms that have reduced the number of specialized business units they contain.

Adapting the theoretical model from Schmidbauer (2017), we devise experiments to explore behavior in this organizational environment. Specifically, we study the incentives of competing agents to strategically communicate about their own favored project's value to a decision-making principal when new projects arrive over time. Agents privately observe the independently realized value of their own project (which is taken to be high or low quality) and recommend their project for adoption or not. Agents cannot observe the quality of each other's projects, the idea being that in modern organizations information is often compartmentalized within highly specialized divisions.² Agent reports are modeled as cheap talk, meaning there is no explicit cost to the agent from lying.

After observing advice from all agents, the principal decides which single project to adopt, if any. If no project is adopted subjects enter the next period with new independently drawn projects and continue indefinitely until one project is adopted. We assume the principal and agent whose project is adopted benefit symmetrically from the outcome of the project, whereas an agent whose project is not adopted receives no benefit.

Under what conditions will agents report their information truthfully? An agent who knows their project is low quality but successfully lies and induces its acceptance nets a small benefit (since the project is of low quality), though this outcome is better for him/her than a competing agent's project being selected. On the other hand, reporting truthfully ensures rejection in this period but leaves open the possibility that a high quality project will be available next period and is accepted then. If the payoff from a high quality project is sufficiently high, or the probability of a high quality project is sufficiently low (so that preemption by a competing agent in the current period is unlikely), an agent with a low quality project will prefer to truthfully reveal his/her information. Thus to some extent agents internalize the allocative distortions from lying, even absent reputational concerns. It is also predicted theoretically that as the number of competing agents rises, the incentive to truthfully report when the project is low decreases. This is because the probability of being preempted by at least one competitor increases in the total number of competitors.

Addressing this question empirically presents several challenges. Organizations may use criteria to make investment decisions beyond the probability of being misled by project managers, which could lead less-advantaged managers to lie more. Moreover, the timing of project availability may not be controllable and project quality may be unobservable. Laboratory experiments allow us to explore the model's predictions in a controlled environment. Experiment 1 involves a single instance of the project selection task described above, with treatments varying core elements of the task environment. Specifically, we vary (i) the probability of receiving a low project, (ii) the value of a low project, and (iii) the number of competing agents. Results are largely in line with comparative statics; agents are more truthful when high quality projects become rarer, when the low projects convey worse returns, and when facing reduced competition. Agent behavior closely mirrors the beliefs that principals hold, again reinforcing the model's predictions.

Findings from Experiment 1 suggest that agents respond to incentives for honesty, though perhaps not as strongly as predicted. This could be due to the one-shot nature of the task, and so agents with more experience in the environment may respond even more strongly to increased incentives for truth-telling. Experiment 2 addresses this by having participants repeat the task a number of times, with random and anonymous rematching between tasks to preclude reputational effects. We again see evidence that agents respond to incentives, as in Experiment 1. For their part, principals largely fall into two groups based on their acceptance decisions. Over half make decisions fully consistent with the model's predictions, accepting a project when the project's agent reports that it is high quality. The majority of remaining principals exhibit a more cautious approach, rejecting all projects initially before eventually

² For example, the operations department may know the benefit that would accrue from buying a machine to further automate production, while the marketing department might know the increase in sales that would occur from a particular advertising campaign, yet neither knows the benefit from the other department's project. In the lobbying context, an environmental lobbyist may know the benefit from a cap on carbon emissions but not the effect of safety improvements to an interstate highway, while the converse is true for a transportation lobbyist.

accepting a project whose agent reports it to be high quality.

Interestingly, we actually see a decline in truth-telling over time in Experiment 2. This is often – though not always – in conjunction with beliefs growing more pessimistic. We conclude our analysis by more carefully exploring the role of beliefs in guiding participant behavior, structured around a novel belief-based equilibrium. Our findings suggest that pessimistic agents – those who believe their competitors will be highly unlikely to report truthfully – are most sensitive to the treatments, while optimistic agents tell the truth more frequently, showing less sensitivity to parameterization.

Our findings are of relevance to both the academic literature and practitioners. Information problems are prevalent within organizations and have implications for optimal organization design. By validating our theoretical predictions with experiments we provide new evidence that organizations may benefit from reducing their core functions (in the model, the number of agents with their own area of expertise), and illustrate one key detriment to organizational size and complexity. In such cases, market exchange and contractual arrangements with external agents to provide some of these tasks may limit the incentive misalignment faced by firms.

The results also demonstrate that the source of cooperative behavior from agents with a short term incentive to be opportunistic can matter. If it is just reputation that is disciplining agents then outsourcing functions should not have an effect on information flows from agents. However, outsourcing functions would provide a benefit if our mechanism is at play. Note too that these two mechanisms are not mutually exclusive: reputation could complement our mechanism by further strengthening the incentive for truthful reporting.

Although we have thus far only discussed our model as it relates to organizational design, our findings also apply to instances of market exchange where more informed suppliers pitch their competing solutions to satisfy a firm's given need. This can take many forms and, importantly, we show that trust in this context need not rely on a reputational mechanism. For example, an insurance company or pension fund might seek investment opportunities, with various hedge funds competing by proposing their own proprietary trading or investing strategy. Suppose the discovery of profitable investment strategies by hedge funds is a random process over time and they receive a percentage of their clients' returns. By the logic of our model, a hedge fund might "hold back" when it knows its current best investment strategy is sub-par because a better investment strategy that would earn it more money might be discovered later, even though this risks losing the client entirely to a competing firm today.

2. Related literature

Experimental research on honesty largely started with tests of sender-receiver models such as Crawford and Sobel (1982), in which the two parties must communicate strategically in the face of various degrees of incentive conflict. Such studies typically find that some senders resist sending advantageous, though dishonest, messages (see Abeler et al., 2019 for an overview and meta-analysis). This has led to a number of theoretical contributions suggesting preferences for truth-telling or being seen as honest that could support these findings as a utility-maximizing behavior (Abeler et al., 2019; Gneezy et al., 2018; Kartik et al., 2007).

Other recent work sheds light on how competitive settings drive dishonest behavior. For example, Tergiman and Villeval (2023) extend our understanding of dishonesty to project selection with biased advisors by showing how reputational concerns change the type of lies people use. In a finding related to our study, they find that senders use fewer detectable lies and increasingly rely on deniable lies under fixed matching compared to random and anonymous rematching. Another related paper, Dato and Nieken (2020), show that related unethical behavior is driven in part by beliefs about unethicality of one's opponents.

In this research, we shift our focus away from individual incentives and preferences for truth-telling, instead studying ways in which honesty impacts organizational function more broadly. For this, we draw on a class of 'competitive cheap talk' models that highlight how organizational growth adds new challenges. We focus on advisor competition without reputational concerns. Specifically, in an environment with one principal and one agent, both parties will prefer the same outcome, allocating resources to the best possible project. However, when multiple agents, each with their own project, must compete for the same limited resources, a new tension arises between agents and the principal in which agents may use misinformation to induce acceptance of their project.

Li et al. (2016) started this literature stream with their one-period model in which two biased agents compete via cheap talk for the funding of a single project. Li (2016) extended that model to a dynamic setting where the principal consults a single agent in each period and must alternate between two agents over time with some known probability. In each, the degree of bias is shown to reduce the set of parameters under which truth-telling can be supported in equilibrium. Schmidbauer (2017) extends this theoretical work to a setting with agents so biased that they only care about their own outcomes, finding that even under such extreme incentives truth-telling equilibria persist. We adapt the model of Schmidbauer (2017), to analyze these truth-telling paths empirically. In the present paper, all $n \geq 2$ agents compete in a multi-period setting and do not internalize any benefit when a competitor's project is adopted (see also Rantakari, 2018 and Schmidbauer, 2019).

This paper contributes to the project selection literature more generally. In Moldovanu and Shi (2013) new projects keep arriving until one is unanimously agreed upon by a committee with members with different preferences, while in Armstrong and Vickers (2010) the principal can delegate the project selection task but is not aware of all the projects available to the agent.³ More generally

³ Other models allow for influential communication with just a single agent through other means such as reputation. For example, Kim (1996) explores how reputation can affect cheap talk over an infinite horizon. Unlike his paper, we do not require infinitely many periods nor do we have ex-post verifiability. Along similar lines (Sobel, 1985) explores reputation in a cheap talk model in which the state is fixed across periods. Finally, in a static model (Chakraborty and Harbaugh, 2007) demonstrate that a single expert who observes all dimensions of the state space can make credible comparative statements even when it would not be credible on a single dimension.

this paper relates to the literature on competition between experts; see for example (Gilligan and Krehbiel, 1989), Krishna and Morgan (2001a,b) as well as Battaglini (2002). In each of these models, however, both experts observe the same state of the world, whereas in the present paper each expert has private information about just one dimension of the state.

