DISCUSSION PAPER SERIES

DP13231
(v. 2)

LOOKING INTO CRYSTAL BALLS: A LABORATORY EXPERIMENT ON REPUTATIONAL CHEAP TALK

Debrah Meloso, Salvatore Nunnari and Marco Ottaviani

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

LOOKING INTO CRYSTAL BALLS: A LABORATORY EXPERIMENT ON REPUTATIONAL CHEAP TALK

Debrah Meloso, Salvatore Nunnari and Marco Ottaviani

Discussion Paper DP13231
First Published 09 October 2018
This Revision 21 November 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Debrah Meloso, Salvatore Nunnari and Marco Ottaviani

LOOKING INTO CRYSTAL BALLS: A LABORATORY EXPERIMENT ON REPUTATIONAL CHEAP TALK

Abstract

We experimentally study cheap talk by reporters motivated by their reputation for being well informed. Evaluators assess reputation by cross checking the report with the realized state of the world. We manipulate the key drivers of misreporting incentives: the uncertainty about the state of the world and the beliefs of evaluators about the strategy of reporters. Consistent with theory, reporters are more likely to report truthfully when there is more uncertainty and when evaluators conjecture that reporters always report truthfully. However, the experiment highlights two phenomena not predicted by standard theory. First, a large fraction of reports is truthful, even when this is not a best response. Second, evaluators have di�culty learning reporters' strategies and overreact to message accuracy. We show that a learning model where accuracy is erroneously taken to represent truthfulness ts well evaluators' behavior. This judgement bias reduces reporters' incentives to misreport and improves information transmission.

JEL Classification: C91, D83

Keywords: Forecasting, experts, reputation, cheap talk, Laboratory experiments

Debrah Meloso - d.meloso@tbs-education.fr
Toulouse Business School

Salvatore Nunnari - salvatore.nunnari@unibocconi.it
Bocconi University and CEPR

Marco Ottaviani - marco.ottaviani@unibocconi.it
Bocconi University and CEPR

Acknowledgements
A previous version of this work circulated under the title \The Mechanics of Reputational Cheap Talk: An Experiment with Crystal Balls." We are grateful to Nageeb Ali, Anthony Kwasnica, Jacopo Perego, Chloe Tergiman and seminar audiences at various seminars and conferences for helpful comments and interesting discussions. Luca Ferocino, Antonio Giannino, Giacomo Saibene, Francesca Zambrini, and Giovanni Montanari provided excellent research assistance. We gratefully acknowledge nancial support from the European Research Council through ERC Grant 295835 (EVALIDEA).

Looking into Crystal Balls: A Laboratory Experiment on Reputational Cheap Talk*

Debrah Meloso
Toulouse Business School

Salvatore Nunnari
Bocconi University

Marco Ottaviani
Bocconi University

November 21, 2019

Abstract

We experimentally study cheap talk by reporters motivated by their reputation for being well informed. Evaluators assess reputation by cross checking the report with the realized state of the world. We manipulate the key drivers of misreporting incentives: the uncertainty about the state of the world and the beliefs of evaluators about the strategy of reporters. Consistent with theory, reporters are more likely to report truthfully when there is more uncertainty and when evaluators conjecture that reporters always report truthfully. However, the experiment highlights two phenomena not predicted by standard theory. First, a large fraction of reports is truthful, even when this is not a best response. Second, evaluators have difficulty learning reporters' strategies and overreact to message accuracy. We show that a learning model where accuracy is erroneously taken to represent truthfulness fits well evaluators' behavior. This judgement bias reduces reporters' incentives to misreport and improves information transmission.

JEL Classification: C72, C91, C92, D83, D91
Keywords: Forecasting; Experts; Reputation; Cheap Talk; Laboratory Experiments.

*A previous version of this work circulated under the title "The Mechanics of Reputational Cheap Talk: An Experiment with Crystal Balls." We are grateful to Nageeb Ali, Anthony Kwasnica, Jacopo Perego, Chloe Tergiman and seminar audiences at various seminars and conferences for helpful comments and interesting discussions. Luca Ferocino, Antonio Giannino, Giacomo Saibene, Francesca Zambrini, and Giovanni Montanari provided excellent research assistance. We gratefully acknowledge financial support from the European Research Council through ERC Grant 295835 (EVALIDEA).

1 Introduction

Forecasting is a thriving industry in economics, finance, and politics. Experts' predictions of future trends constitute the basis of policy prescriptions, investment decisions, and firm management. As a consequence, forecaster accuracy is actively monitored, and forecasters who attain outstanding reputation face remarkable career prospects.[1] The strength of these incentives might lead one to believe that reputation motives and market forces should ensure performance and truthfulness of forecasters. As reported by Keane and Runkle (1998), "since financial analysts' livelihoods depend on the accuracy of their forecasts [...], we can safely argue that these numbers accurately measure the analysts' expectations."

Building on models of career concerns by Holmström (1999) and herding by Scharfstein and Stein (1990), this belief has been challenged by a large theoretical literature positing that forecasters are economic agents who make strategic choices and may be reluctant to release truthful information that could be considered inaccurate. In its basic structure (Ottaviani and Sørensen, 2006a), the strategic situation studied in this literature generates a game of *reputational cheap talk* between a *reporter* and an *evaluator*. The reporter privately observes a signal about a state of the world and reports a message to the evaluator. The informativeness of the reporter's signal is uncertain and initially unknown to both the reporter and the evaluator. The evaluator assesses the informativeness of the reporter's signal on the basis of the reporter's message and the realized state of the world. The objective of the reporter is to maximize the reputation for being well informed, according to the assessment made by the evaluator. Variants of the reputational cheap talk game have been extensively used in the applied theory literature to model the strategic incentives for reputation management in the context of financial recommendations, news reporting, and communication in organizations (Trueman, 1994; Ehrbeck and Waldmann, 1996; Graham, 1999; Ottaviani and Sørensen, 2001, 2006a,b; Lamont, 2002; Prat, 2005; Gentzkow and Shapiro, 2006; Levy, 2007; Visser and Swank, 2007; Deb et al., 2018).

Testing the predictions of these theories has proven challenging. For example, it is hard to measure, let alone manipulate, the information available to reporters and evaluators. This leaves us with many interesting open questions: Do these models accurately predict behavior? Do reporters misreport available information to appear competent? How does

[1] For example, Alan Greenspan and Lawrence Meyer ran successful consulting firms offering forecasting services before becoming key members of the Board of Governors of the Federal Reserve Bank. In a different domain, after successfully calling the winner in 49 states in the 2008 U.S. presidential election, Nate Silver was named one of The World's 100 Most Influential People by Time, licensed his blog for publication in the New York Times, and eventually sold it to ESPN.

this depend on uncertainty and evaluators' expectations? Are evaluators able to interpret forecasts or are they naïve? This paper develops an experimental framework to address these questions. With our innovative design, we are able to measure and exogenously manipulate the information structure, the degree of uncertainty over the forecasted variable, the evaluators' expectations, and the evaluators' learning process.

A first challenge for the experimental design is to find a simple way to implement the information structure posited by the theory. We meet this challenge by developing a novel urn scheme with nested balls, which builds on the classic urn paradigm pioneered by Anderson and Holt (1997). As in the classic setting, the private signal observed by the reporter corresponds to the color of a ball drawn from an urn (either blue or orange). We innovate by introducing an inner core inside the outer shell of each ball; the color of the core—which is not visible to the reporter—represents the state of the world (either blue or orange). We then capture the informativeness of the signal about the state by constructing urns containing a different composition of nested balls: an informative urn contains balls whose shell and core have the same color; in an uninformative urn, instead, the color of the shell gives no indication of the color of the core.

A second challenge for the experimental design is to find a way to control for evaluators' beliefs about reporter truthfulness, a key driver of reporting incentives. Indeed, these beliefs are difficult to pin down for both the experimenter analyzing the data and reporters engaged in the experimental task. This is due to the existence of multiple equilibria and to the high level of strategic sophistication required for equilibrium play in these games. We meet this challenge by designing treatments where human reporters face computerized evaluators and by introducing a novel model of Bayesian learning which determines how the beliefs of some computerized evaluators evolve over time.

In particular, to study reporters' strategic incentives, we employ eight different treatments, manipulating two crucial dimensions of the game: the common prior belief on the state of the world, q, and the evaluators' expectations. We consider two values of q, chosen to generate different predictions about reporter behavior. We vary how we control for evaluators' beliefs through four games: a game with *computerized* evaluators programmed to believe that all reporters always report *truthfully* (CT); a game with *computerized* evaluators programmed to believe that the fraction of reporters who report truthfully is *uniformly* distributed (CU); a game with *computerized* evaluators whose beliefs are programmed to evolve according to Bayesian *learning*, based on the outcome of past individual interactions with reporters (CL); and a game with *human* evaluators who have *free* beliefs about reporters' behavior (HF).

Our design isolates two forms of uncertainty about evaluators' behavior that may affect reporters' choices: when facing human evaluators, reporters are uncertain both about the beliefs evaluators may hold and about whether they best reply to held beliefs. Games CT and CU eliminate both forms of uncertainty by fixing evaluators' beliefs and, essentially, transform the reputational cheap talk game in a decision problem for reporters. Game CL eliminates only the latter form of uncertainty—since, there, evaluators best reply to evolving beliefs—but allows us to control for learning on the side of evaluators. The tractable learning process we develop for game CL—a noisy Bernoulli learning model whose conjugate prior is a generalized Beta distribution—is of independent interest and also serves as a benchmark to study learning by human evaluators.[2]

According to the theory of reputational cheap talk, prior belief about the state of the world and evaluators' beliefs about reporter truthfulness jointly determine which, among two forces driving reporting incentives, should prevail. Intuitively, these two forces emerge because the reporter uses the signal to simultaneously update beliefs about the state of the world and about the informativeness of the signal. To understand the consequences of the double role of the reporter's signal, suppose that the prior about the state of the world is unbalanced in favor of blue (as is always the case in our experiment) and that the evaluator believes that the reporter truthfully reports the observed shell (as in our CT game). If the evaluator's assessment were formed just on the basis of the report but without observing the state of the world, the reporter would always want to report a signal corresponding to the most likely state (blue), so as to improve the evaluator's assessment that the signal was informative (blue shells are more abundant in the informative urn only). This first force pushes toward misreporting whenever the reporter's signal corresponds to the least likely state of the world (that is, he observes an orange shell). But the evaluator's assessment is also based on the realized state of the world (the core). The second force then comes into play because the reporter's signal always increases the belief that the state of the world will be equal to the observed signal (that is, the core observed by the evaluator will have the same color as the observed shell). Since a matching signal and state of the world are indicative of an informative urn, this second force creates an incentive for the reporter to truthfully report the observed signal.

[2]For other uses of computerized opponents in experimental economics, where a game is reduced to a decision problem, see Roth and Murnighan (1978), Fehr and Tyran (2001), Esponda and Vespa (2014), and Koch and Penczynski (2018). In none of these experiments, computerized players are programmed to learn about the strategy of their human opponents on the basis of their experience. The learning model and the experimental methodology we introduce are portable to other games of asymmetric information.

If the evaluator believes that the reporter is always truthful, the second force prevails when the reporter, upon observing an orange shell, thinks that the core is more likely to be orange than blue. This is the case when the prior is mildly unbalanced; then there exists a separating Bayesian Nash equilibrium in which the reporter truthfully reports the observed shell and the evaluator believes in the report. If, instead, the prior is strongly unbalanced, there is only a pooling equilibrium in which the reporter reports a blue shell even when observing an orange shell and the evaluator disregards the report.

Empirically, we find that, as predicted by the theory, reporters are more likely to truthfully report information supporting the *ex-ante* unlikely state when they are less certain about the state of the world and when evaluators always expect them to report truthfully. We also find that human evaluators appropriately use the information they receive about signal informativeness (report and core): assessments based on different report and core combinations are ranked in the order predicted by theory.

At the same time, the experiment highlights phenomena that are not accounted by the standard theory. First, around a third of reports are truthful, even when this is not a best reply and even when reporters have gained experience with the task. Cluster analysis reveals that this can be partly attributed to a subset of subjects who is insensitive to the uncertainty about the state of the world and reports truthfully under both values of q.

Second, we find human evaluators' behavior to be incompatible with beliefs based on correct inference about reporters' strategies in different environments. In particular, we find that human evaluators incorrectly weigh reporter accuracy—that is, predicting the state of the world by matching the color of the core—and inaccuracy—that is, mispredicting the state of the world. Accuracy and inaccuracy carry information about signal informativeness only when reporters are truthful; instead, evaluators' assessments are more reactive to accuracy and inaccuracy exactly when reporters are predicted and found to be less truthful (when the prior is strongly unbalanced). This cannot be simply attributed to confusion about the rules of the game or inability to do Bayesian updating: we show that this result is driven by a subset of subjects whose assessments are indistinguishable from those of a Bayesian learner when the inference does not depend on reporters' strategies. This behavior may reflect a higher willingness of evaluators to reward accuracy or punish inaccuracy at a cost (for example, because of reciprocal motives, as in Rabin 1993 and Dufwenberg and Kirchsteiger 2004), when they detect more misreporting. Alternatively, we show that a learning model where accuracy is erroneously taken to represent truthfulness and inaccuracy to represent misreporting also generates the above-mentioned overreaction.

In turn, evaluators' observed behavior reduces reporters' incentive to misreport. Consider the treatment with strongly unbalanced priors. In the unique equilibrium of the game, the reporter always misreports and, thus, a Bayesian evaluator's assessment of the signal informativeness equals the balanced prior, regardless of the signal accuracy. In the experiment, the average assessment following an inaccurate (accurate) blue report is 30% (56%). While misreporting remains the best response, the gain with respect to truth-telling is dampened by the harsh punishment associated with an inaccurate report. If reporters' choices are stochastic rather than deterministic and respond continuously to incentives (as in a random utility model or a quantal response equilibrium), this leads to a positive incidence of truth-telling, even among subjects who are not averse to lying.