We also join a growing body of experimental research seeking to provide empirical insight into organizational structure, function and information asymmetries. Recent studies provide empirical support to theoretical models exploring the role of information in decentralized versus centralized firm structure (Hamman and Martínez-Carrasco, 2022; Evdokimov and Garfagnini, 2019). Hamman and Martínez-Carrasco (2022) analyze the joint decision of managers to decentralize and which agents type to hire in a task uncertainty environment where agents have better information. Evdokimov and Garfagnini (2019) design an experimental environment based on Alonso et al. (2008) and Rantakari (2008), finding that communication closely matches equilibrium predictions. However, communicated information was suboptimally incorporated by decision makers, resulting in deviations from efficiency under both decentralized and centralized organizational structure. Brandts and Cooper (2021) experimentally test some of the underlying assumptions of these models, finding that agents are more truthful in vertical communication than predicted, which enables efficient coordination under centralized control. In this paper, we evaluate strategic communication under competition for resources without considering reputational concerns. We similarly see overly truthful vertical communication in several treatments, though truthfulness does respond to incentives as predicted by our model. Overall, we join these studies in contributing to our empirical understanding of how communication affects organizational structure and function.

3. Theory and predictions

3.1. Model with biased advice

We adapt the model of Schmidbauer (2017) but discretize the state space. There are $n+1$ players: n agents and a single decision-maker (DM). Each agent i has access to a single project whose profitability θ_i is $\$l$ with probability $r \in (0, 1)$ and $\$h$ otherwise, where $0 \leq l < h$. Agent i privately observes the value of θ_i , which is not observed by the DM nor agent $j \neq i$. It is common knowledge that all realizations are independent and identically distributed across agents.

After observing θ_i , each agent simultaneously sends a message $m \in M = \{\text{High}, \text{Low}\}$ to the DM. Communication is modeled as cheap talk, thus there is no explicit cost to sending either message regardless of the true state. Next, the DM decides which single project to accept, if any.⁴ Projects are assumed to be indivisible and so this decision as respects each agent's project is binary. If no projects are accepted, the players proceed to the next period in which new i.i.d. draws are available. Thus all realizations are independent across agents and time. The game continues indefinitely until a single project is adopted, and we assume that previously rejected projects cannot be brought back. Payoffs in the stage game are as follows:

$$U_{DM} = \begin{cases} \theta_i & \text{if agent } i\text{'s project is adopted} \\ \$0 & \text{if no project was adopted} \end{cases}$$

$$U_{Agent\ i} = \begin{cases} \theta_i & \text{if agent } i\text{'s project is adopted} \\ \$0 & \text{otherwise} \end{cases}$$

There is no time discounting and players are expected utility maximizers.⁵ Thus the DM prefers to accept the project with the highest θ irrespective of which agent generated it, while each agent only prefers that his own project is accepted. In particular, agent i internalizes no benefit at all if agent j 's project is adopted, $i \neq j$.

3.2. Theoretical predictions

We proceed by defining what we mean by truth-telling before finding conditions under which agents will truthfully reveal their information.

Definition 1. A strategy is truth-telling if the agent with an l project sends message Low and the agent with an h project sends message High. An equilibrium is truth-telling if each agent uses a truth-telling strategy.

In an equilibrium in which all agents use a truth-telling strategy, the DM's best response is to reject when receiving the low message and accept when receiving the high message. Since the DM is indifferent between any agents that report high, it is a best response to randomly pick a project to adopt from those who so reported. By our assumption that the DM treats agents symmetrically it follows that a truth-telling equilibrium must be of the form just described.

We now look for a truth-telling equilibrium. Since an agent with a high project never has incentive to misreport it as low, it suffices to confirm that an agent with a low project receives a weakly higher payoff from reporting low than falsely reporting high, given that all other agents truthfully report. To understand the reporting incentives, first observe that before learning one's type an agent's ex-ante expected payoff in a truth-telling equilibrium is simply $\frac{h}{n}$ since each agent is equally likely to "win" (either now or in a future period) and the payoff from winning is h . However, upon learning his project is low an agent's payoff in the truth-telling

⁴ Our theoretical findings would not qualitatively change if $k > 1$ projects could be adopted provided that $k < n$.

⁵ We specify risk-neutrality for analytical convenience. In Appendix A we show that if agents are risk-averse the equilibrium structure remains qualitatively unchanged but that incentives must be sharper to induce an agent to truthfully report bad news.

equilibrium falls to $r^{n-1} \left(\frac{h}{n}\right)$, since in this case he only achieves expected payoff $\frac{h}{n}$ if the game proceeds to the next period, which in turn only occurs when all other agents currently have a low project. On the other hand, by deviating from a truth-telling equilibrium an agent with a low project can increase his probability of winning but only achieves payoff equal to his true type l when doing so. Truth-telling is weakly more profitable to the low type agent whenever

$$l \times \underbrace{\left(r^{n-1} + \frac{\binom{n-1}{1}(1-r)r^{n-2}}{2} + \dots + \frac{(1-r)^{n-1}}{n} \right)}_{\text{Low type's probability of winning from lying}} \leq h \times \underbrace{\left(\frac{r^{n-1}}{n} \right)}_{\text{Low type's prob. of winning from truth-telling}} \tag{1}$$

The term in large parentheses on the left hand side is the probability of acceptance when reporting a high project, given all other players truthfully report: with probability r^{n-1} all others have a low project and so a high report is accepted for sure, with probability $\binom{n-1}{1}(1-r)r^{n-2}$ exactly one other player reports high and so a high report is accepted with probability $\frac{1}{2}$ (due to the tie-breaking rule), and so on. It can be seen by comparing this term to the corresponding term on the right hand side of (1) that indeed the probability of winning the game is higher for a low type when lying than telling the truth.⁶ Offsetting this, though, is the lower payoff l from lying than the payoff h that would be obtained if one were to eventually win by truth-telling.

Rewriting the summation term on the left hand side of (1) more succinctly as $\sum_{i=0}^{n-1} \binom{n-1}{i} \frac{r^{n-1-i}(1-r)^i}{1+i}$ and recognizing that this equals $\frac{1-r^n}{n(1-r)}$ allows us to substitute and rearrange terms to obtain

$$l \leq \frac{r^{n-1}}{1-r^n} (1-r)h, \tag{2}$$

which gives us our equilibrium condition.⁷ Using that $1-r^n = (1-r) \sum_{j=0}^{n-1} r^j$ and rearranging once again we obtain our incentive compatibility condition

$$0 \leq \left(\frac{r^{n-1}}{\sum_{j=0}^{n-1} r^j} \right) h - l \equiv IC(l, h, r, n). \tag{4}$$

In summary, IC represents the increase in the expected payoff to a low type agent from truthfully reporting instead of lying, assuming all others truthfully report. Thus $IC \geq 0$ is the necessary and sufficient condition such that a truth-telling equilibrium exists. In our experimental implementation we will also use the fact that the larger is IC the stronger is the incentive of a low project agent to truthfully report, or equivalently the costlier a mistake it is to lie.

Proposition 1. *A truth-telling equilibrium exists if and only if $IC(l, h, r, n) \geq 0$. The strength of the incentive the low type has to report truthfully, IC , is increasing in h and r while decreasing in l and n .⁸*

To better understand the comparative statics of the proposition, consider the incentive to tell the truth by an agent who has a low project. Doing so earns a positive payoff only if the game continues to the next period, which requires all other agents to have a low project. Since the probability of this is decreasing in the total number of agents, and increasing in the prior probability of a low project, the incentive for truth telling decreases in n but increases in r . Next, by continuing to the next period the agent will possibly earn h in the future but gives up l now. Thus a more valuable high project strengthens the incentive for truth telling while a more valuable low project weakens it.

Finally, we have deliberately excluded reputational concerns to focus on the mechanism of the model. However, we acknowledge that in many real-world settings, reputational concerns could interact with the incentives we analyze. For example, suppose an agent who lies about a project that is accepted incurs a reputational cost f .⁹ The effect of this on the incentive compatibility condition is straightforward: in line (1), the reputation cost f is subtracted from l on the left-hand side of the inequality, and following the algebra through line (4) becomes $0 \leq \left(\frac{r^{n-1}}{\sum_{j=0}^{n-1} r^j} \right) h - l + f$. The value of IC is increased by f and thus the parameter space supporting

⁶ This can be easily seen since the first term on the left hand side, r^{n-1} , is itself greater than the term on the right hand side, $\frac{r^{n-1}}{n}$.

⁷ Schmidbauer (2017) shows that when the state space is a continuum a symmetric equilibrium entails each agent sending the “accept” message if and only if his project exceeds the symmetric threshold t , which must satisfy

$$t = \frac{\delta F(t)^{n-1}}{1 - \delta F(t)^n} \int_t^\infty \theta dF(\theta), \tag{3}$$

where F is the cumulative distribution function of θ and $\delta \in [0, 1]$ is the time discounting factor. In equilibrium the threshold type is indifferent to continuing (by recommending rejection) or stopping (by recommending acceptance), the payoffs to which are proportional to the right and left hand sides of line (3), respectively. In our setting there is no time discounting ($\delta = 1$), the threshold value t is the low project’s value l , the probability of being at or below this threshold is simply r , and integrating over all types exceeding l consists of merely the project value h multiplied by its probability $1-r$.

⁸ Proof in Appendix C, subsection C.1.

⁹ This cost can be endogenized in a repeated game context. For instance, suppose that an agent who has lied in a prior game will not have a project accepted in the event that more than one agent has reported having a high project. That is, a previous lie results in losing any ties that might occur in the future since the agent is deemed less credible. If so, a truth-telling equilibrium may still exist since an agent who previously lied still has some opportunity for acceptance (when all others report low) but will earn a lower expected payoff than if ties were broken evenly. This reduction in expected payoff can be regarded as the value f .

a truth telling equilibrium has grown larger, making truth-telling easier to sustain. It also follows that the sign of the comparative statics in [Proposition 1](#) remains unchanged.

In the following section, we vary core elements of this theoretical environment to inform our experimental design and test the model's comparative statics. We focus on the probability r of receiving a low project, the value l of a low quality project, and the degree of competition among agents n .

3.3. Equilibrium selection

It is well known that in communication games a plethora of equilibria exist. For example, in the present model there is always a babbling equilibrium in which the DM randomly selects a project to immediately adopt and all agents always report their projects are High. Above we found the necessary and sufficient conditions such that a non-babbling pure strategy equilibrium exists, and when such an equilibrium exists we select it since an equilibrium with meaningful communication is arguably focal. This equilibrium has several other desirable properties over the babbling equilibrium.

Proposition 2. *The truth telling equilibrium of [Proposition 1](#) dominates the babbling equilibrium in the following sense: (i) ex-ante Pareto dominance, (ii) interim Pareto dominance, and (iii) first order stochastic dominance of the players' payoffs.*¹⁰

Another source of multiple equilibria arises from the fact that two equilibria can be outcome equivalent yet differ with respect to the specific messages used to induce any outcome.¹¹ Since we believe equilibria in which the literal meaning of a message corresponds to its equilibrium meaning are more intuitive, we restrict attention to these in our analysis.

Finally, in [Schmidbauer \(2017\)](#) it is shown that asymmetric as well as non-stationary equilibria may exist. Owing to the difficulty subjects likely would have coordinating on such equilibria, we focus here on equilibria that are symmetric and stationary. That is, we select equilibria in which each agent uses the same reporting function in each period and when the DM is indifferent between accepting some agents' projects he uniformly mixes between them.

4. Experiment 1

4.1. Design

We conducted all sessions in the xs/fs laboratory at Florida State University. Each session involves between 18 and 24 subjects. Sessions were programmed in the experimental software z-Tree ([Fischbacher, 2007](#)) and subject recruitment was done through ORSEE ([Greiner, 2015](#)). Sessions last on average for 45 min, with mean earnings of \$15 per subject.

In each session, subjects are divided into groups of three (two-agent treatments) or four (three-agent treatments). Each group consists of one subject in the role of Decision Maker (DM) and two or three in the role of Agent, with roles fixed throughout the session. Groups played a project-selection task one time, with the session concluding after all DMs selected a project for their group.

In the task, agents each receive a project with privately known quality, High or Low. These projects are assigned by an independent and identically distributed random draw for each agent, with probability r of receiving a low quality project. The probability of receiving a high project will therefore be $1 - r$ (subjects were informed of all parameter values at the start of the session). Once agents see their project, they each send a private binary message to the DM, recommending either that their project be chosen or that it be declined. Specifically, agents choose a message from the set {Low, High}. While the Agents are selecting their message, we elicit beliefs from the DM regarding the percentage of Agents in the session receiving a low quality project who will send a message of Low. We similarly elicit these beliefs from Agents after they send their message, while the DM is making their decision. The DM views both messages and then chooses either project or declines both. If both projects are declined, the period will "reset", with agents receiving new independently drawn projects and repeating the process, without additional belief elicitation. The task will conclude after a project has been accepted, or after twenty unsuccessful attempts.¹² The payoffs from a high project are $\$h$ and for a low project are $\$l$ for both the DM and selected agent.

Once all triads have completed the task, the session concludes with a brief demographic questionnaire and a measure for risk preference. The risk elicitation task is the lossless version from [Eckel and Grossman \(2008\)](#). Earnings from the project selection task are combined with earnings from the risk elicitation. Upon completing the questionnaire items, participants were paid privately and left the laboratory. In experiment one, we have 277 participants across six treatments for a total of 87 groups. The summary statistics of our demographic variables are in Appendix D, Table D.1.

¹⁰ Proof in Appendix C, subsection C.2.

¹¹ For example, a fully separating equilibrium in which an l -type agent reports Low and an h -type agent reports High is outcome equivalent to an equilibrium in which type l reports High and type h reports Low, since the true state is learned by the DM in either case.

¹² This threshold was never reached in any experimental session.

Table 1

Experiment 1 overview.

Treatment	Parameters & IC					Participation			Decisions by	
	r	$\$h$	$\$l$	Agents	IC_{agent}	Sessions	N	Groups	Manager	Agents
Baseline	.90	10	2	2	2.74	2	45	15	79	158
Varying r	.75	10	2	2	2.29	2	45	15	50	100
	.60	10	2	2	1.75	2	45	15	22	44
Varying l	.90	10	3.50	2	1.24	2	39	13	47	94
	.90	10	5	2	-0.26	2	39	13	57	114
Varying n	.90	10	2	3	0.99	3	64	16	37	111

Notes. The first section specifies the parameters and incentive compatibility constraint for each treatment. The second section shows the number of sessions, participants and groups for each treatment. The third section focuses on the number of decisions made by participants in the roles of managers and agents.

4.2. Treatments

We conduct six treatments using this environment, each designed to test the model's predictions. The parameters for each treatment are summarized in Table 1, which also reports the Agent's IC and the number of participants and choices made for each treatment. In the first three treatments, we vary the value of r , while keeping the values for $\$h$ and $\$l$ at \$10 and \$2, respectively and have two agents. In the next set of treatments, we raise $\$l$ to examine how higher values of the low outcome influence willingness to deviate from truthful advice. Following this, we explore a treatment that adds a third agent, preserving a direct comparison to the baseline treatment.

Recall that truth telling is supported in equilibrium whenever the IC is positive. Therefore, we expect to see truthful reporting in the three treatments that vary r , though the strength of the incentive for truth-telling rises as r rises. When we vary l , we predict truthful reports in only the $l = 2$ and $l = 3.5$ treatments, though we expect to see lower levels of truthful reporting in $l = 3.5$ than any of the treatments that vary r . With a negative IC, the model does not predict truth-telling in $l = 5$. Adding a third agent still supports a truthful reporting equilibrium, but the incentive compatibility is much weaker than in the baseline.

4.3. Results

Our analysis follows the order of treatments in Table 1, focusing on behavior of participants in the role of agent.¹³ Agent participants made a total of 584 message decisions, averaging 3.4 decisions per participants. We focus primarily on the behavior of agents who receive a low-quality project and the signal they send to the DM. In other words, our main level of analysis is whether the agents tell the truth when they have the opportunity to lie. We take each decision individually in the graphical analysis and cluster them by agent in the regression analysis. Fig. 1 summarizes the findings. All pairwise comparisons are two-tailed Mann-Whitney tests, with regression results providing further support in Table 2.