These results have important implications for models of reputational cheap talk and the design of markets for professional forecasting. Overall, our experiment suggests that current models of reputational cheap talk correctly capture reporters' behavior but might be missing important elements in the way evaluators process the available information or reward reporters for their advice. It also suggests that making experts' ex-post accuracy (rather than experts' advice) a salient element of the information available to clients might ameliorate forecasters' performance and the transmission of information.[3]

While there is a large experimental literature on the statistical herding models of Banerjee (1992) and Bikhchandani et al. (1992),[4] this is the first paper testing experimentally the basic building block of the reputational herding model of Scharfstein and Stein (1990) with a single sender. Other recent experimental papers in this area (Fehrler and Hughes, 2018; Renes and Visser, 2018; Mattozzi and Nakaguma, 2019) focus on situations with multiple senders. Fehrler and Hughes (2018) and Mattozzi and Nakaguma (2019) experimentally examine the role of transparency when career-concerned experts, differently from our setup, are privately informed about the informativeness of their own signal and make a decision on behalf of the evaluator. In a similar setting, Renes and Visser (2018) consider the case in which the committee experts care both about their reputation and the quality of collective decision making. Koch et al. (2009) and Irlenbusch and Sliwka (2006) conduct experiments based on Holmström's (1999) career concerns model. Contrary to our experiments, where we manipulate the experts' incentives to misreport, the main experimental treatment of these

[3] An example of such focus on accuracy is TipRanks (www.tipranks.com), a dataset of analysts, hedge fund managers, financial bloggers, and corporate insiders. The site uses Natural Language Processing algorithms to aggregate and analyze financial data online.

[4] Anderson and Holt (1997) pioneered the investigation of informational cascades in the laboratory. Their work was extended, among others, by Hung and Plott (2001), Kübler and Weizsäcker (2003), Çelen and Kariv (2004), Goeree et al. (2007), and Eyster et al. (2015). Anderson and Holt (2008) provide an excellent review.

works is the information available to the evaluator before making his assessment.

With the exception of Fehrler and Hughes (2018) and Renes and Visser (2018), where experts are allowed to send messages to each other, in these experiments there is no communication among agents. Thus, differently from our work, in none of these studies do experts send a message to evaluators about the information they possess.[5]

While in our setting the sender cares about his reputation, in the experimental literature testing models of cheap talk,[6] the sender cares about a decision taken by the receiver, with the preference alignment between the sender and the receiver as the key driver of information transmission. The experimental literature on "naïve advice" also explores the determinants and consequences of advice transmitted from informed senders who internalize, at least partially, the receiver's well being (Schotter, 2003; Schotter and Sopher, 2007; Chaudhuri et al., 2009; Çelen et al., 2010). Given that the reporter's utility depends on the beliefs of the evaluator about the reporter's type, reputational cheap talk is a psychological game (Geanakoplos et al., 1989; Battigalli and Dufwenberg, 2009). In typical applications of psychological games, the effect of beliefs on utility is mediated by emotions, thus reducing the appeal of direct belief manipulation via computerized agents (our work) or payoff distributions (Khalmetski, 2016; Ederer and Stremitzer, 2017), in favor of the elicitation and communication of beliefs (Ellingsen et al., 2010), or their indirect manipulation (Dufwenberg and Gneezy, 2000; Charness and Dufwenberg, 2006).

The paper proceeds as follows: Section 2 introduces the theory underpinning our experimental investigation and explains the design of the experiment. Section 3 derives in detail the testable hypotheses we take to the laboratory. Section 4 presents the experimental results. Section 5 concludes. The Appendix develops a novel and tractable learning model that combines a generalized Beta distribution with a noisy Bernoulli outcome; this learning model plays a key role in the experimental design and analysis, but is also of independent interest. The Online Appendix contains proofs, supplementary empirical results, and full experimental instructions.

[5]In some treatments of Renes and Visser (2018), experts can transmit to evaluators a statement about their confidence in the committee decision.

[6]See Blume et al. (Forthcoming) for a comprehensive review.

2 Model and Experimental Design

2.1 Model

We consider a Bayesian game of reputational cheap talk between a *reporter* and an *evaluator*. The model has been explicitly designed to capture the key issues from the reputational cheap talk literature, while at the same time keeping it simple enough to investigate its predictions in the laboratory.

The reporter and the evaluator are uncertain about a state of the world (corresponding in the experiment to the color of the *core* of the ball), which can be either *blue* or *orange*, $c \in \{b, o\}$. The common prior belief is unbalanced towards state b, $\Pr(c = b) = q \geq 1/2$.[7] The reporter privately observes a signal about the state (the color of the *Shell*), which can be either *Blue* or *Orange*, $S \in \{B, O\}$. There are two types of reporters (*urns* from which signals are drawn), $u \in \{I, U\}$: reporters with $u = I$ receive perfectly *informative* signals, meaning that core and Shell color are equal for sure. Reporters with $u = U$ receive perfectly *uninformative* signals, meaning that regardless of the color of the Shell, the core is equally likely to be blue or orange. The reporter and the evaluator are uncertain about the signal's informativeness and, ex-ante, believe that both possibilities are equally likely, $\Pr(u = I) = \Pr(u = U) = 1/2$.

After observing the signal, the reporter sends a report to the evaluator, $R \in \{B, O\}$. The reporter is unable to prove the signal received, so R is a *cheap talk* message. In the theoretical analysis as well as in the experiment, we constrain the reporter to report $R = B$ after observing $S = B$ and we allow *misreporting* only after observing $S = O$.[8] Thus, the reporter has two possible strategies:

- *Misreporting* (M): always report $R = B$, regardless of the signal.
- *Truth-telling* (T): report $R = B$ when $S = B$; report $R = O$ when $S = O$.

After observing the report R and the state of the world c, the evaluator assesses the likelihood that the reporter was informed (i.e., that the signal was drawn from the informative urn): $\Pr(u = I|R, c) = p_{Rc}$. The reporter benefits from being perceived as informed with a payoff proportional to this *assessment*. We assume the reporter to be risk neutral, with expected utility from either strategy proportional to the expected evaluator's assessment,

[7] Since the model is perfectly symmetric with respect to c, this is without loss of generality.

[8] Following observation of $S = B$, for $q \geq 1/2$, there is no belief of the evaluator about the reporter's strategy for which the reporter finds it optimal to report $R = O$. Thus, to simplify the analysis and the experimental task, we do not allow for this kind of misreporting.

$E[p_{Rc}]$. The reporter, who does not know the state of the world when making the report, perceives the evaluator's assessment as a random variable taking the value p_{Rb} if the state of the world is b, and p_{Ro} if the state of the world is o. Thus, if the reporter sends $R = O$, the evaluator's assessment will be either p_{Ob} or p_{Oo}; if the reporter sends $R = B$, the reporter induces assessments p_{Bb} or p_{Bo}.

The evaluator's objective is to make an accurate assessment of the reporter's informativeness. This assessment depends on the evaluator's belief about the probability the reporter is truthful, denoted by f. Conditional on this belief f, on the received report R, and on the observed state c, the evaluator has an incentive to make the most accurate possible assessment.[9]

2.2 Experimental Implementation

The previous section describes the game played in one period of our experiment. In this section, we further specify experimental subject payoffs, the structure of a session (composed of several periods), and the structure of the entire experiment (composed of several sessions and treatments).

2.2.1 Structure of a Session and Payoffs

Periods and Blocks. An experimental session consists of 4 blocks of 16 periods (for a total of 64 periods) to allow learning. Each reporter is randomly re-matched with an evaluator at the beginning of each period. The value of q is fixed during a block but it changes from one block to the next (within-subject treatment variation), so that each value occurs in two non-consecutive blocks. We use the term *first block* to indicate the block in which a given value of q is encountered for the first time, and *second block* to indicate the block in which this same q is encountered for the second time.

[9] The reduced-form payoffs we posit can be derived by appending a second period in which the same game is played again, following a construction formulated in Holmström (1999) and further developed by Scharfstein and Stein (1990). Before the second period starts, the report sent in the first period and the realized state are publicly observed by at least two evaluators who compete to hire the reporter. Given that this second period is also the last, the reporter has no incentives to lie and so can be safely assumed to report truthfully. When hiring a reporter, each evaluator obtains a decision payoff that increases in the informativeness of the reporter, because a more informed reporter truthfully sends a more informed report. This justifies evaluators in our setting being paid by the accuracy of their assessment. Because of competition among evaluators, the reporter is paid the expected value of her information, thus justifying that the compensation of reporters is the evaluator's assessment of their informativeness.

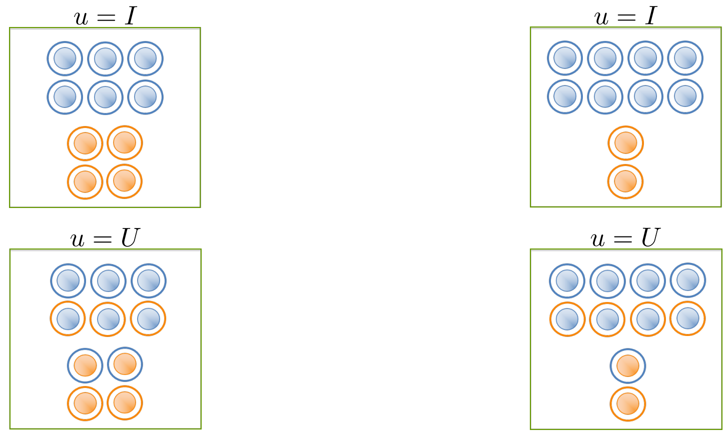

(a) $q = 6/10$ (b) $q = 8/10$

Figure 1: Informative and uninformative urns with a mildly unbalanced prior—6 balls have a blue core—in panel (a), and a strongly-unbalanced prior—8 balls have a blue core—in panel (b).

Information Structure. The urns and nested balls structure is implemented with a computerized toss of a fair coin to determine the urn type, $u \in \{I, U\}$, and computerized selection of a nested ball from the chosen urn. Draws are independent across reporter-evaluator pairs and periods. Urns in our experiment are composed of 10 balls each, among which $Q = 10q$ have a blue core and the remainder an orange core. Figures 1a and 1b, which are given to experimental subjects as part of their instructions, represent the urn composition for, respectively, $q = 6/10$ and $q = 8/10$. Consider our *mildly unbalanced* prior of $q = 6/10$ (Figure 1a). The informative urn in this case, is composed of six balls with a blue shell and a matching blue core and of four balls with an orange shell and a matching orange core. The uninformative urn also contains six balls with a blue core and four balls with an orange core; however, among the six balls with a blue core, only three have a blue shell and, similarly, among the four balls with an orange core, only two have an orange shell. Notice that the uninformative urn always contains five balls with an orange shell and five with a blue shell. As the prior probability of the blue state of the world increases to generate a *strongly unbalanced* prior, $q = 8/10$ (Figure 1b), the number of balls with a blue core increases to eight in both urns, but the number of balls with a blue shell only increases in the informative urn.

Choices. As in the model, reporters in the experiment can choose to truthfully report the received signal always, or to misreport by sending a report $R = B$ when they observe $S = O$. We use the strategy method to elicit reporters' strategies as a plan of action set out before they observe the drawn ball's shell. Reporters in the experiment thus choose one of the following two plans of action:

1. If I see a BLUE shell, I will report: "The shell is BLUE".
   If I see an ORANGE shell, I will report: "The shell is ORANGE".

2. If I see a BLUE shell, I will report: "The shell is BLUE".
   If I see an ORANGE shell, I will report: "The shell is BLUE".

The first plan corresponds to truth-telling and the second plan to misreporting. After receiving the report and observing the realized state of the world (core of the drawn ball), evaluators are asked to assess the probability that the ball was drawn from the informative urn.[10]

Payoffs. At the beginning of each block of periods, the reporter receives a budget of €4. In each period, the reporter pays an operating fee of €0.25 and obtains a payoff equal to €P, where $P \in [0, 1]$ represents the evaluator's assessment of the probability that the ball was drawn from the informative urn. The operating fee allows for negative payoffs in periods where received assessments are particularly low, which may stimulate best reply behavior. The block budget of €4 ensures that reporters will never have negative final payoffs, since negative payoffs are hard to enforce in the laboratory.[11] We compensate human evaluators using a binarized scoring rule (Harrison et al., 2013; Hossain and Okui, 2013). Once the scoring rule is binarized, it is optimal for evaluators to truthfully report their personal assessment, even if they are risk averse or do not have expected utility preferences, provided that they prefer a binary lottery that assigns a higher probability to the larger reward to one that assigns a lower probability to the same reward. An exact description of this scoring rule can be found in the experimental instructions in Online Appendix C.

Feedback. At the end of each period, each reporter receives individual feedback about the color of the core of the ball, the type of urn from which the ball was drawn, and the evaluator's assessment (that is, the reporter's payoff). Similarly, at the end of each period, each evaluator receives feedback on the type of urn from which the ball was drawn and the evaluator's payoff. We give this feedback to reporters and evaluators to allow them to gain experience and learn how to play the game. Given that evaluators observe neither the

[10]To assist human evaluators in their assessment, the software provides a slider and computes the payoff for different provisional assessments made by subjects.

[11]See Feltovich (2011) for a thorough discussion of the use of negative payoffs in economics experiments.

strategy of reporters nor the color of the shell, feedback given to evaluators is also meant to help them learn the strategy used by reporters.

2.2.2 Experimental Treatments and Sessions

Experimental Treatments. Experimental treatments vary along two dimensions, between two values of the prior probability of the state of the world, and between four forms of experimental control of evaluator beliefs, thus yielding eight distinct treatments. The prior probability of the state of the world is varied between a *mildly unbalanced* prior with $q = 6/10$, and a *strongly unbalanced* prior with $q = 8/10$. Experimental control of evaluator beliefs, f, is varied between four *treatment games*, described below. The value of q varies within subject, since in each session subjects are confronted with both, while the treatment game varies across subject (that is, one treatment game per session).

Treatment Games. The following treatment games are implemented:

- HF *(Human Free)*. Evaluators are human subjects with *free* beliefs about reporters' behavior.

- CT *(Computerized Truthful)*. Given a received report and an observed core, computerized evaluators in this game assess the probability that a drawn ball comes from an informative urn as if they believed reporters are always truthful.

- CL *(Computerized Learning)*. Computerized evaluators in this game are programmed to learn about reporter truthfulness through experience. Evaluators are endowed with a learning model that assumes there is a true, fixed, but unknown fraction f of reporters that report truthfully. Evaluators initially hold a uniformly-distributed belief over the true value of f and update their belief with signals given by the triple of information that any evaluator (human or not) has after an interaction with a reporter: the received report, the observed core, and the true informativeness of the urn. Some triples constitute positive signals, others negative signals, and others are noisy signals of the reporter's truthfulness. Bayesian updating of beliefs using these signals generates an Exton-generalized Beta conjugate learning model (see Section 4.2.3 and the Appendix).

- CU *(Computerized Uniform)*. Given a received report and an observed core, computerized evaluators in this game assess the probability that a drawn ball comes from

Game	Sessions	Decisions (R-E Pairs)		Subjects	
		$q = 6/10$	$q = 8/10$	Reporters	Evaluators
HF	4	1,504	1,504	47	47
CT	2	1,504	1,504	47	0
CU	2	1,504	1,504	47	0
CL	4	1,536	1,536	48	0

Table 1: Experimental Design. R-E stands for Reporter-Evaluator.

an informative urn as if they believed the probability that reporters are truthful were uniformly distributed, $f \sim U[0,1]$. This treatment is intermediate between CT and CL: like in CT, beliefs are fixed; like in CL, the starting belief is a uniform distribution over f.

Treatments and Sessions. We committed ex-ante to the described treatments. We ran at least two sessions with the same treatment game, and a total of 12 sessions. Each session lasted around two hours. To control for order effects in learning, we ran an equal number of sessions with order 6868—initial block had $q = 6/10$—and order 8686—initial block had $q = 8/10$. The number of sessions, subjects and decisions for each treatment game is reported in Table 1.

Subjects and Realized Earnings. All experimental sessions were conducted at Bocconi Experimental Laboratory for the Social Sciences (BELSS). Subjects' age ranged between 18 and 30 (with an average of 21) and 45.76% of participants were female. Subjects participated in only one session and maintained their role of reporter or evaluator throughout the whole session. Average earnings (including a €5 show-up fee) were €36.05 for reporters and €39.66 for evaluators. Full experimental instructions were read out with the help of a video. They are reported in Online Appendix C. The experiments were programmed using zTree (Fischbacher, 2007).

3 Theoretical Analysis and Experimental Hypotheses

Equilibrium predictions, carefully presented in this section, constitute an important baseline for our experiment. Another important baseline is that of reporter incentives and best-reply behavior given the parameters q and f that we manipulate across treatments. We commence with this analysis.

Reporter Gain from Misreporting. Our analysis of reporter best-reply behavior centers around the gain a reporter expects to obtain if she misreports instead of telling the truth after observing signal $S = O$, as a function of q and f. This magnitude, which we denote $\Delta_{EU}(q, f)$, captures all expected payoff differences between the two strategies, since truth-telling and misreporting differ only if the signal is O. After observing $S = O$, the reporter weighs evaluator assessments given a report and a state of the world, using posterior probability q_O that the state of the world is b. Noticing that q_O depends on the prior, q, and that evaluator assessments, p_{Rc}, depend on their belief about reporter truthfulness, f, we obtain

$$\Delta_{EU}(q, f) = \underbrace{[(1 - q_O(q))\, p_{Bo}(f) + q_O p_{Bb}(f)]}_{=EU^M} - \underbrace{[(1 - q_O(q))\, p_{Oo}(f) + q_O p_{Ob}(f)]}_{=EU^T}, \quad (1)$$

where EU^M and EU^T denote the reporter's expected payoff from misreporting or truth-telling after observing an Orange shell.

3.1 Evaluator's Assessments

If the evaluator could observe the signal and the state of the world, he would always update his prior belief about signal quality based on whether signal and state matched or not. Specifically, a matching signal and state of the world is twice as likely to occur if the signal is informative than if it is uninformative, and a mismatch between signal and state of the world is impossible if the signal is informative. This logic applies without qualifiers when the evaluator receives a report $R = O$, since reporters are not allowed to use this report falsely. Thus, in this case, received report equals observed shell, and the posterior from a match is $p_{Oo} = 2/3$, and from a mismatch is $p_{Ob} = 0$.

Instead, when the evaluator receives a report $R = B$, he knows the signal may have been $S = O$, since the reporter may have misreported. The observation of *report* and state of the world, thus, differs in its informational content from the observation of *signal* and state of the world. Specifically, given report B, state of the world c, and belief f, by Bayes' rule the evaluator's assessment that the reporter is informed is

$$p_{Bc} = \Pr(u = I | R = B, c, f) = \frac{\Pr(R = B | c, u = I, f)\,\frac{1}{2}}{\Pr(R = B | c, u = I, f)\,\frac{1}{2} + \Pr(R = B | c, u = U, f)\,\frac{1}{2}}. \quad (2)$$

The evaluator's assessments after a Blue report crucially depend on the likelihood that the reporter is truthful, f. A belief that the reporter may misreport, $1 - f > 0$, dampens the evaluator's favorable inference about informativeness following a match between report and state, p_{Bb}, as well as the unfavorable inference following a mismatch, p_{Bo}. From (2), we have

$$p_{Bb} = \frac{1}{\frac{3}{2} + (1-f)\frac{1}{2}} \text{ and } p_{Bo} = \frac{(1-f)}{\frac{1}{2} + (1-f)\frac{3}{2}}.$$

The evaluator's assessments are most extreme and informative when the reporter is always truthful, $f = 1$. If the evaluator expects that the reporter truthfully reports $R = O$ following $S = O$ with lower probability f, the information content of the report is reduced and the assessments that the urn is informative following reports $R = B$ are dampened toward $p = 1/2$. When the reporter is believed to always misreport, $f = 0$, the posterior assessments must be equal to the prior, $p_{Bb} = p_{Bo} = 1/2$. We conclude:

Proposition 1 (Ranking of Evaluator's Assessments) ***(a)*** *Assessment p_{Bb} is strictly increasing in f, with $p_{Bb} \in [1/2, 2/3]$; p_{Bo} is strictly decreasing in f, with $p_{Bo} \in [0, 1/2]$.* ***(b)*** *The evaluator's assessments are weakly ranked: $p_{Oo} \geq p_{Bb} \geq p_{Bo} \geq p_{Ob}$ for any $f \in [0, 1]$. If $f \in (0, 1)$, this ranking is strict: $p_{Oo} > p_{Bb} > p_{Bo} > p_{Ob}$. If $f = 0$, $p_{Oo} > p_{Bb} = p_{Bo} > p_{Ob}$. If $f = 1$, $p_{Oo} = p_{Bb} > p_{Bo} = p_{Ob}$.*

3.2 Reporter's Incentives

Equation (1) and derivations in the previous section tell us that the two main drivers of reporters' behavior are the common prior belief on the state of the world, q, and the evaluator's belief on the reporter's strategy, f. Our experimental design exploits this observation by varying q and manipulating f via computerized evaluators in three of our four treatment games. The following analysis of the way in which these two factors affect the reporter's propensity to misreport, underlies our experimental hypotheses. In Section 3.3 we derive the equilibrium f as a function of q.

Comparative Statics with Respect to q. When the state of the world is more likely to be b, the reporter is more confident that a report $R = B$ will match the state and that a report $R = O$ will not, regardless of the signal received. As q increases, the posterior belief that the state of the world is b, q_O, increases and, thus, a reporter who misreports expects a favorable assessment, p_{Bb}, to become more likely relative to an unfavorable one, p_{Bo}. On the one hand, this means that the expected utility from misreporting increases in q, given that $p_{Bb} \geq p_{Bo}$ regardless of $f \in [0, 1]$. On the other hand, as q increases, truth-telling becomes

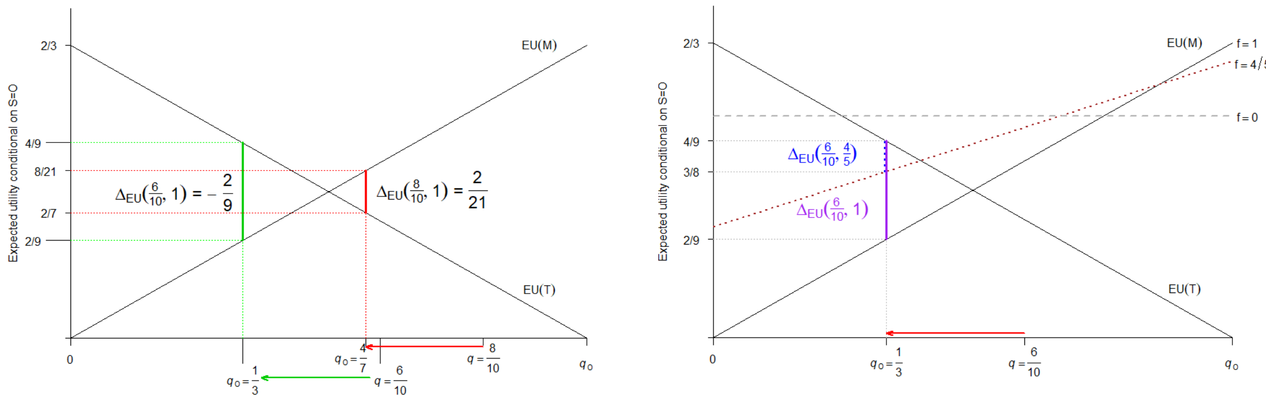

(a) Changing q, fixed $f = 1$.

(b) Changing f, fixed $q = 6/10$.

Figure 2: Reporter's expected gain from misreporting. Comparative statics with respect to q and f.

less attractive, because the evaluator is more likely to assess p_{Ob} and less likely to assess p_{Oo}, where $p_{Oo} > p_{Ob}$ for all $f \in [0, 1]$. More generally:

Proposition 2 (Comparative Statics with Respect to q) *The reporter's incentive to misreport (i.e., the expected gain from misreporting with respect to truth-telling) is strictly increasing in q for any $f \in [0, 1]$. Furthermore, the reporter strictly prefers to misreport rather than reporting truthfully if and only if $q > (4 - f) \,/\, [4(2 - f)]$.*

To illustrate, suppose $f = 1$. Figure 2a determines the reporter's best reply for the values of q used in our experiment. Represented functions are the reporter's expected payoff after observing $S = O$ and choosing to truthfully report $R = O$ (downward sloping segment) or to misreport (upward sloping segment), as a function of the probability that the state is b, q_O. When the prior is mildly unbalanced ($q = 6/10$), the posterior belief that the state is b is $q_O = 1/3 < 1/2$ and the reporter's gain from misreporting after seeing signal O is $\Delta_{EU}\,(6/10, 1) = -2/9$ (the unconditional gain, before observing $S = O$, is $-1/10$). When the prior is strongly unbalanced ($q = 8/10$), the reporter's posterior belief is $q_O = 4/7 > 1/2$, the gain from misreporting is $\Delta_{EU}\,(8/10, 1) = 2/21$ (unconditionally, before observing $S = O$, it is $1/30$), and, thus, the reporter prefers to misreport.

Comparative Statics with Respect to f. Next, turn to the effect of the evaluator's belief f that the reporter is truth-telling on the reporter's expected gain from misreporting. In general, the effect is ambiguous. While the gain in state o decreases with f, the gain in state b increases with f. Depending on q and on the evaluator's initial belief, increased trust (an increase in f) may either reduce or enhance the reporter's incentives to misreport. Nonetheless, for both values of q used in the experiment, incentives to misreport are lowest when the evaluator is strictly trusting. More generally, we have:

Proposition 3 (Comparative Statics with Respect to f) *If $q \in [1/2, 9/10]$, the reporter's incentive to misreport (i.e., the expected gain from misreporting with respect to truth-telling) is lowest when $f = 1$.*

The evaluator's belief about the reporter's truthfulness affects the expected assessment from misreporting only through p_{Bb} and p_{Bo}. When f decreases, by Proposition 1, both assessments p_{Bb} and p_{Bo} are dampened toward the prior, given that the signal is now less informative. A smaller f reduces the assessment when the core is blue and the report is accurate (p_{Bb}), but improves the assessment when the core is orange and the report is inaccurate (p_{Bo}). When the prior q is sufficiently low (and thus q_O is also low), the expected payoff of misreporting, $EU^M = (1 - q_O)\, p_{Bo} + q_O p_{Bb}$, assigns more weight to p_{Bo}, given that observing an orange core is more likely. Since p_{Bo} is decreasing in f, for low values of q, EU^M is decreasing in f, so that the incentives to misreport are minimized at $f = 1$. After a certain threshold of q, the expected gains from misreporting become increasing in f and are minimized at $f = 0$.[12]

Figure 2b illustrates changing reporter gains from misreporting for different evaluator beliefs, f, and a mildly unbalanced prior, $q = 6/10$. The continuous upward sloping line represents EU^M when $f = 1$. The expected payoff from misreporting becomes flatter as evaluator beliefs become less trusting, since p_{Bb} and p_{Bo} are dampened toward the prior 1/2. The dotted line represents EU^M when $f = 4/5$, and the flat dashed line represents EU^M when $f = 0$ (in this case evaluator assessments always equal the prior). For $q = 6/10$, the gain from misreporting increases as f decreases. The reporter is better off truthfully reporting when $f = 1$ (or $f = 4/5$), but prefers misreporting when $f = 0$.