We see significantly more truth-telling in the baseline than either of the other two treatments that vary r , in line with the predictions of the model, though we find no significant differences when comparing the 0.6 and 0.75 treatments to each other (left-hand panel of Fig. 1, $p = 0.32$). Both $r = 0.6$ and $r = 0.75$ are significantly different from $r = 0.9$ ($p = 0.04$ and $p = 0.00$, respectively). The middle panel shows that increasing the payment of the low project significantly decreases the truth-telling probability, as expected.¹⁴ Comparing these results with varying r , it appears that agents may react closer to the model predictions when the treatment is related to payments than to probabilities.

The right-hand panel of Fig. 1 explores the effects of increased agent competition when $l = 2$. Adding an agent reduces truth-telling, as predicted ($p = 0.03$). Agent beliefs (represented by triangular markers) about others reporting truthfully when receiving a low project are in line with our behavioral results, which we return to later in our analyses.

Table 2 reinforces the results discussed so far regarding truth-telling (See Table D.3 in appendix for regression results on beliefs). Models 1 and 2 show the causal reduction in truth-telling as we lower r . Models 3 and 4 show clearly that increasing l significantly reduces truthful reports, and models 5 and 6 confirm the reduction in truthful reporting under increased agent competition when $l = 2$.

5. Experiment 2

Results from Experiment 1 show that the model has predictive power, as treatments largely find support for the comparative statics. This is true for both behavior and beliefs of participants in each role. These results suggest that organizational structure can be engineered to effectively promote truthful communication in the presence of extremely biased agents.

However, one critical question that Experiment 1 cannot address is whether behavior may converge as participants gain experience. Because the session ended once a project was accepted, we cannot test for learning, which is commonly discussed as a critical element in improving decision making in organizations. As such, repeated play that allows for learning but precludes reputation-building may provide a stronger test of the model. Would participants converge even closer to the equilibrium predictions if they learn to play the game better? Experiment 2 addresses this question.

¹³ We briefly examine the decision making of Principals in our experiments in Section 7.1.

¹⁴ We have $p = 0.05$ and $p = 0.00$ comparing $l = 2$ to $l = 3.5$ and $l = 5$, respectively. Comparing $l = 3.5$ and $l = 5$ gives a p -value of 0.00.

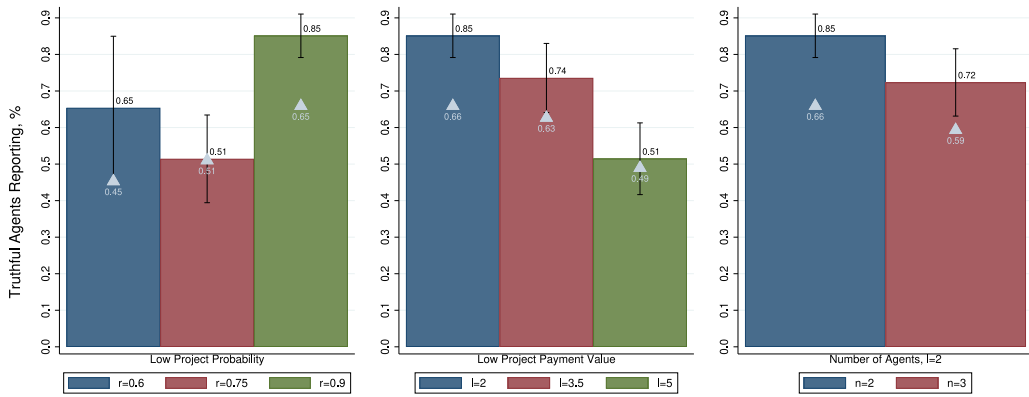


Fig. 1. Truth-telling by Treatment, Experiment 1.

Notes. Summary of honest reports by agents, varying r in upper left panel (a), varying l in the upper right panel (b), varying n with $l = 2$ in lower left panel (c) and varying n with $l = 3.5$ in lower right panel (d). Beliefs are indicated by gray triangles for each treatment.

Table 2
Regressions of truth-telling, Experiment 1.

	Varying r		Varying l		Varying n	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
$r = 0.6$	-1.107*	-0.372				
	(0.61)	(0.70)				
$r = 0.75$	-1.686***	-1.204**				
	(0.58)	(0.61)				
$l = 3.5$			-0.720	-1.055*		
			(0.53)	(0.55)		
$l = 5$			-1.685***	-2.257***		
			(0.54)	(0.59)		
$n = 3$					-0.782	-0.884*
					(0.53)	(0.51)
Constant	1.743***	2.522***	1.743***	0.946	1.743***	1.648*
	(0.37)	(0.86)	(0.37)	(0.70)	(0.37)	(0.85)
Controls	No	Yes	No	Yes	No	Yes
Obs.	237	237	331	331	235	235
Pseudo R_2	0.0987	0.1499	0.0838	0.1242	0.0240	0.0682

Notes. The dependent variable is the truth-telling probability. Logit regressions with clustered standard errors by agent. Odd columns include individual controls of age, ethnicity, career orientation, year at the university and risk tolerance.

* p -value < 0.1.
 ** p -value < 0.05.
 *** p -value < 0.01.

5.1. Design

Experiment 2 mirrors the design of Experiment 1, with one difference. Participants now play the project selection task 10 times, with anonymous random rematching between each iteration. Roles are fixed throughout the session, but agent IDs are randomly re-assigned for each task and participants are informed that they will not interact in the same group more than once (see Appendix G for sample instructions). In this way, we give agents and principals the ability to learn from feedback in early tasks while reducing confounding reputational concerns.

5.2. Treatments

We conduct five treatments, similar to those in Experiment 1. Treatment parameters and IC values are given in Table 3. Again, positive IC values support a truth-telling equilibrium. In experiment 2, we have data from 154 participants in total across all treatments for a total of 650 triads and 1723 opportunities to select a project.¹⁵ As in Experiment 1, the summary statistics of our demographic variables are in Appendix D, Table D.2.

¹⁵ In only two of these opportunities did the participants not select a project in any of the 20 rounds they had available. As a consequence, their payoffs were zero in those cases.

Table 3
Experiment 2: Treatment overview.

Treatment	Parameters & IC					Participation			Decisions by	
	r	$\$h$	$\$l$	Agents	IC_{agent}	Sessions	N	Groups	Manager	Agents
Baseline	.90	10	2	2	2.74	2	36	120	417	834
Varying r	.60	10	2	2	1.75	2	45	150	249	498
Varying l	.90	10	3.50	2	1.24	2	42	140	403	806
Varying n	.90	10	2	3	0.99	2	48	120	338	1014
	.90	10	3.50	3	-0.51	2	48	120	316	948

Notes. The first section specifies the parameters and incentive compatibility constraint for each treatment. The second section shows the number of sessions, participants and groups for each treatment. The third section focuses on the number of decisions made by participants in the roles of managers and agents.

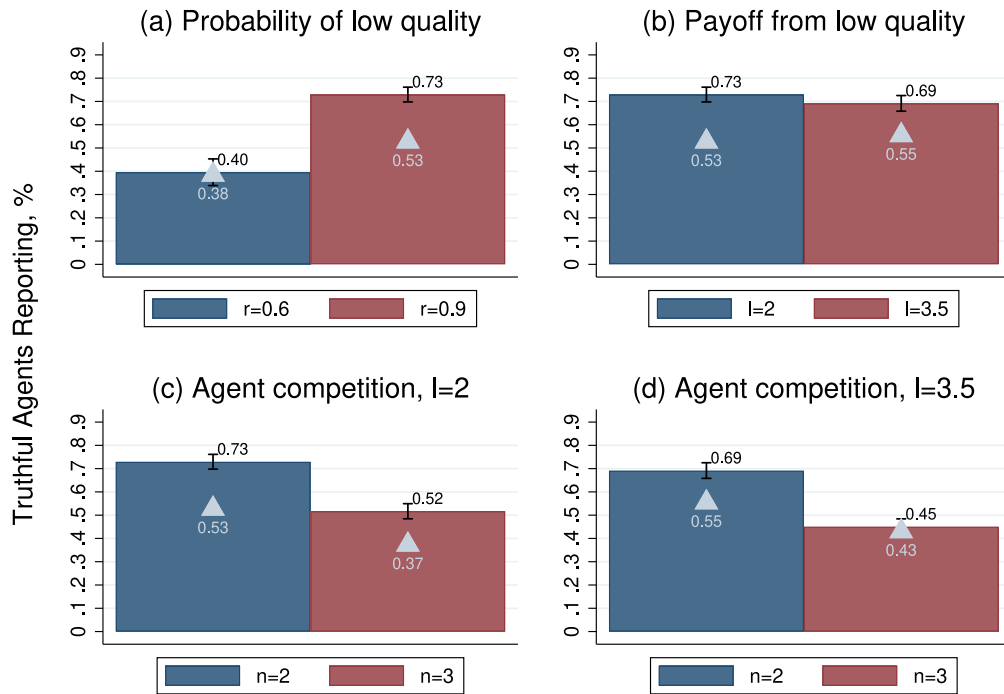


Fig. 2. Truth-telling by Treatment, Experiment 2.

Notes. Summary of honest reports by agents, varying r in upper left panel (a), varying l in the upper right panel (b), varying n with $l = 2$ in lower left panel (c) and varying n with $l = 3.5$ in lower right panel (d). Beliefs are indicated by gray triangles for each treatment.

5.3. Results

Our main level of analysis is whether the agents tell the truth when they have the opportunity to lie.¹⁶ We take each decision individually in the graphical analysis and now cluster them by agent and period in the regression analysis. We see modest intertemporal effects across treatments, and so first present results aggregated over all ten rounds before discussing trends. Fig. 2 shows aggregated truth-telling behavior for all treatments of Experiment 2. As in Experiment 1, we see that lowering r reduces truthful reports (see panel (a) in the upper left, $p = 0.00$). As in Experiment 1, we observe a reduction in truth-telling as we increase l , though here it is weaker on aggregate (panel (b), $p = 0.1069$). Finally, we see clear and strongly significant reductions in truth-telling with increased agent competition, regardless of the value of l (lower two panels, all $p = 0.00$). In Experiment 2, the results fit the model robustly. Once again, agent beliefs are very similar to observed behavior.

We find similar results from the regression analysis in Table 4.¹⁷ Models 1 and 2 show a consistent and significant negative effect on honesty as the probability to receive a low project decreases from $r = 0.9$ to $r = 0.6$. Models 3 and 4 confirm the very modest impact of changing l reported in Fig. 2. Only when including participant controls do we see marginally significant reductions in truth-telling from raising the low quality project value, though the effect is negative in both specifications. Models 5 through 8 reinforce the finding that increased agent competition substantially reduces honest reporting, in line with the model's predictions.

¹⁶ Again, see Section 7.1 for a brief note on the decision making of Principals in our experiments.

Table 4
Regressions of truth-telling, Experiment 2.

	Varying r		Varying l		Varying number of agents			
	Model 1	Model 2	Model 3	Model 4	Low payment = 2		Low payment = 3.5	
	Model 5	Model 6	Model 7	Model 8				
$r = 0.6$	-1.447*** (0.38)	-1.713*** (0.40)						
$l = 3.5$			-0.142 (0.28)	-0.422* (0.23)				
3 Agents					-0.935** (0.474)	-0.990** (0.47)	-1.005*** (0.37)	-0.811** (0.33)
Constant	1.500*** (0.40)	0.724*** (0.74)	1.611* (0.89)	1.397 (0.88)	2.874** (1.17)	2.511 (1.58)	3.360*** (0.95)	3.051*** (1.00)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	1041	1041	1491	1491	1673	1673	1598	1598
Clusters	54	54	52	52	60	60	64	64
Pseudo R_2	0.0826	0.2023	0.0115	0.1310	0.0396	0.0859	0.0079	0.0521

Notes. The dependent variable is the truth-telling probability. Logit regressions with error clustered by subject and period — each period the participants play a new game with newly randomly assigned partners. Odd columns include individual controls of age, ethnicity, career orientation, year at the university and risk tolerance.

* p -value < 0.1.

** p -value < 0.05.

*** p -value < 0.01.

Throughout the treatments reported here, we see truthful reporting moving in the direction predicted by the model. For variations in n , the changes in honesty correspond directly to similar changes in the overall quality of accepted projects. More truth-telling leads to fewer instances in which principals accept a Low quality project that was said to be High quality, as shown in Fig. 3, panels (c) and (d).

Interestingly, this relationship between honesty and efficiency does not hold as we vary r and l . While truth-telling almost doubles as the probability of receiving a low project grows, it does not follow that accepted projects will be more likely to be of high quality, as some workers may strategically report untruthfully, making their low quality projects more likely to be accepted, as shown in Fig. 3, top-left panel (a). Clearly, had we seen similar levels of truthful reports between treatments, the difference in accepted project quality would be much greater. Thus, we observe that a higher probability of receiving a low quality project increases truth-telling, but not enough to overcome the starker project environment, and so we ultimately see a greater number of low quality projects accepted under high values of r . Similarly, when the payment related to the low quality project is greater, we observe fewer low quality projects accepted after a lie (Fig. 3, top-right panel (b)). Yet more low quality projects are accepted overall, even if the agents truthfully report their low quality projects, as shown in Figure E5 in the appendix.

Altogether, we find that increased truth-telling has an efficiency-enhancing effect on project selection when the number of competitors decrease. Efficiency is unchanged when moving the payoff from the low quality project. Similarly, there is no efficiency-enhancing effect from changing the probability of receiving a low quality project. This corresponds to comparisons of average earnings, as shown in Appendix E, Figure E6.

Comparing Experiment 2 results with those from Experiment 1, we see somewhat lower levels of truth-telling on average, but a larger difference in treatments varying r (a difference of 20 percentage points in Experiment 1 versus 33 percentage points in Experiment 2) and n (for $l = 2$, the gap grows from 13 percentage points in Experiment 1 to 21 percentage points in Experiment 2. For $l = 3.5$, the gap changes from negative 7 percentage points in Experiment 1 to 24 percentage points in Experiment 2). If we restrict our comparison with Experiment 1 to the first round decisions in Experiment 2, we find almost indistinguishable levels of truth-telling between experiments for all treatments.

Experiment 2, then, illustrates that experience leads to a very gradual reduction in agent honesty. We observe muted but identifiable reductions in truth-telling from round 1 to round 10. Figures E.1 through E.4 in appendix show time trends for each treatment, both of truthful reports and beliefs. Reported beliefs of agents are consistent with these patterns, though it is notable that beliefs are always more pessimistic than reports in each round.

Truthful reporting seems to coevolve with agent beliefs, which raises an important question. Do beliefs play a causal role in truth-telling, or do beliefs evolve in response to past observations? We explore this connection further in the next section.

¹⁷ We include more detailed regression results for each set of treatments in Appendix D, all of which remain unchanged in any meaningful way from those presented here.




Fig. 3. Number of low project acceptances from dishonest messages, Experiment 2.

Notes. The figure plots the number of Low projects accepted per principal when the corresponding agent signaled that their project was High, over all ten rounds played.

6. The role of beliefs

Given the observed results in Experiment 2, one may wonder how important beliefs are in determining the truthfulness of agent reports. Prior studies have examined how positive organizational culture can impact individual behaviors such as trust (Canning et al., 2020), cooperative behavior (Ugwu and Igbende, 2017), and innovation (Sameer, 2018). Here, we look instead at how individuals’ optimistic or pessimistic beliefs about other agents’ behavior affect organizational performance.¹⁸

To shed light on this question, we classify the agents as ‘optimistic’ or ‘pessimistic’, depending on their beliefs. We re-examine our theoretical model by explicitly considering agent beliefs about others’ truthful reporting, and in doing so we find a mixed strategy equilibrium.¹⁹ We briefly illustrate the mixed strategy equilibrium below not to present it as a viable alternative to the predictions in Section 3.1, but rather to identify a beliefs threshold that allows us to categorize agents. We find a beliefs threshold for those treatments with positive incentive compatibility constraint such as those agents with beliefs above the threshold are considered optimistic, or pessimistic otherwise. This theory-based classification provides us a new lens through which to view our results and allows us to contribute to the literature on the interaction between beliefs and organizational function.