3.3 Equilibrium Analysis

The analysis of the reporter's best reply hints at the structure of the equilibria. If the common prior belief about the state of the world is such that the reporter's best reply to an

[12] We give exact values for all mentioned thresholds in the proof provided in Supplementary Appendix A.

evaluator with perfectly trusting beliefs is to report truthfully, then such behavior can be sustained in equilibrium. Otherwise, only misreporting can be sustained in equilibrium.[13]

Proposition 4 (Equilibria) *When the prior belief about the state is mildly unbalanced,* $q \in [1/2, 3/4]$*, there are three Bayesian Nash equilibria: (i) a separating equilibrium in which the reporter reports truthfully, (ii) a pooling equilibrium in which the reporter misreports, and (iii) a hybrid mixed-strategy equilibrium (MSE) in which the reporter reports truthfully with probability:*

$$f^*(q) = \frac{8q-4}{4q-1} \tag{3}$$

and misreports with complementary probability, $1 - f^*$*. When the prior belief is strongly unbalanced,* $q \in [3/4, 1]$ *there is only a pooling equilibrium in which the reporter misreports.*

Figure 3 illustrates the intuition behind Proposition 4. For mildly unbalanced prior probabilities of the state, $q \in [1/2, 3/4]$, the thick dashed line corresponds to the inverse of equation 3, and thus represents the hybrid mixed strategy equilibrium of the game: the belief, f, that for a given prior, q, makes the reporter exactly indifferent between misreporting and truth-telling. The arrows indicate that this MSE is unstable. If the evaluator's belief about the reporter's truthfulness is slightly larger than $f^*(q)$ the reporter is better off reporting truthfully. Thus, points to the right and below the dashed curve constitute the basin of attraction of the separating equilibrium. If, instead, the evaluator's belief about the reporter's truthfulness is smaller than $f^*(q)$, the reporter is better off misreporting. Thus, points to the left and above the dashed curve constitute the basin of attraction of the pooling equilibrium. For strongly unbalanced priors, $q \in [3/4, 1]$, only the pooling equilibrium exists and all evaluator beliefs lie in the basin of attraction of this equilibrium.

Figure 3 also shows that as f decreases, the set of priors for which the reporter prefers to misreport increases. Intuitively this happens because, when the evaluator expects the reporter to be more likely to misreport, report $R = B$ contains less information about S, so that the evaluator becomes less able to make inference about u, the reporter's type. The assessment after report $R = B$ is thus dampened towards $1/2$, the prior probability that the signal is informative. This in turn reduces the potential loss from misreporting, which accrues if the state is o. Thus, misreporting may prove profitable even when state o has a high probability (low values of q).

[13]See also Ottaviani and Sørensen (2001) Lemma 1. We use Harsanyi's Bayesian Nash equilibrium notion since the choices made by the reporter and the evaluator are strategically simultaneous. The reason is that the evaluator observes only the report but not the reporter's strategy, even though the reporter's choice of a report precedes the evaluator's choice of an assessment.

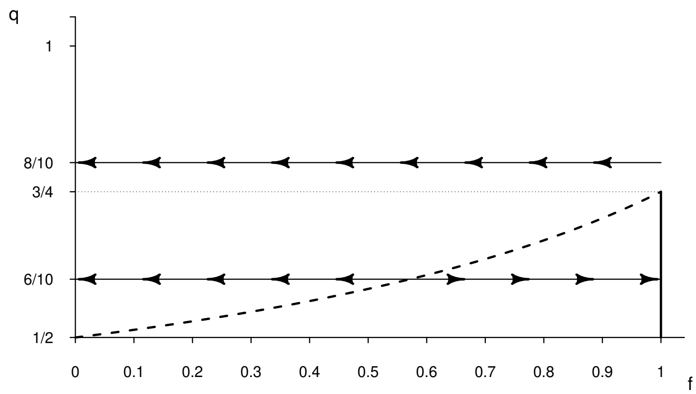

Figure 3: Equilibrium Set and Basin of Attraction.

3.4 Testable Hypotheses

The theory delivers the following testable hypotheses:

Reporters' Behavior

HP1 *Reporters are more likely to misreport when there is less uncertainty about the state of the world (that is, when q is larger).*

Proposition 2 predicts that the incentives to misreport increase in q. As the fraction of balls with a blue core increases, the fraction of balls with a Blue shell increases in the informative urn but remains unchanged in the uninformative urn. Hence, an Orange shell becomes a stronger indication that the urn is uninformative, giving the reporter a stronger incentive to misreport. Furthermore, as discussed in Proposition 4, when $q = 8/10$ only the pooling equilibrium (where the reporter misreports) exists, while $q = 6/10$ admits multiple equilibria.

HP2 *Reporters are least likely to misreport when the evaluator believes all reports are truthful, $f = 1$.*

The effects of a change in f on the incentives to misreport are ambiguous. On the one hand, as f increases, the payoff from misreporting when the core is Blue increases, because a Blue report is a stronger indication that the reporter observed a Blue shell and, thus, that the urn was informative. On the other hand, as f increases, the payoff from misreporting when the core is Orange decreases, because a stronger belief that the shell is truly Blue more strongly indicates that the urn was uninformative. According to Proposition 3, nonetheless, for the values of q used in our experiment, the loss from misreporting increases faster in f than the gain, so that the net gain reaches a minimum at $f = 1$.

Evaluators' Behavior

HP3 *Assessments are ranked: evaluators believe reporters are more likely to be informed after observing any accurate report than after observing any inaccurate report. If the received report is one that can only be made by a truthful reporter (that is, $R = O$), an accurate report leads to the highest belief that the reporter is informed, and an inaccurate report leads to the lowest belief that the reporter is informed (that is, $p_{Oo} \geq p_{Bb} \geq p_{Bo} \geq p_{Ob}$).*
Matching shell and core is the strongest indication that the urn is informative. We have $p_{Oo} \geq p_{Bb}$ because a matching orange report and core can only result if the shell is orange; instead, matching blue report and core may result from a Blue shell or from a misreported Orange shell. Next, $p_{Bb} \geq p_{Bo}$ because the evaluator is more confident that the reporter is informed when the report matches the state. Finally, p_{Ob} is the lowest assessment because an Orange report perfectly reveals an Orange shell, which combined with a blue core results in $p_{Ob} = 0$.

HP4 *Assessments after reports that can only be made by a truthful reporter (that is, $R = O$), do not depend on evaluators' beliefs about the reporter's strategy (that is, f). Assessments after reports that can be made by both truthful and misreporting reporters (that is, $R = B$), are more sensitive to the report's accuracy when evaluators believe that reporters are more likely to report truthfully (that is, are further from the 50% prior when f is larger).*
An Orange report perfectly reveals that the shell was Orange, resulting in an assessment independent of f. When the evaluator expects the reporter to be less likely to misreport (larger f), a Blue report contains more information about the observed shell, so that the evaluator makes a better inference about informativeness u. According to Proposition 1, the assessment after a Blue report then moves further away from 1/2, the prior probability that the signal is informative.

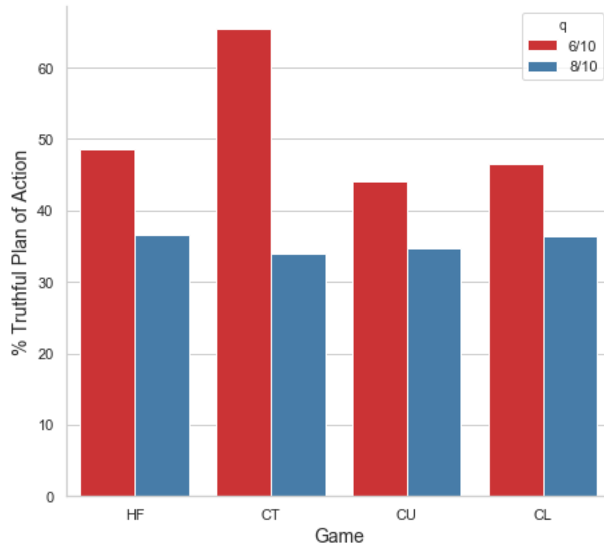

Figure 4: Reporters' Behavior by Game and Prior Beliefs on State (q). Experienced Subjects.

4 Experimental Results

4.1 Reporters' Behavior

Our experimental treatments are explicitly designed to investigate the effect of both q, the common belief about the state of the world, and f, the evaluator's belief about the reporter's strategy, on the reporters' incentives to misreport. We can thus use the behavior observed in our experimental games to test hypotheses HP1 and HP2 from Section 3.4. Throughout the Results section, we focus on experienced subjects, that is, on decisions belonging to the second block for each treatment.[14] Figure 4 shows the fraction of truthful plans of action by treatment. In all treatment games, the incidence of truth-telling is larger with $q = 6/10$ than with $q = 8/10$. The treatment with the largest incidence of truth-telling is game CT with $q = 6/10$ (65%) while the treatment with the least incidence of truth-telling is game CT with $q = 8/10$ (34%). In the rest of this section, we investigate whether these treatment effects are statistically significant. Whenever we state a result is significant, unless otherwise indicated, we refer to significance at the 1% level.

[14]Our findings are unchanged if we use all decisions. We present summary statistics for first-block decisions and discuss the effect of experience in Online Appendix B.

Pr[Reporter Chooses Truthful Plan of Action]				
	(1)	(2)	(3)	(4)
$q = 8/10$	-0.12	-0.31	-0.09	-0.10
	(0.02)	(0.02)	(0.02)	(0.02)
Constant	0.49	0.65	0.44	0.47
	(0.05)	(0.03)	(0.04)	(0.04)
Game	HF	CT	CU	CL
N	1472	1504	1504	1536

Table 2: The Effect of Prior Belief on Reporter's Behavior. Random effects GLS regressions. Experienced subjects. Each subject is a panel and periods are times within a panel. Standard errors in parentheses (clustered at the session level in columns 1 and 4). Results are robust to using random effects Probit regressions or OLS/Probit regressions with subject fixed effects.

4.1.1 Effect of Prior Beliefs on State (q)

Table 2 shows estimates of the effect of holding strongly unbalanced ($q = 8/10$) rather than weakly unbalanced priors ($q = 6/10$) about the state of the world on the probability reporters choose the truthful plan of action. We use random effects panel regressions to account for the fact that each individual makes multiple decisions in a single game. Moreover, we cluster standard errors at the session level to account for potential interdependencies between observations that come from random re-matching of subjects between periods in a session (when the reporters are matched with human evaluators, as in HF, or with computerized evaluators with strategies evolving over time, as in CL). All coefficients reported in Table 2 have the hypothesized sign and are significant at the 1% level.

FINDING 1: In all games, reporters are more likely to misreport with strongly unbalanced priors ($q = 8/10$) than with mildly unbalanced priors ($q = 6/10$). This provides evidence in favor of HP1.

4.1.2 Effect of Evaluators' Beliefs on Reporters' Strategy (f)

In games CU and CT, we exogenously manipulate evaluators' beliefs about reporters' strategies. Knowing that truth-telling incentives are maximal in treatment CT (HP2) gives a benchmark for comparison also for treatments CL and HF, where we do not control beliefs. We can thus test how reporters respond to the change in incentives due to a shock to their opponents' strategies by comparing reporter behavior across games. Table 3 shows estimates of the effect of the game on the probability the reporter chooses the truthful plan of action,

	Pr[Reporter Chooses Truthful Plan of Action]							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CT		0.21	0.19	0.17		-0.01	-0.02	-0.03
		(0.06)	(0.06)	(0.07)		(0.06)	(0.06)	(0.07)
CU	-0.21		-0.02	-0.04	0.01		-0.02	-0.02
	(0.06)		(0.06)	(0.07)	(0.06)		(0.06)	(0.07)
CL	-0.19	0.02		-0.02	0.02	0.02		-0.01
	(0.06)	(0.06)		(0.06)	(0.06)	(0.06)		(0.06)
HF	-0.17	0.04	0.02		0.03	0.02	0.00	
	(0.07)	(0.07)	(0.06)		(0.07)	(0.07)	(0.06)	
Constant	0.65	0.44	0.47	0.49	0.34	0.35	0.36	0.37
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Baseline	CT	CU	CL	HF	CT	CU	CL	HF
q	0.6	0.6	0.6	0.6	0.8	0.8	0.8	0.8
N	3008	3008	3008	3008	3008	3008	3008	3008

Table 3: The Effect of Evaluators' Beliefs on Reporter's Behavior. Random effects GLS regressions. Experienced subjects. Each subject is a panel and periods are times within a panel. Standard errors in parentheses.

keeping prior beliefs about the state of the world constant. Columns (1)–(4) report estimates for $q = 6/10$, while columns (5)–(8) report estimates for $q = 8/10$.

FINDING 2: With mildly unbalanced priors ($q = 6/10$), reporters are less likely to misreport when evaluators believe they report truthfully ($f = 1$, i.e., in game CT). This provides evidence in favor of HP2.

Column (1) in Table 3 shows that, with $q = 6/10$, there is significantly less misreporting in CT than in HF, CL, and CU. This provides evidence in favor of HP2. Column (5) in the same table shows that, with $q = 8/10$, there is no evidence of differential behavior between CT and the other games. This is not in line with predictions but the theory does predict a smaller effect for $q = 8/10$. For example, consider CU and CT—the two treatments where we exogenously manipulate f and the comparison is, thus, sharper. The difference in the expected gain from misreporting in a single period of the two games is €0.21 with $q = 6/10$ but only €0.10 with $q = 8/10$. The other columns in Table 3 show that we do not find any other significant difference between reporters' behavior in any other pair of games for any q.

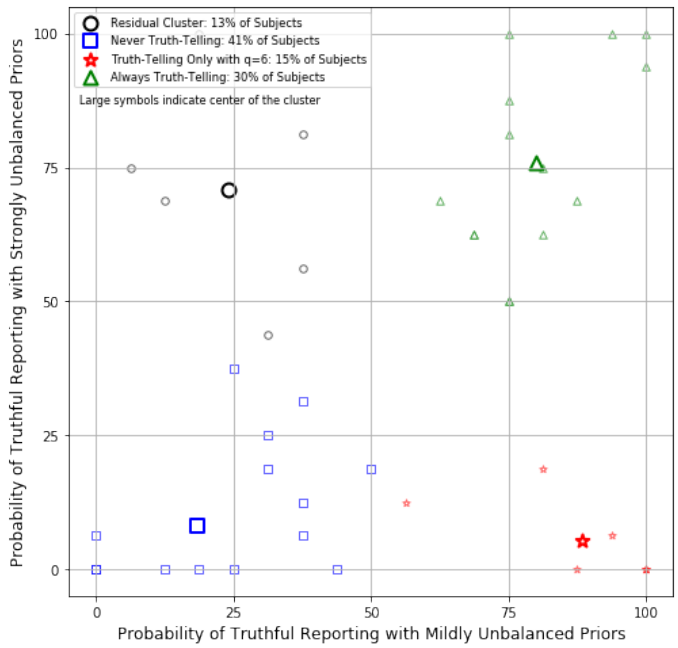

Figure 5: Reporters' Strategies Grouped in Clusters. Game HF. Experienced Subjects.

4.1.3 Behavioral Heterogeneity

The experimental results discussed above show that reporters, on average, adopt strategies which are consistent with the logic of our reputational cheap talk model. At the same time, aggregate data might hide heterogeneity in individual behavior. To investigate this possibility, we focus on game HF and use *k-means clustering* analysis of reporters' strategies.[15]

We use a subject as one observation and we define each subject's strategy as a two-dimensional vector including the probability of truthful reporting when $q = 6/10$ and the probability of truthful reporting when $q = 8/10$. The results indicate that about 90% of subjects can be classified into three clusters, whose representative strategies are displayed in Figure 5.[16] Two clusters, accounting for 56% of subjects, exhibit behavior which is consistent with equilibrium predictions: subjects belonging to one cluster, which accounts for 41% of observations, misreport most of the time regardless of the belief on the state of the world;

[15] K-means clustering is a common unsupervised learning methodology used to group observations according to their similarity in a multidimensional space of observable characteristics (see MacQueen 1967, Hartigan 1975, Hastie et al. 2005, and Murphy 2012; for a recent use in experimental economics, see Fréchette et al. 2019). The procedure randomly selects k points in the space of observable characteristics to be the centers of k clusters; each observation is then associated with the closest center, and centers are iterated on to minimize the total within cluster variance. This procedure is repeated 10 times with 10 different random cluster centers; if the final clusters are different, the algorithm selects the best result.

[16] K-means clustering requires the choice of the number of clusters at the outset. As customary, we determined the number of clusters for evaluators and reporters with the elbow method. The figure also displays a "residual cluster" that gathers all the remaining, harder-to-categorize, subjects.

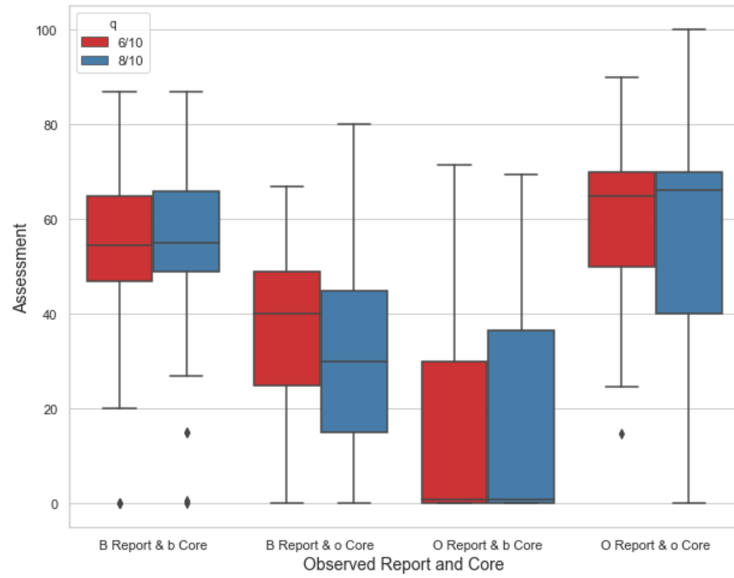

Figure 6: Human Evaluators' Assessments. Boxplots by Observed Report-Core Pair and Prior Beliefs on State (q). Game HF. Experienced Subjects.

subjects belonging to a second cluster, which accounts for 15% of observations, misreport most of the time when they hold strongly unbalanced priors on the state of the world and report truthfully most of the time when they hold mildly unbalanced priors. Another cluster, which includes 30% of subjects, displays truthful reporting regardless of the priors on the state of the world, that is, including when priors are strongly unbalanced and the unique equilibrium prescribes misreporting. While inconsistent with equilibrium predictions, the behavior of these subjects can be accounted for by lying aversion (Gneezy et al., 2018) or preferences for truth-telling (Abeler et al., 2019) and contributes to explain the large incidence of truth-telling we observe in the aggregate data.

4.2 Human Evaluators' Behavior

We study evaluators' behavior in game HF, the only one with human evaluators. Given the incentives of human evaluators, we expect them to truthfully reveal their best assessment about the probability that the drawn ball in any period came from the informative urn. With this in mind, evaluators' assessments should be affected only by three variables: the received report, the observed color of the core of the ball, and the belief about reporters' truthfulness, f. While we do not exogenously set any of these variables, we can exploit the variation in the realization of reports and states of the world, as well as the indirect effect of the common prior, q, on evaluators' beliefs, f, for our hypotheses tests. We thus use

	N	Average	Median	Theory
Blue Report, Blue Core	383	54.4	54.5	[50, 66.7]
Blue Report, Orange Core	194	35.4	40.0	[0, 50]
Orange Report, Blue Core	54	15.2	0.1	0
Orange Report, Orange Core	105	58.6	65.0	66.7

Table 4: Human Evaluators' Assessments, $q = 6/10$. Game HF, experienced subjects.

	N	Average	Median	Theory
Blue Report, Blue Core	522	56.4	55.0	[50, 66.7]
Blue Report, Orange Core	117	30.2	30.0	[0, 50]
Orange Report, Blue Core	59	16.8	1.0	0
Orange Report, Orange Core	38	58.7	66.2	66.7

Table 5: Human Evaluators' Assessments, $q = 8/10$. Game HF, experienced subjects.

behavior observed in game HF to test hypotheses HP3 and HP4 from Section 3.4

4.2.1 Effect of Observing Different Reports and Cores

Figure 6, Table 4 and Table 5 show summary statistics for human evaluators' assessments of the probability the urn is informative. Each boxplot (in the figure) and each row (in the tables) are for a different pair of observed report and observed core. Table 4 focuses on the treatment with mildly unbalanced prior ($q = 6/10$), while Table 5 focuses on the treatment with strongly unbalanced prior ($q = 8/10$).

FINDING 3: Evaluators' assessments are strictly ranked: $p_{Oo} > p_{Bb} > p_{Bo} > p_{Ob}$**. This provides evidence in support of HP3. Moreover, following an Orange report, assessments in the experiment are indistinguishable from assessments by a Bayesian evaluator,** $p_{Oo} = 2/3$ **and** $p_{Ob} = 0$**.**

We compare the whole distribution of assessments following each report and core pair with Kolmogorov-Smirnov and Wilcoxon-Mann-Whitney tests. For both $q = 6/10$ and $q = 8/10$, evaluators are most confident that the signal is informative (that is, that the reporter is well informed) when they observe an orange report and an orange core and least confident when they observe an orange report and a blue core. Observing a blue report and a blue core makes evaluators more confident than observing any inaccurate report and observing a blue report and an orange core makes them more confident than observing an orange report and a blue core and less confident than observing any accurate report. All differences are statistically significant at the 1% level.[17] This evidence supports hypothesis HP3.

The median assessment following an orange report and an orange core (65.0 for $q = 6/10$ and 66.2 for $q = 8/10$) and the median assessment following an orange report and a blue core (0.1 with $q = 6/10$ and 1.0 with $q = 8/10$) are indistinguishable from the assessments made by a Bayesian evaluator (respectively, 66.6 and 0). Note that the assessments of a Bayesian evaluator following an orange report do not depend on beliefs about the reporters' strategy. On the other hand, the assessments of a Bayesian evaluator following a blue report do depend on these beliefs. The average and median assessments after a blue report given by our human evaluators are consistent with some belief $f \in [0, 1]$. We explore this in further detail in Sections 4.2.2 and 4.2.3 below.

4.2.2 Effect of Prior Beliefs on the State (q)

Table 6 shows estimates of the effect of holding strongly unbalanced ($q = 8/10$) rather than weakly unbalanced priors ($q = 6/10$) about the state of the world on human evaluators' assessments of the probability the urn is informative. Each column focuses on a different report and core pair. From the perspective of evaluators who are trying to assess the informativeness of the urn, the only possible difference between the two treatments lies in the strategy adopted by reporters: both theoretically and empirically, reporters are less likely to report their signal truthfully with $q = 8/10$ than with $q = 6/10$. A Bayesian evaluator who is aware of this differential behavior should give assessments that are further from the 50% prior (in the sense of rewarding to a larger extent a report matching the state and punishing to a larger extent a report mismatching the state) with $q = 6/10$ than with $q = 8/10$. At the same time, any difference in evaluators' beliefs about the strategy adopted by reporters with $q = 8/10$ and with $q = 6/10$ should not affect assessments following an orange report.

FINDING 4: Evaluators' assessments after orange reports do not depend on q. This is in line with HP4. Assessments after blue reports are further from the prior with $q = 8/10$ than with $q = 6/10$. This is in line with HP4 only if evaluators (incorrectly) believe reporters are more likely to misreport with $q = 6/10$.

[17] The only exception is the difference between assessments following a blue report and a blue core and assessments following an orange report and an orange core with $q = 8/10$: the p-value of the Kolmogorov-Smirnov test is 0.001 but the p-value of the Wilcoxon-Mann-Whitney test is 0.918.

	(1)	(2)	(3)	(4)
	Dependent Variable: Evaluator's Assessment			
$q = 8/10$	1.87	-4.34	1.52	-0.62
	(0.65)	(1.42)	(1.66)	(2.28)
Constant	54.99	34.46	15.53	59.20
	(1.69)	(1.89)	(3.13)	(2.40)
Report	Blue	Blue	Orange	Orange
Core	Blue	Orange	Blue	Orange
N	905	311	113	143

Table 6: Human Evaluators' Assessments as a Function of q. Random effects GLS regressions. Experienced subjects. Each subject is a panel and periods are times within a panel. Standard errors clustered at the session level in parentheses.

Columns (3) and (4) in Table 6 show that evaluators' assessments following an orange report are not significantly affected by the prior belief about the state of the world. On the other hand, the assessed likelihood that the urn is informative is significantly larger with $q = 8/10$ after seeing an accurate blue report and significantly lower with $q = 8/10$ after seeing an inaccurate blue report. This suggests that human evaluators are more sensitive to information with $q = 8/10$ than with $q = 6/10$. This can be rationalized by a belief reporters are more likely to report their signal truthfully with $q = 8/10$ than with $q = 6/10$. Indeed, this is confirmed by the structural estimation of human evaluators' beliefs about reporters' strategies reported in Table 7. For each human evaluator, the estimated f is found as the minimizer of the sum of the squared distances from the Bayesian posteriors for all assessments this subject makes following a blue report in a given treatment. Each evaluator makes 16 assessments in each treatment (we consider only experienced evaluators—second block of 16 periods). The median estimated probability that reporters are truthful is 50% with $q = 6/10$ and 65% with $q = 8/10$. Evaluators' estimated beliefs are positively and significantly affected by the treatment according to a Tobit (p-value 0.004) or linear (p-value 0.037) regression with subjects fixed effects. As discussed in Section 4.1.1 and Finding 1, this perception is not in line with reporters' actual behavior in game HF. In fact, consider a Bayesian evaluator who best responds to reporters' observed behavior, that is, an evaluator who forms assessments using the empirical values of f. In the treatment with $q = 6/10$, this evaluator would estimate a 57% chance the signal is informative when observing a blue report and a blue core; and a 40% chance the signal is informative when observing a blue report and an orange core. In the treatment with $q = 8/10$, this evaluator would estimate

	N	Average	1st Quartile	2nd Quartile	3rd Quartile	Theory
$q = 6/10$	47	0.43	0	0.50	0.80	[0,1]
$q = 8/10$	47	0.51	0.15	0.65	0.85	0

Table 7: Human Evaluators' Estimated f. Experienced subjects. Game HF.

a 55% chance the signal is informative when observing a blue report and a blue core; and a 44% chance the signal is informative when observing a blue report and an orange core. The observed assessments are close to the best responses with one exception: assessments following an inaccurate blue report with $q = 8/10$ are much more pessimistic than they should (the average and median assessment being 30%).