6.1. A mixed strategy equilibrium

Consider our baseline model in Section 3.1 and now conjecture that instead of truth-telling agents report truthfully with probability ρ when they receive a low project and probability 1 when they receive a high project. See Fig. 4. Given agents report this way the DM has beliefs $Pr(\theta = h | m = \text{high}) = \frac{1-r}{1-r\rho}$ when receiving a high message and beliefs $Pr(\theta = l | m = \text{low}) = 1$ when receiving a low message, and so as before it is a best response to accept when the message is high and reject when it is low. Also, define $\bar{h} = \frac{1-r}{1-r\rho}h + \frac{r(1-\rho)}{1-r\rho}l$ as the expected value of the project conditional on an agent reporting high.

¹⁸ Hamman and Martínez-Carrasco (2022) take a similar approach to show how cognitive traits like risk tolerance and deliberative thinking affect organizational structure and performance. Other research uses self-reported measures of optimism to show that optimistic firm members perform better (Hough et al., 2020) and optimistic executives invest more (Ikeda et al., 2021; Graham et al., 2013), though this may reduce performance if managers are overconfident (Hung and Tsai, 2020).

¹⁹ The belief associated with this equilibrium defines the threshold for each treatment. A mixed strategy equilibrium exists in our model owing to the discreteness of our state space while an analogous equilibrium does not exist in Schmidbauer (2017) which has a continuous state space. Importantly, this equilibrium preserves our comparative statistics with respect to the main variables under analysis.

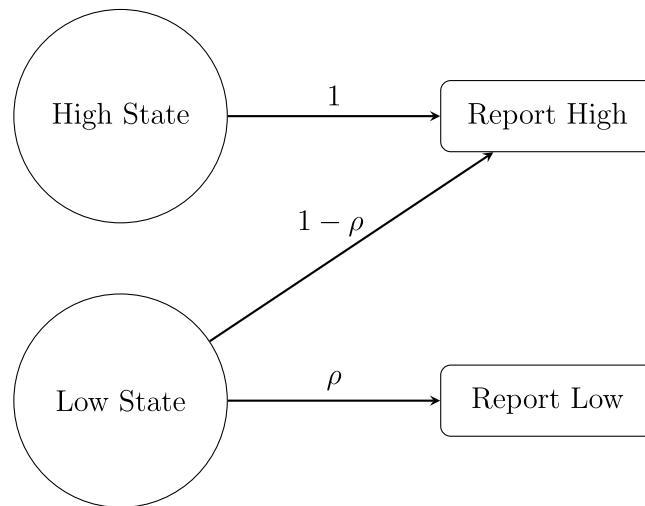


Fig. 4. A mixed strategy equilibrium in which high states are always truthfully reported but low states are truthfully reported with probability ρ .

In order for mixing to be a best response the low type agent must be indifferent in the current period between the pure strategies of reporting low or high. From an agent’s perspective, given the conjectured behavior of others, let $A(n, r, \rho)$ denote the probability of acceptance when reporting a high project.²⁰ The incentive compatibility condition becomes

$$l \times A = (r\rho)^{n-1} \left((1-r\rho)A\bar{h} + (r\rho)^n \left((1-r\rho)A\bar{h} + (r\rho)^n (\dots) \right) \right) \tag{5}$$

$$l \times A = (r\rho)^{n-1} \sum_{j=0}^{\infty} (1-r\rho)A\bar{h} (r\rho)^j$$

$$l \times A = \frac{(r\rho)^{n-1}(1-r\rho)A\bar{h}}{1-(r\rho)^n}$$

$$0 = \frac{(r\rho)^{n-1}}{1-(r\rho)^n} (1-r\rho)\bar{h} - l \equiv IC_2. \tag{6}$$

On the right hand side of the first line $1-r\rho$ represents the probability that the agent will report high next period (either because the project is high, $1-r$, or it is low but reported high, $r(1-\rho)$). Otherwise, the game either ends because some other agent’s project was adopted, resulting in a payoff of 0 for all other agents, or with probability $(r\rho)^n$ all agents report low and so the game continues yet another period, and so on. The remaining steps simply rearrange the expression. When $IC_2 > 0$ an agent strictly prefers to report high, when $IC_2 < 0$ the agent prefers to report low, and when $IC_2 = 0$ the agent is indifferent as required in a mixed strategy equilibrium.

Several important aspects of this characterization warrant mentioning. First, note that IC_2 exactly coincides with the prior IC condition in the pure strategy truth-telling equilibrium found in line (2), after replacing r with $r\rho$ and h with \bar{h} . This is intuitive because here the probability an agent reports low is $r\rho$ (he must be both low and mix on the low message) whereas in the pure strategy equilibrium the probability of a low report was r (the probability of being low, which is always reported as low in the truth telling equilibrium). Since an agent is concerned with the probability that other agents send the low message, the equivalence follows. Next, the value an agent can expect to receive in the future when reporting high and winning the contest is h in the pure strategy equilibrium and \bar{h} in the mixed strategy equilibrium. Then given the similarities of the IC conditions and since \bar{h} is increasing in h , l , and ρ , most of the comparative statistics on IC_2 are the same as IC found in line (4) (which is an equivalent formulation from line (2)) for the main treatment variables.

Proposition 3. *There exists a symmetric mixed strategy equilibrium in which each agent reports truthfully when high and when low reports truthfully with probability ρ^* , where ρ^* solves line (6), if and only if $IC > 0$ (where IC is defined in line (4)). The probability ρ^* is increasing in l and n , decreasing in h , and ambiguous with respect to r .²¹*

²⁰ Suppose there are i high reports coming from the $n-1$ other agents. This occurs when j of those agents are truly high and from among the remaining $n-1-j$ agents that are low, $i-j$ of them lie and report high. Summing over all such j from 0 to i and noting that the probability of lying is independent across agents gives $\sum_{j=0}^i \binom{n-1}{j} (1-r)^j r^{n-1-j} \binom{n-1-j}{i-j} (1-\rho)^{i-j} \rho^{n-1-j-(i-j)}$. Since ties are broken evenly the probability the focal agent’s project is accepted when there are i other reports of high is $\frac{1}{1+i}$ and thus the overall probability of acceptance is

$$A(n, r, \rho) \equiv \sum_{i=0}^{n-1} \frac{1}{1+i} \left(\sum_{j=0}^i \binom{n-1}{j} (1-r)^j r^{n-1-j} \binom{n-1-j}{i-j} (1-\rho)^{i-j} \rho^{n-1-i} \right).$$

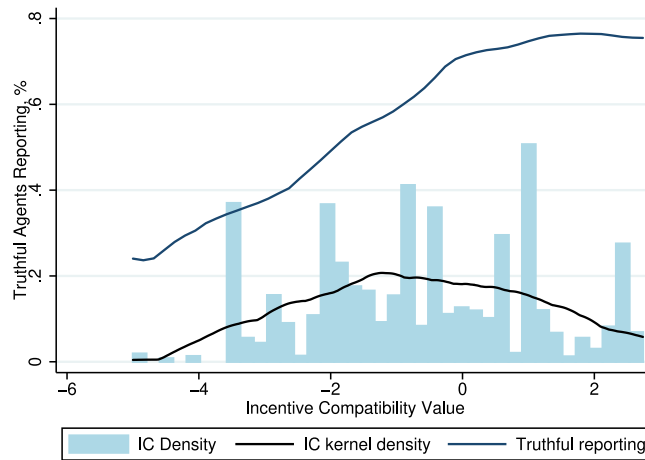


Fig. 5. Truth-telling and IC_2 constraints given beliefs by round.
 Notes. Plots of the mean truthful reporting in all the experiments by the incentive compatibility of the agents by round when they receive a bad quality project.

It is easy to show that this mixed strategy equilibrium is Pareto dominated by the pure strategy truth-telling equilibrium established in Proposition 1. In the truth-telling equilibrium any winning project is type h , implying the DM’s payoff is h . Given n agents each agent has $\frac{1}{n}$ chance of winning and thus has expected payoff $\frac{h}{n}$. In contrast, in the mixed strategy equilibrium of Proposition 3 the winning project is sometimes h but sometimes l , and so the DM’s expected payoff is less than h . Similarly, each agent’s expected payoff is $\frac{1}{n}$ of the DM’s payoff and so agents are worse off too. Put simply, lies are sometimes told and believed in the mixed strategy equilibrium whereas this does not happen in the pure strategy equilibrium.