4.2.3 Explanation for Discrepancy: Learning Model of Evaluators' Behavior

To shed light on the origins of this discrepancy, we investigate whether human evaluators' assessments are consistent with the learning model we endowed computerized evaluators with in game CL. As described in more detail in the Appendix, this learning model has the following elements:

- The "true" fraction f of reporters that report truthfully is unknown but constant.
- The starting point of beliefs is the uniform distribution, $f \sim U(0, 1)$.
- Evaluators learn from each interaction with a reporter. At the end of a period, the belief is updated based on the information available to human evaluators in game HF:
  - the received report, $R \in \{B, O\}$,
  - the observed core, $c \in \{b, o\}$,
  - the true informativeness of the signal, $u \in \{I, U\}$.

This information is easily summarized as a triple, (R, c, u), taking on eight distinct values. One of these triples, (O, b, I), cannot occur, while the remaining seven can be grouped into four events considered to be positive, negative, neutral, or muddy signals about the reporters' truthfulness. A negative signal tells the evaluator that the reporter he met this period misreported for sure; a positive signal tells him that the reporter was truthful for sure; a neutral signal bears no information; a muddy signal assigns positive but distinct likelihood to

Event	Triples – (R, c, u)	Likelihood
Positive	(O, o, I) (O, o, U) (O, b, U)	$\frac{3-2q}{4} f$
Negative	(B, o, I)	$\frac{1-q}{2} (1 - f)$
Muddy	(B, o, U) (B, b, U)	$\frac{1}{4} (2 - f)$
Neutral	(B, b, I)	$\frac{q}{2}$

Table 8: Events and likelihoods used for evaluators' updating in the Bayesian learning model.

both truth-telling and misreporting. Table 8 summarizes these events and their likelihoods given a prior belief, f, about the reporters' truthfulness. Proposition 5 in the Appendix characterizes the distribution of an evaluator's beliefs over f as a function of the number of negative, positive and muddy signals observed thus far. In the same Appendix, Proposition 6 characterizes the expected assessments given by an evaluator who receives a Blue report and observes either an orange or a blue core as a function of his experience.

Table 9 reports summary statistics for assessments given by computerized evaluators who update their beliefs on f as described above and have the same experience as human evaluators in game HF. Table 10 reports estimates on the effect of holding strongly unbalanced rather than weakly unbalanced priors about the state of the world on these computerized evaluators' assessments of the probability the urn is informative. Assessments following an accurate blue report are significantly less generous and assessments following an inaccurate blue report are significantly less punitive with $q = 8/10$ than with $q = 6/10$. This is in line with computerized evaluators believing (correctly) that human reporters are more likely to misreport with $q = 8/10$ than with $q = 6/10$.

FINDING 5: The behavior of human evaluators is not consistent with a learning model that posits they initially believe $f \sim U(0, 1)$ and update beliefs according to Bayes' rule based on the evolution of the interaction with reporters.

Another possibility is that evaluators have limited attention and do not use all the information available to them but focus on a salient piece of information, that is, whether reports were accurate or inaccurate (that is, on whether they matched or not the core). If this is the case, evaluators might naïvely infer reporters' strategies from report accuracy. In particular, we modify our learning model and assume that evaluators take any accurate report as a positive signal (that is, a signal that the reporter he met this period was truthful for sure) and any inaccurate report as a negative signal (that is a signal that the reporter he met this

	N	Average	Median	Theory
$q = 6/10$				
Blue Report, Blue Core	383	57.3	57.2	[50, 66.7]
Blue Report, Orange Core	194	37.2	38.2	[0, 50]
$q = 8/10$				
Blue Report, Blue Core	522	56.1	55.6	[50, 66.7]
Blue Report, Orange Core	117	38.0	39.1	[0, 50]

Table 9: Hypothetical assessments made by computerized evaluators who initially believe $f \sim U(0, 1)$, learn in a Bayesian way depending on interactions, and have the same experience as humans in game HF.

	(1)	(2)
Dependent Variable: Computer Assessment		
$q = 8/10$	-1.34	3.30
	(0.13)	(0.57)
Constant	57.39	37.21
	(0.26)	(0.61)
Report	Blue	Blue
Core	Blue	Orange
N	905	311

Table 10: Random effects GLS regressions. Experienced 'Bayesian' computers. Each computer is a panel and periods are times within a panel. Standard error in parentheses.

period misreported for sure).

Table 11 reports summary statistics for assessments given by 'naïve' or 'behavioral' computerized evaluators who update their beliefs on f as described above and have the same experience as human evaluators in game HF. Table 12 reports estimates on the effect of holding strongly unbalanced rather than weakly unbalanced priors about the state of the world on these computerized evaluators' assessments of the probability the urn is informative. As is the case for human evaluators, assessments following an accurate blue report are significantly more generous and assessments following an inaccurate blue report are significantly more punitive with $q = 8/10$ than with $q = 6/10$. This is in line with computerized evaluators believing (incorrectly, and similarly to human evaluators) that human reporters

	N	Average	Median	Theory
$q = 6/10$				
Blue Report, Blue Core	383	59.8	59.6	[50, 66.7]
Blue Report, Orange Core	194	33.2	33.9	[0, 50]
$q = 8/10$				
Blue Report, Blue Core	522	61.6	61.6	[50, 66.7]
Blue Report, Orange Core	117	27.7	29.3	[0, 50]

Table 11: Hypothetical assessments made by computerized evaluators who initially believe $f \sim U(0,1)$, learn 'behaviorally' depending on interactions, and have the same experience as humans in game HF.

	(1)	(2)
Dependent	Variable:	Computer Assessment
$q = 8/10$	1.80	-5.59
	(0.09)	(0.54)
Constant	59.72	32.72
	(0.19)	(0.61)
Report	Blue	Blue
Core	Blue	Orange
N	905	311

Table 12: Random effects GLS regressions. Experienced 'behavioral' computers. Each computer is a panel and periods are times within a panel. Standard error in parentheses.

are less likely to misreport with $q = 8/10$ than with $q = 6/10$.

FINDING 6: The behavior of human evaluators is consistent with a learning model that posits they initially believe $f \sim U(0,1)$ and learn depending on interactions with reporters but mistakenly consider report accuracy (inaccuracy) as a perfect signal of truthful reporting (misreporting).

4.2.4 Behavioral Heterogeneity

As for reporters, aggregate data might hide heterogeneity in individual behavior. To investigate this possibility, we use k-means clustering analysis of human evaluators' strategies. We define each subject's strategy as a six-dimensional vector including the average assessment

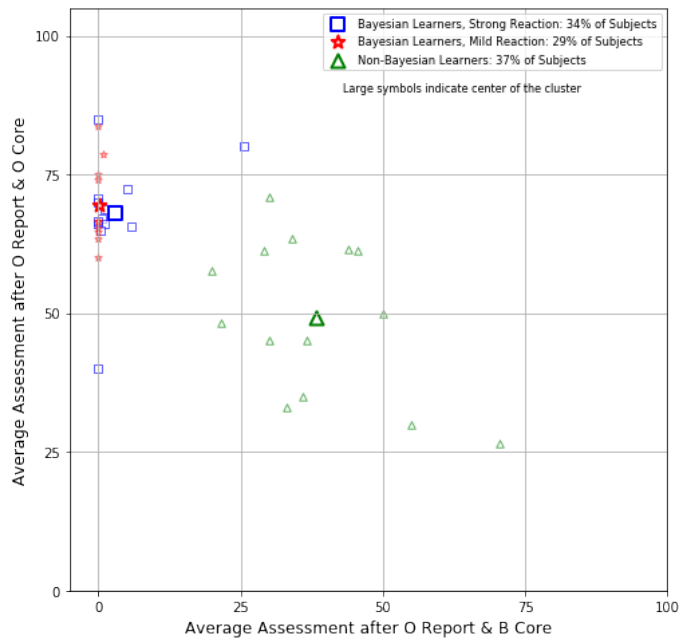

Figure 7: Human Evaluators' Strategies Grouped in Clusters, O Report

when the report is orange and the core is blue (unconditional on the belief on the state of the world, as q does not matter for inference in this case), the average assessment when the report is orange and the core is orange (again, unconditional on q), the average assessment when the report is blue and the core is blue for each q, and the average assessment when the report is blue and the core is orange for each q.[18] The results indicate that subjects can be classified into three clusters, whose representative strategies are displayed in Figure 7 (for assessments following an orange report), Figure 8 (for assessments following an inaccurate blue report), and Figure 9 (for assessments following an accurate blue report).

The first two clusters, which encompass 63% of subjects, have one feature in common: assessments following an orange report — which is indicative of a truthful plan of action — are consistent with those of a Bayesian learner. The modal strategy in these two clusters is to assess a 0% chance the urn is informative following an inaccurate orange report and a 67% chance the urn is informative following an accurate orange report. We label these subjects as Bayesian learners. The third cluster is, instead, composed of subjects who fail to make these basic inferences: their average assessments following an orange core are dispersed and, in general, do not reward (punish) reporters sufficiently in case of an accurate (inaccurate) report. We label these subjects as non-Bayesian learners. Interestingly, the two groups of

[18]We exclude from the analysis evaluators who do not make at least one assessment in each of these six instances. This leaves us with 38 out of 47 subjects.

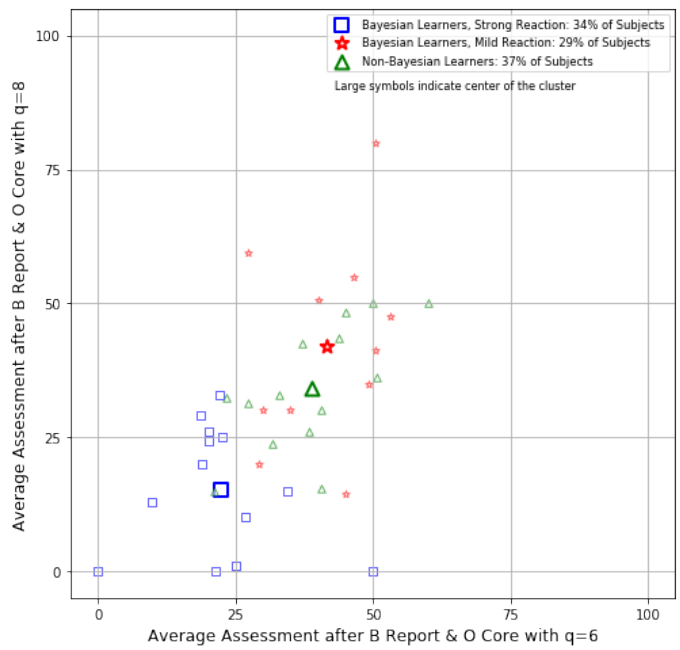

Figure 8: Human Evaluators' Strategies Grouped in Clusters, Inaccurate B Report

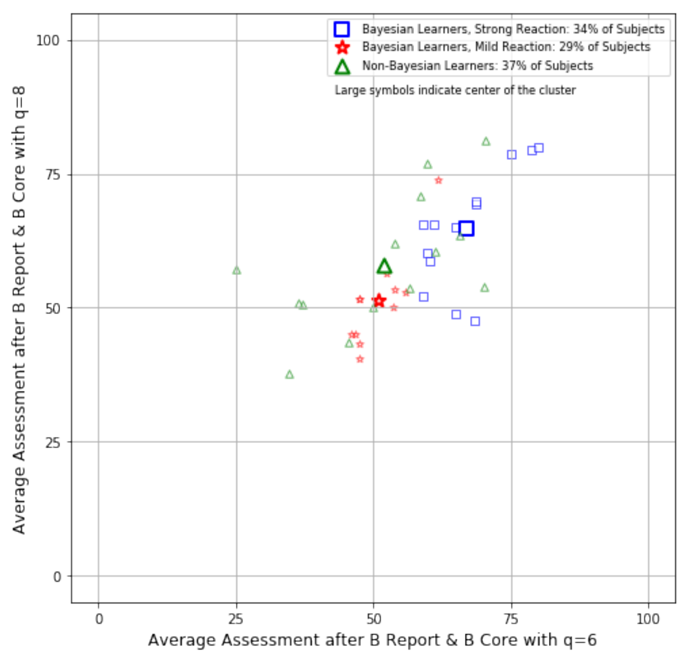

Figure 9: Human Evaluators' Strategies Grouped in Clusters, Accurate B Report

Bayesian learners differ in the strategies following a blue report. While their assessments remain consistent with Bayesian inference (under some belief on reporters' strategy), subjects belonging to the first cluster do not update their prior much in either direction and their assessments do not depend on the prior beliefs on the state of the world: the typical assessments are around 51 after an accurate blue report and around 42 with an inaccurate blue report. On the other hand, subjects belonging to the second cluster respond strongly to the observed accuracy of a blue report, especially in the sense of punishing an inaccurate report: the modal assessments when the report is accurate are 52 with $q = 6/10$ and 58 with $q = 8/10$; when the same report is inaccurate, the typical response is to assess a 39% chance of an informative urn with $q = 6/10$ and only a 34% chance with $q = 8/10$. Inconsistent with best reply to the empirical rate of truth-telling, these subjects respond more strongly to the accuracy of a blue report in the treatment with $q = 8/10$ than in the treatment with $q = 6/10$ and are, thus, responsible for the pattern we documented in the previous section.

5 Conclusion

This paper presents a laboratory experiment designed to test a widely applied model of reputational cheap talk where a reporter wants to convince an evaluator of being well informed.

A first innovation in the experimental design consists in the introduction of nested balls, where the color of the inner core corresponds to the realization of the state while the color of the outer shell corresponds to the noisy signal. A second, innovation consists in controlling for strategic behavior and learning on the side of the evaluator in a number of intermediate experiments in which we computerize evaluators by programming them to best reply to expectations about the reporter's behavior. In treatments CT and CU these beliefs are fixed throughout the experiment, while, in treatment CL, computerized evaluators' expectations are updated according to a generalized Beta learning model that we characterize in the Appendix. We analyze the outcomes of these experiments and use them as baselines to study play in the full game where also evaluators are human subjects with unrestricted beliefs.

Within the context and the sample of our experiment, we find that reporters realize how their strategic incentives are affected by the uncertainty about the phenomenon they are asked to forecast (q) and by the evaluators' expectations (f) and learn to best reply even when confronted with the noisy behavior of human evaluators. On the other side of the game, human evaluators find it difficult to assess the quality of the information available to reporters and to learn how the strategies used by reporters change in different environments. The noisier evaluation that results from human evaluators, in turn, reduces the reporters' incentives for misreporting and ameliorates information transmission. Overall, our experiment suggests that current models of reputational cheap talk are accurate in modeling reporters' behavior but might be missing important elements in the way evaluators process the available information or reward reporters for their advice.

Our experiment has also broader implications for sender-receiver games. In a comprehensive review, Blume et al. (Forthcoming) note that, "from its inception, the experimental literature on strategic information transmission with a single sender and a single receiver documents systematic over-communication, relative to the most informative equilibria [...]. Some of this over-communication can be accounted for by bounded-rationality approaches, like learning and level-k reasoning." To relate our experiment to this literature, we focus on the treatment with human senders, human receivers and a unique equilibrium in which senders misreport.[19] As in the experiments reviewed by Blume and coauthors, we find evidence of over-communication: truthful reporting is chosen around a third of the time, even among subjects who have gained a considerable amount of experience. Replacing human receivers with computerized receivers does not affect senders' behavior (Finding 2). This suggests that over-communication cannot be fully explained by lying costs: senders should be less averse to misreport the available information when misreporting does not affect other subjects' ability to form accurate assessments and accrue higher earnings.

More plausibly, over-communication, both in our game and in other sender-receiver games, could be due to the difficulty of both senders and receivers to learn the strategy of their opponents. Our experiment shows that this is a hurdle even in a setting where the space of messages is minimal and language is codified. Indeed, thanks to the tractable model of Bayesian learning we introduce, we show that receivers in our experiment do not learn senders' strategies as a statistician would do. In turn, this can impair senders' learning on receivers' strategies. To test this hypothesis directly, we committed ex-ante to a series of novel experimental treatments with computerized receivers whose beliefs are fixed or evolve in a Bayesian way with experience and who use these beliefs optimally to make decisions. Our attempt to manipulate receivers' strategies and their learning does not reduce the amount of over-communication we observe in the experiment, possibly because of bounded rationality on the side of senders and because the treatments we devised do not reduce sufficiently the complexity of the game. Future work should continue to investigate sender-receiver games along this direction, for example, by using computerized receivers in other games of strategic communication or developing treatments with computerized senders.

[19]In our experiment, the sender has two choices. As discussed in Section 3, with $q = 6/10$, both actions are part of an equilibrium so there is no behavior which can be classified as over- or under-communication. On the other hand, with $q = 8/10$, there is a unique equilibrium in which the sender misreports. In this case, choosing the truthful plan of action can be regarded as over-communication relative to the most informative equilibrium.

Appendix: Generalized Beta-Noisy Bernoulli Model

The classic Bernoulli-Beta model—generalized in this appendix—characterizes the evolution of the belief that a (possibly biased) coin, which, when tossed, gives Tails with frequency f and Heads with complementary frequency. Suppose that the prior of f is Beta distributed with parameters α and β. After observing a toss of the coin, the posterior of f is still Beta distributed, with parameters $\alpha' = \alpha + 1$ and $\beta' = \beta$ if Tails is observed and with parameters $\alpha' = \alpha$ and $\beta' = \beta + 1$ if Heads is observed. Thus, the Beta distribution is a conjugate prior with respect to Bernoulli-trial learning.

In the application we consider in this paper, the fraction of truthful reporters, f, is an unknown parameter. Each period of the experiment in which evaluators and reporters interact corresponds to a trial giving the evaluator the opportunity to learn about f. The information received by evaluators at the end of each period's trial, however, does not exactly correspond to the observation whether the reporter drawn from the population is truthful or not. The feedback evaluators receive consists, instead, of a triple, comprising the Report, the observed core, and the true informativeness of the urn from which the ball was drawn, (R, c, u).

As presented in Table 8, all feedback triples with $R = O$ perfectly reveal that the reporter is truthful, whereas the feedback triple involving an informative urn, an orange core, and a Blue report, (B, o, I), is a perfect signal of a misreporting reporter. These signals are, thus, equivalent to observing Heads or Tails in a Bernoulli trial. However, triples involving a Blue report and an uninformative urn, are imprecise: when the urn is uninformative, a Blue report may originate from a truthful or a misreporting reporter, regardless of the color of the core. The conditional probabilities of these realizations depend on the fraction of truthful reporters, so that these *muddy* signals contain some information.

Below we generalize the basic Beta learning model to allow for learning from Bernoulli trials with imprecise signals of trial outcomes. As we show, the generalized Beta distribution introduced by Exton (1976) is a conjugate prior with respect to noisy Bernoulli sampling (Proposition 5). We then compute the expected value of a class of functions of random variables with Exton generalized Beta distribution (Proposition 6). To the best of our knowledge, our characterization of this conjugate model is novel to the literature and constitutes an additional, free-standing, contribution of our paper. The paper put this model to work to program learning by computerized evaluators; this model can be applied more generally to model learning about a fixed frequency in other settings involving imperfect signals. This model also proves useful as a benchmark for our analysis of the learning behavior of human subjects.

Noisy Bernoulli Experiment. The underlying parameter is the unknown probability $\theta \in (0,1)$ that a Bernoulli trial gives $i = 1$. A noisy signal $j \in \{0, \ldots, J\}$ about the outcome of the Bernoulli trial $i \in \{0,1\}$ is observed, rather than the outcome of the underlying Bernoulli trial. Realization j of this noisy Bernoulli experiment has conditional probability $\pi_{j|1}$ when the outcome of the trial is $i = 1$, and $\pi_{j|0}$ when the outcome of the trial is $i = 0$, where $\pi_{0|1} = \pi_{J|0} = 0$ capture the possibility that signal realizations $j = 0$ and $j = J$ perfectly reveal the outcome of the trial.[20]

Noisy Bernoulli sampling consists of K independent repetitions of this noisy Bernoulli experiment, where σ_j denotes the number of times signal j is realized.

Exton Generalized Beta. The Exton Generalized Beta probability density function of a random variable $0 \leq x \leq 1$ is given by

$$g(x; v_1, v_2, d_1, \ldots, d_H, \delta_1, \ldots, \delta_H) = \frac{x^{v_1-1}(1-x)^{v_2-1}}{\mathrm{B}(v_1, v_2)} \times \frac{(1-\delta_1 x)^{d_1-1} \ldots (1-\delta_H x)^{d_H-1}}{F_D^{(H)}(v_1; 1-d_1, \ldots, 1-d_H; v_1+v_2; \delta_1, \ldots, \delta_H)}, \tag{4}$$

with $v_1, v_2 > 0$, where

$$F_D^{(H)}(a; \mu_1, \ldots, \mu_H; c; \gamma_1, \ldots, \gamma_H) = \sum_{i_1, \ldots, i_H = 0}^{\infty} \frac{(a)_{i_1+\ldots+i_H} (\mu_1)_{i_1} \cdots (\mu_H)_{i_H}}{(c)_{i_1+\ldots+i_H}} \frac{\gamma_1^{i_1}}{i_1!} \cdots \frac{\gamma_H^{i_H}}{i_H!}$$

is the fourth Lauricella function with $H \in \mathbb{N}$ and positive real parts of a and $c-a$ (v_1 and v_2 in the Exton generalized Beta), the notation $(\cdot)_h$ denotes the Pochhammer symbol, defining the function

$$(b)_h = \begin{cases} 1 & \text{if } h = 0 \\ b(b+1)\ldots(b+h-1) & \text{if } h > 0, \end{cases}$$

and

$$\mathrm{B}(v_1, v_2) = \frac{\Gamma(v_1)\Gamma(v_2)}{\Gamma(v_1+v_2)}$$

[20] If, in addition, $\pi_{0|0} = 0$ or $\pi_{J|1} = 0$ no such perfectly revealing outcomes are possible. Also, note that in the degenerate case with $\pi_{0|0} = \pi_{J|1} = 1$ we are back to the original Bernoulli trial. The special case with $J = 2$ corresponds to Warner's (1965) randomized response model. See also Winkler and Franklin (1979).

is the Beta function with $v_1, v_2 \in \mathbb{Z}_+$, with

$$\Gamma(\nu) = \int_0^{\infty} t^{\nu-1} e^{-t} \mathrm{d}t$$

denoting the Gamma function.

Proposition 5 (Conjugation) *The Exton Generalized Beta distribution is conjugate with respect to noisy Bernoulli sampling.*

Proof. With noisy Bernoulli sampling, the probability of observing signal j conditional on θ is $Pr\left(j|\,\theta\right) = \theta\pi_{j|1} + \left(1-\theta\right)\pi_{j|0}$. Knowing that $j=0$ and $j=J$ are precise signals of $i=0$ and $i=1$, respectively, we have

$$\begin{aligned}
\Pr(j = 0|\,\theta) &= (1-\theta)\,\pi_{0|0} \\
\Pr(j = 1|\,\theta) &= \theta\pi_{1|1} + (1-\theta)\,\pi_{1|0} \\
&\vdots \\
\Pr(j = J|\,\theta) &= \theta\pi_{J|1}
\end{aligned}$$

and, thus, the likelihood of the sample with signal frequencies $\sigma_0, \sigma_1, \ldots, \sigma_J$, is given by

$$l\left(\sigma_0, \ldots, \sigma_J|\,\theta\right) = \left[\pi_{0|0}^{\sigma_0} \times \ldots \pi_{J-1|0}^{\sigma_{J-1}} \times \pi_{J|1}^{\sigma_J}\right] \theta^{\sigma_J} \left(1-\theta\right)^{\sigma_0} \prod_{j=1}^{J-1} \left[1 - \left(1 - \frac{\pi_{j|1}}{\pi_{j|0}}\right)\theta\right]^{\sigma_j}. \qquad (5)$$

Notice that it is not necessary that there be precise signals ($j=0$ and $j=J$), since all ensuing steps will follow through if $\sigma_0 = \sigma_J = 0$ always.

Bayesian updating from a prior $g\left(\theta;\cdot\right)$, after observing sample $\sigma_0, \sigma_1, \ldots, \sigma_J$, yields posterior

$$g\left(\theta;\cdot|\,\sigma_0, \ldots, \sigma_J\right) = \frac{g\left(\theta;\cdot\right) l\left(\sigma_0, \ldots, \sigma_J|\,\theta\right)}{\int_t g\left(t;\cdot\right) l\left(\sigma_0, \ldots, \sigma_J|\,t\right) dt}.$$

Assume that θ has an Exton generalized Beta prior, $g\left(\theta; v_1, v_2, d_1, \ldots, d_H, \delta_1, \ldots, \delta_H\right)$, with $H \geq J$, and parameters $\delta_h = 1 - \frac{\pi_{h|1}}{\pi_{h|0}}$ for $h = 1, \ldots, J-1$.[21] The numerator of the above expression is given by

[21]This is without loss of generality, since it suffices to set $d_h = 1$ to obtain a prior with less than $J+1$ factors, or whose original parameters, δ_h, differ from $1 - \frac{\pi_{h|1}}{\pi_{h|0}}$ for all h. As is clear from the ensuing steps, factors are added via the likelihood, as signals are sampled.

$$
\begin{aligned}
& g\left(\theta; v_1, v_2, d_1, \ldots, d_H, \delta_1, \ldots, \delta_H\right) l\left(\sigma_0, \ldots, \sigma_J \middle| \theta\right) \\
= & \frac{\left[\theta^{v_1-1}\left(1-\theta\right)^{v_2-1} \prod_{h=1}^{H}\left(1-\delta_h\theta\right)^{d_h-1}\right]\left[\left(\pi_{0|0}^{\sigma_0}\ldots\pi_{J-1|0}^{\sigma_{J-1}}\pi_{J|1}^{\sigma_J}\right)\theta^{\sigma_J}\left(1-\theta\right)^{\sigma_0}\prod_{j=1}^{J-1}\left(1-\delta_j\theta\right)^{\sigma_j}\right]}{\mathrm{B}\left(v_1, v_2\right) F_D^{(H)}\left(v_1; 1-d_1, \ldots, 1-d_H; v_1+v_2; \delta_1, \ldots, \delta_H\right)} \\
= & \frac{\left(\pi_{0|0}^{\sigma_0}\ldots\pi_{J-1|0}^{\sigma_{J-1}}\pi_{J|1}^{\sigma_J}\right)}{\mathrm{B}\left(v_1, v_2\right) F_D^{(H)}\left(v_1; 1-d_1, \ldots, 1-d_H; v_1+v_2; \delta_1, \ldots, \delta_H\right)} \times \\
& \left[\theta^{v_1+\sigma_J-1}\left(1-\theta\right)^{v_2+\sigma_0-1}\prod_{j=1}^{J-1}\left(1-\delta_j\theta\right)^{d_j+\sigma_j-1}\prod_{h=J}^{H}\left(1-\delta_h\theta\right)^{d_h-1}\right]
\end{aligned}
$$