Finally, a similarity between the mixed strategy equilibrium and an alternative model with behavioral agents is insightful to explore. Suppose each agent has independent probability ρ of being a strategic agent and $1 - \rho$ of being a behavioral type that always report High regardless of the state. These types are redrawn each period, in addition to the project types.²² Then conditional on being a strategic type with a low project, reporting truthfully will result in a non-zero payoff only if (i) every other player is a strategic type, and (ii) every other player has a low project. The probability of this is $(r\rho)^{n-1}$, which is precisely the probability that appears in line (5) above. Thus with appropriate assumptions we see that the existence of behavioral types that always report high can be accommodated within the model.

6.2. Relationship between beliefs and reports

Extending our results beyond equilibrium play we can use each participant’s self-reported beliefs of others’ behavior ρ to calculate that participant’s IC_2 . Recalling an agent strictly prefers to report high only if $IC_2 > 0$ we can then determine a participant’s best response given his conjectured (possibly out of equilibrium) behavior of the other players.

In Fig. 5, we observe a positive relationship between the value of IC_2 and average truthful reporting. The slope is much steeper in the negative side of IC_2 (64% of the sample) and flatter in the positive side of IC_2 (36% of the sample).²³ Once again, we observe that participants tell the truth more often than predicted when their IC_2 is negative and they were above 70% of truth-telling when IC_2 was positive.²⁴

We use the equilibrium ρ^* as a threshold to theoretically classify agents as either optimistic (with beliefs above the threshold) or pessimistic.²⁵ Given this classification, we analyze the differences in truthful reporting in Fig. 6. First, note that in all treatments the optimistic agents report truthfully more often than the pessimistic agents. Second, the treatment effects are higher and more significant among the pessimistic agents, while effect sizes are more muted among optimistic agents. Third, the difference between

²¹ Proof in Appendix C, subsection C.3.

²² This assumption shuts down any learning motive by the DM since an agent’s strategic type is not constant over time. Otherwise, the DM may reject all projects for several periods to update on each agent’s (non) strategic type. Further, a strategic agent who has received a high project for several periods may falsely report low to avoid being perceived as non-strategic. Such incentives would complicate matters considerably.

²³ The results are similar if we compile the graph at the participant level, as reported in Figure E.5 in Appendix E. In this case, 58% of the sample by person is on the negative side of IC_2 and 42% of the sample by person is on the positive side of IC_2 . Intuitively, pessimistic participants play more rounds of each game.

²⁴ For this graph, we use the treatments where we originally have a positive IC as required by Proposition 3. We replicate this graph with the original IC by treatment without considering beliefs. We also obtained a positive slope but not as steep as in the case of Fig. 5.

²⁵ The threshold in beliefs for the different treatments with originally positive IC are in Table D.8 in the appendix as well as the proportion of optimistic and pessimistic agents using this threshold and the median of beliefs per treatment.

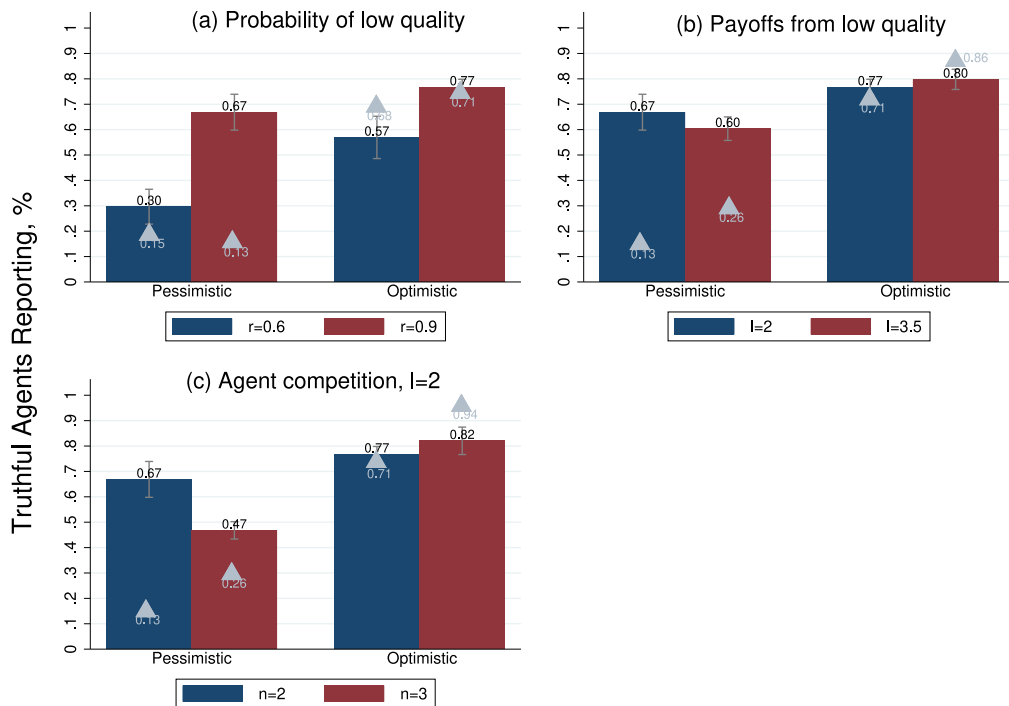


Fig. 6. Truth-telling and beliefs. Notes. Plots of the mean truthful reporting in all the experiments using the predicted threshold split by observation.

reporting and beliefs is striking among pessimistic agents, who drastically underpredict truth-telling from others.

Comparing the treatments that vary the probability of a low project, agents in $r = 0.9$ report more truthfully regardless of their beliefs. In particular, we observe a high percentage of truthful reporting from the pessimistic agents in this treatment, despite the fact that their average beliefs are similar to the pessimistic agents in $r = 0.6$. Comparing the treatments that vary the value of low payments and the number of participants, we do not observe statistical differences among optimistic agents. However, when pessimistic agents face a higher payment from the low project or a higher number of competitors, they significantly decrease truthful reporting, even though their average beliefs are not as low. Thus, we observe behaviors closer to beliefs among optimistic agents, but the treatments play a more important role among pessimistic agents.²⁶

These results suggest that the organizational structure can play a direct role in inducing more truthful reporting. However, environmental conditions that impact beliefs can also significantly impact the effectiveness of organizational decisions.

7. Discussion

While our analyses focus on behavior of the agents in our experiments, it is illuminating to briefly examine the behavior of principals before a broader discussion on future avenues of inquiry.

7.1. Principals' behavior

The model makes clear predictions of principals' decision-making in our separating equilibria. That is, in either our pure strategy truth-telling equilibrium or the threshold equilibrium discussed in the prior section, a principal should accept the project of any agent who reports that their project is High quality. In the data, we see that this does not strictly hold. Principals accept the first report of High in 58% of cases in Experiment 2 and 52% of cases in Experiment 1 (See Appendix F for more details).

The acceptance data allow us to identify several clear types of principals.²⁷ As noted, over half of our principals accept the first High signal they see, which is fully consistent with the model. Another subset of principals starts off cautiously, rejecting any message in at least the first round or so before eventually accepting a message of High. We refer to these as "cautious" principals, and they make up between 20% (Experiment 2) and 25% (Experiment 1) of our sample. A third type of principal seems to behave cautiously with respect only to agents who always signal High. These principals reject all messages for a number of rounds to see

²⁶ This result is true even if we use the median to split the sample in each treatment, as opposed to the computed threshold, or if we use the data at the person level instead of person by round. These results are shown in Appendix E in Figures E.8, E.9 and E.10 in the appendix.

²⁷ Again, specifics are relegated to Appendix F for space consideration.

that they are faced with an agent who always says High and one who occasionally says Low. After identifying this differing behavior, they accept a signal of High from the agent who more frequently signaled Low in earlier rounds. This small fraction of principals makes up between 11% (Experiment 2) and 14% (Experiment 1) of our sample.

Clearly, then, our principals are behaving in ways that closely align with the model's prediction, though with some heterogeneity in early-round behavior that seems best attributable to a degree of caution on their part.

7.2. Future directions

Our results up to this point emphasize that the organizational structure has an impact on the level of misreporting of agents with privileged information. The model highlights certain parameters through which truthful reporting can be incentivized to a greater or lesser degree. However, our findings raise two important issues that are worth considering further in future research. First, an agent's beliefs about others' behavior is an important factor in determining truthful reporting decisions. The previous section shows that pessimistic agents actually respond more strongly to the comparative statics in our treatments than do optimistic agents. Second, we observe truthful reporting between 40 and 50% of the time in treatments where the babbling equilibrium is predicted. In treatments with positive IC values, this rises to between 70 and 80%. Thus, agents are further away from the expected babbling equilibrium than they are from the expected truth-telling equilibrium, by treatment.

There are two possible behavioral explanations for this that warrant further study. First, an agent's behavior may depend on their beliefs because agents react to the expected behavior of others — in other words, an agent who anticipates more truthful reports from others may reciprocate by telling the truth themselves. This could explain why reporting by optimistic agents is almost independent of the varying IC values in our treatments.

Second, agents may be averse to dishonesty or feel a moral cost from lying. If agents vary with respect to the magnitude of this cost, we need not observe full truth-telling or full lying, as suggested by the model. However, the proportion of dishonest agents with a low project would be expected to decrease as the financial incentives to truth-telling rise, as our results show. The latter may explain why we observe more truth-telling than expected in treatments with negative IC values. In fact, these two mechanisms may even interact, such that pessimistic beliefs about others' honesty may actually lower an agent's lying costs. Taken together, these mechanisms may explain why pessimistic agents react more to parameter changes than optimistic agents. Understanding these mechanisms and their interaction with organizational structure is an important research avenue that we look forward to exploring in future work.

8. Conclusion

In this paper we experimentally investigate a model of communication between informed agents and an uninformed decision maker when agents compete for resources over time. We find evidence broadly consistent with the theoretical predictions that even when agents are completely biased to prefer their own projects, it is still possible to see truthful reports emerge as an equilibrium phenomenon.

Furthermore, our experimental results support the model's predictions that aspects of the organizational environment can successfully encourage honesty. Specifically, we suggest two strategies that firms can take to encourage truthful communication from biased agents. On one hand, increasing the disparity in earnings based on realized project quality can induce agents to be patient — passing on low quality opportunities to await stronger ones in the future. This beneficial aspect of modest inequality in earnings has been shown in other organizational settings to encourage coordination (Winter, 2004; Goerg et al., 2010). Additionally, while divisional competition can encourage innovation, our results suggest that such competition may best be constrained somewhat when project quality cannot be determined *ex ante* by firm executives. Especially in areas where high quality projects arise only rarely, these elements of organizational structure can substantially improve outcomes.

Both the theoretical model and empirical results show that beliefs play a crucial role in supporting truth-telling behavior. The gradual decline in these beliefs directly interferes with the ability of honest agents to improve organizational outcomes. This raises important questions for future research on how to maintain more optimistic (or even more realistic) beliefs on the part of decision makers and agents. Is it simply a case of “once bitten, twice shy”, in which an early experience with an unexpected negative outcome proves too difficult for decision makers to overcome? If so, it would be valuable to determine whether such early outcomes are more likely to cause impatience on the part of decision makers, or induce increased dishonesty on the part of agents. We see these as fruitful areas of research in continuing to explore the impacts — both causes and effects — of honest communication within organizations.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2025.106936>.

Data availability

Data will be made available on request.

References

- Abeler, J., Nosenzo, D., Raymond, C., 2019. Preferences for truth-telling. *Econometrica* 87 (4), 1115–1153.
- Alonso, R., Dessein, W., Matouschek, N., 2008. When does coordination require centralization? *Am. Econ. Rev.* 98 (1), 145–179.
- Armstrong, M., Vickers, J., 2010. A model of delegated project choice. *Econometrica* 78 (1), 213–244.
- Battaglini, M., 2002. Multiple referrals and multidimensional cheap talk. *Econometrica* 70 (4), 1379–1401.
- Brandt, J., Cooper, D.J., 2021. Managerial leadership, truth-telling, and efficient coordination.
- Canning, E.A., Murphy, M.C., Emerson, K.T., Chatman, J.A., Dweck, C.S., Kray, L.J., 2020. Cultures of genius at work: Organizational mindsets predict cultural norms, trust, and commitment. *Pers. Soc. Psychol. Bull.* 46 (4), 626–642.
- Chakraborty, A., Harbaugh, R., 2007. Comparative cheap talk. *J. Econom. Theory* 132 (1), 70–94.
- Crawford, V.P., Sobel, J., 1982. Strategic information transmission. *Econometrica* 1431–1451.
- Dato, S., Nieken, P., 2020. Gender differences in sabotage: the role of uncertainty and beliefs. *Exp. Econ.* 23, 353–391.
- Eckel, C.C., Grossman, P.J., 2008. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *J. Econ. Behav. Organ.* 68 (1), 1–17.
- Evdokimov, P., Garfagnini, U., 2019. Communication and behavior in organizations: An experiment. *Quant. Econ.* 10 (2), 775–801.
- Fischbacher, U., 2007. Z-tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10 (2), 171–178.
- Gilligan, T.W., Krehbiel, K., 1989. Asymmetric information and legislative rules with a heterogeneous committee. *Am. J. Political Sci.* 459–490.
- Gneezy, U., Kajackaite, A., Sobel, J., 2018. Lying aversion and the size of the lie. *Am. Econ. Rev.* 108 (2), 419–453.
- Goerg, S.J., Kube, S., Zultan, R., 2010. Treating equals unequally: Incentives in teams, workers' motivation, and production technology. *J. Labor Econ.* 28 (4), 747–772.
- Graham, J.R., Harvey, C.R., Puri, M., 2013. Managerial attitudes and corporate actions. *J. Financ. Econ.* 109 (1), 103–121.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1 (1), 114–125.
- Hamman, J.R., Martínez-Carrasco, M.A., 2022. Managing uncertainty: An experiment on delegation and team selection. *Organ. Sci.* 34 (6), 2272–2295.
- Hough, C., Sumlin, C., Green, K.W., 2020. Impact of ethics, trust, and optimism on performance. *Manag. Res. Rev.* 43 (9), 1135–1155.
- Hung, M.-W., Tsai, W.-H., 2020. Managerial optimism, CEO retention, and corporate performance: evidence from bankruptcy-filing firms. *J. Econ. Financ.* 44, 506–527.
- Ikeda, N., Inoue, K., Sugitani, S., 2021. Managerial optimism and corporate investment behavior. *J. Behav. Exp. Financ.* 30, 100492.
- Kartik, N., Ottaviani, M., Squintani, F., 2007. Credulity, lies, and costly talk. *J. Econom. Theory* 134 (1), 93–116.
- Kim, J.-Y., 1996. Cheap talk and reputation in repeated pretrial negotiation. *Rand J. Econ.* 787–802.
- Krishna, V., Morgan, J., 2001a. Asymmetric information and legislative rules: Some amendments. *Am. Political Sci. Assoc.* 95 (02), 435–452.
- Krishna, V., Morgan, J., 2001b. A model of expertise. *Q. J. Econ.* 116 (2), 747–775.
- Li, Z., 2016. Essays on Cheap Talk (Ph.D. thesis). Ohio State University.
- Li, Z., Rantakari, H., Yang, H., 2016. Competitive cheap talk. *Games Econom. Behav.* 96, 65–89.
- Moldovanu, B., Shi, X., 2013. Specialization and partisanship in committee search. *Theor. Econ.* 8 (3), 751–774.
- Rantakari, H., 2008. Governing adaptation. *Rev. Econ. Stud.* 75 (4), 1257–1285.
- Rantakari, H., 2018. Good things come to those who wait? Project selection under advocacy and durable projects. *Work. Pap.*
- Sameer, Y.M., 2018. Innovative behavior and psychological capital: Does positivity make any difference? *J. Econ. Manag.* (32), 75–101.
- Schmidbauer, E., 2017. Multi-period competitive cheap talk with highly biased experts. *Games Econom. Behav.* 102, 240–254.
- Schmidbauer, E., 2019. Budget selection when agents compete. *J. Econ. Behav. Organ.* 158, 255–268.
- Sobel, J., 1985. A theory of credibility. *Rev. Econ. Stud.* 52 (4), 557–573.
- Tergiman, C., Villeval, M.C., 2023. The way people lie in markets: detectable vs. deniable lies. *Manage. Sci.* 69 (6), 3340–3357.
- Ugwu, F.O., Igbende, D.A., 2017. Going beyond borders: Work centrality, emotional intelligence and employee optimism as predictors of organizational citizenship behavior. In: Wakefield, J. (Ed.), *Cogent Psychol.* 4 (1).
- Winter, E., 2004. Incentives and discrimination. *Am. Econ. Rev.* 94 (3), 764–773.