The denominator being the integral on θ of the above expression, all terms independent of θ (first factor in the final expression) exactly match in numerator and denominator and, thus, cancel out to give

$$
\begin{aligned}
& g\left(\theta; v_1, v_2, d_1, \ldots, d_H, \delta_1, \ldots, \delta_H \middle| \sigma_0, \ldots, \sigma_J\right) \\
= & \frac{\theta^{v_1+\sigma_J-1}\left(1-\theta\right)^{v_2+\sigma_0-1}\prod_{j=1}^{J-1}\left(1-\delta_j\theta\right)^{d_j+\sigma_j-1}\prod_{h=J}^{H}\left(1-\delta_h\theta\right)^{d_h-1}}{\int_0^1 t^{v_1+\sigma_J-1}\left(1-t\right)^{v_2+\sigma_0-1}\prod_{j=1}^{J-1}\left(1-\delta_j t\right)^{d_j+\sigma_j-1}\prod_{h=J}^{H}\left(1-\delta_h t\right)^{d_h-1}\mathrm{d}t} \\
= & g\left(\theta; v_1+\sigma_J, v_2+\sigma_0, d_1+\sigma_1, \ldots, d_{J-1}+\sigma_{J-1}, d_J, \ldots, d_H, \delta_1, \ldots, \delta_H\right),
\end{aligned}
$$

where the last expression is again an Exton-generalized Beta function with parameters updated by the sample, and the last equality follows from the integral representation of the fourth Lauricella function, given below (Lauricella, 1893, p.149):

$$
F_D^{(H)}\left(a; \mu_1, \ldots, \mu_H; c; \gamma_1, \ldots, \gamma_H\right) = \tfrac{1}{\mathrm{B}(a,c-a)} \int_0^1 t^{a-1}\left(1-t\right)^{c-a-1}\prod_{h=1}^{H}\left(1-\gamma_h t\right)^{-\mu_h}\mathrm{d}t. \qquad (6)
$$

■

Proposition 6 (Expectation) *If the random variable x follows an Exton Generalized Beta distribution with $v_1, v_2 \in \mathbb{Z}_+$, the expectation of the function*

$$
\varphi\left(x; k, \zeta_0, \zeta_1, \ldots, \zeta_{S,1}, \ldots, z_S\right) = \zeta_0 x^k \prod_{s=1}^{S}\left(1-\zeta_s x\right)^{z_s-1} \qquad (7)
$$

is

$$E\left[\varphi\left(x\right)\right]=\zeta_0\frac{\Gamma(v_1+v_2)\Gamma(v_1+k)}{\Gamma(v_1)\Gamma(v_1+v_2+k)}\frac{F_D^{(H+S)}(v_1+k,1-d_1,\ldots,1-d_H,1-z_1,\ldots,1-z_S;v_1+v_2+k;\delta_1,\ldots,\delta_H,\zeta_1,\ldots,\zeta_H)}{F_D^{(H)}(v_1,1-d_1,\ldots,1-d_H;v_1+v_2;\delta_1,\ldots,\delta_H)}. \qquad (8)$$

Proof. Using (4) and (7) and collecting terms, we have

$$\begin{aligned}
E\left[\varphi\left(x\right)\right] &= \int_0^1 \varphi\left(x\right)g\left(x\right)\,\mathrm{d}x \\
&= \int_0^1 \frac{\zeta_0 x^{v_1-1+k}\left(1-x\right)^{v_2-1}\prod_{h=1}^{H}\left(1-\delta_h x\right)^{d_h-1}\prod_{s=1}^{S}\left(1-\zeta_s x\right)^{z_s-1}}{\mathrm{B}\left(v_1,v_2\right)F_D^{(H)}\left(v_1,1-d_1,\ldots,1-d_H;v_1+v_2;\delta_1,\ldots,\delta_H\right)}\mathrm{d}x \\
&= \frac{\zeta_0}{\mathrm{B}\left(v_1,v_2\right)}\frac{\int_0^1 x^{v_1+k-1}\left(1-x\right)^{v_2-1}\prod_{h=1}^{H}\left(1-\delta_h x\right)^{d_h-1}\prod_{s=1}^{S}\left(1-\zeta_s x\right)^{z_s-1}\mathrm{d}x}{F_D^{(H)}\left(v_1,1-d_1,\ldots,1-d_H;v_1+v_2;\delta_1,\ldots,\delta_H\right)} \\
&= \zeta_0\frac{\Gamma(v_1+v_2)\Gamma(v_1+k)}{\Gamma(v_1)\Gamma(v_1+v_2+k)}\frac{\int_0^1 x^{v_1+k-1}\left(1-x\right)^{v_2-1}\prod_{h=1}^{H}\left(1-\delta_h x\right)^{d_h-1}\prod_{s=1}^{S}\left(1-\zeta_s x\right)^{z_s-1}\mathrm{d}x}{\mathrm{B}\left(v_1+k,v_2\right)F_D^{(H)}\left(v_1,1-d_1,\ldots,1-d_H;v_1+v_2;\delta_1,\ldots,\delta_H\right)}
\end{aligned}$$

where the last equality follows from re-writing the Beta function as

$$\begin{aligned}
\frac{1}{\mathrm{B}\left(v_1,v_2\right)} &= \frac{\Gamma\left(v_1+v_2\right)}{\Gamma\left(v_1\right)\Gamma\left(v_2\right)} \\
&= \frac{\Gamma\left(v_1+v_2\right)}{\Gamma\left(v_1\right)\Gamma\left(v_2\right)}\frac{\Gamma\left(v_1+k\right)\Gamma\left(v_2\right)}{\Gamma\left(v_1+v_2+k\right)}\frac{\Gamma\left(v_1+v_2+k\right)}{\Gamma\left(v_1+k\right)\Gamma\left(v_2\right)} \\
&= \frac{\Gamma\left(v_1+v_2\right)}{\Gamma\left(v_1\right)\Gamma\left(v_2\right)}\frac{\Gamma\left(v_1+k\right)\Gamma\left(v_2\right)}{\Gamma\left(v_1+v_2+k\right)}\frac{1}{\mathrm{B}\left(v_1+k,v_2\right)} \\
&= \frac{\Gamma\left(v_1+v_2\right)\Gamma\left(v_1+k\right)}{\Gamma\left(v_1\right)\Gamma\left(v_1+v_2+k\right)}\frac{1}{\mathrm{B}\left(v_1+k,v_2\right)}.
\end{aligned}$$

Using once more the integral representation of the Lauricella function given in 6, we replace the numerator and conclude that

$$E\left[\varphi\left(x\right)\right]=\zeta_0\frac{\Gamma(v_1+v_2)\Gamma(v_1+k)}{\Gamma(v_1)\Gamma(v_1+v_2+k)}\frac{F_D^{(H+S)}(v_1+k,1-d_1,\ldots,1-d_H,1-z_1,\ldots,1-z_S;v_1+v_2+k;\delta_1,\ldots,\delta_H,\zeta_1,\ldots,\zeta_H)}{F_D^{(H)}(v_1,1-d_1,\ldots,1-d_H;v_1+v_2;\delta_1,\ldots,\delta_H)}.$$

■

We conclude this appendix by applying these results to the learning process used for computerized evaluators in treatment CL, where the signals and their probabilities conditional on the unknown parameter f, are given in Table 8. Let truth-telling correspond to outcome $i=1$ in the noisy Bernoulli trial, and the three possible signals, $j=0,1,2$, be the negative, the muddy, and the positive signal, respectively. Their likelihoods according to Table 8 are

$$\begin{aligned}
\Pr(j=0|\,f) &= (1-f)\,\frac{1-q}{2} \\
\Pr(j=1|\,f) &= f\left(\frac{1}{4}\right)+(1-f)\left(\frac{1}{2}\right) \\
\Pr(j=2|\,f) &= f\,\frac{3-2q}{4},
\end{aligned}$$

so that $\delta_1 = 1-\frac{1/4}{1/2} = 1/2$. The uniform prior used in our application, corresponds to parameters $v_1 = v_2 = d_1 = 1$, and either $H = 1$ or $d_h = 1$ for all h, of the Exton generalized Beta function. Proposition 5 tells us that after a sample of n negative signals, p positive signals, and m muddy signals, uncertainty about f is given by the density

$$\begin{aligned}
g\left(f;p+1,n+1,m+1,1/2\right) &= \frac{f^p\left(1-f\right)^n\left(1-\frac{1}{2}f\right)^m}{\mathrm{B}\left(p+1,n+1\right)F_D^{(1)}\left(p+1;-m;p+n+2;\frac{1}{2}\right)} \\
&= \frac{f^p\left(1-f\right)^n\left(1-\frac{1}{2}f\right)^m}{\int_0^1 x^p\left(1-x\right)^n\left(1-\frac{1}{2}x\right)^m\mathrm{d}x}.
\end{aligned}$$

Recall that the assessment of an evaluator who receives a Blue report and observes either an orange or a blue core, is a function of f—either $p_{Bb}(f)$ or $p_{Bo}(f)$. Thus, the expected assessments given by an evaluator whose experience so far is the sample (n, m, p) are respectively $E\left[p_{Bb}(f)\right]$ and $E\left[p_{Bo}(f)\right]$, under the Exton generalized Beta density with parameters $v_1 = p+1$, $v_2 = n+1$, $d_1 = m+1$, and $\delta_1 = 1/2$. Recall that

$$\begin{aligned}
p_{Bb}(f) &= \frac{1}{\frac{3}{2}+(1-f)\,\frac{1}{2}} = \frac{1}{2}\left(1-\frac{1}{4}f\right)^{-1}, \text{ and} \\
p_{Bo}(f) &= \frac{(1-f)}{\frac{1}{2}+(1-f)\,\frac{3}{2}} = \frac{1}{2}\,(1-f)\left(1-\frac{3}{4}f\right)^{-1},
\end{aligned}$$

which means these functions satisfy the assumptions of Proposition 6.

References

Abeler, J., D. Nosenzo, and C. Raymond (2019). Preferences for truth-telling. *Econometrica 87*(4), 1115–1153.

Anderson, L. and C. Holt (1997). Information cascades in the laboratory. *American Economic Review 87*(5), 847–862.

Anderson, L. and C. Holt (2008). Information cascade experiments. In C. Plott (Ed.), *Handbook of Experimental Economics Results*, Volume 1, pp. 335–343. Elsevier.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics 107*(3), 797–817.

Battigalli, P. and M. Dufwenberg (2009). Dynamic psychological games. *Journal of Economic Theory 144*(1), 1–35.

Bikhchandani, S., D. Hirschleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Blume, A., E. K. Lai, and W. Lim (Forthcoming). Strategic information transmission: A survey of experiments and theoretical foundations. In M. Capra, R. Croson, M. Rigdon, and T. Rosenblat (Eds.), *Handbook of Experimental Game Theory*. Cheltenham, UK and Northampton, MA: Edward Elgar Publishing.

Çelen, B. and S. Kariv (2004). Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review 94*(3), 484–498.

Çelen, B., S. Kariv, and A. Schotter (2010). An experimental test of advice and social learning. *Management Science 56*(10), 1687–1701.

Charness, G. and M. Dufwenberg (2006). Promises and partnership. *Econometrica 74*(6), 1579–1601.

Chaudhuri, A., A. Schotter, and B. Sopher (2009). Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice. *Economic Journal 119*(534), 91–122.

Deb, R., M. M. Pai, and M. Said (2018). Evaluating strategic forecasters. *American Economic Review 108*(10), 3057–3103.

Dufwenberg, M. and U. Gneezy (2000). Measuring beliefs in an experimental lost wallet game. *Games and Economic Behavior 30*, 163–182.

Dufwenberg, M. and G. Kirchsteiger (2004). A theory of sequential reciprocity. *Games and economic behavior 47*(2), 268–298.

Ederer, F. and A. Stremitzer (2017). Promises and expectations. *Games and Economic Behavior 106*, 161–178.

Ehrbeck, T. and R. Waldmann (1996). Why are professional forecasters biased? agency versus behavioral explanations. *Quarterly Journal of Economics 111*(1), 21–40.

Ellingsen, T., M. Johannesson, S. Tjøtta, and G. Torsvik (2010). Testing guilt aversion. *Games and Economic Behavior 68*, 95–107.

Esponda, I. and E. Vespa (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics 6*(4), 180–202.

Eyster, E., M. Rabin, and G. Weizsäcker (2015). An experiment on social mislearning. CEPR Discussion Paper 11020.

Fehr, E. and J.-R. Tyran (2001). Does money illusion matter? *American Economic Review 91*(5), 1239–1262.

Fehrler, S. and N. Hughes (2018). How transparency kills information aggregation: Theory and experiments. *American Economic Journal: Microeconomics 10*(1), 181–209.

Feltovich, N. (2011). The effect of subtracting a constant from all payoffs in a hawk-dove game: Experimental evidence of loss aversion in strategic behavior. *Southern Economic Journal 77*(4), 814–826.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*, 171–178.

Fréchette, G., A. Lizzeri, and J. Perego (2019). Rules and commitment in communication: An experimental analysis. *Unpublished Manuscript*.

Geanakoplos, J., D. Pearce, and E. Stacchetti (1989). Psychological games and sequential rationality. *Games and Economic Behavior 1*(1), 60–79.

Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of Political Economy 114*(2), 280–316.

Gneezy, U., A. Kajackaite, and J. Sobel (2018). Lying aversion and the size of the lie. *American Economic Review 108*(2), 419–53.

Goeree, J. K., T. R. Palfrey, B. W. Rogers, and R. D. McKelvey (2007). Self-correcting information cascades. *Review of Economic Studies 74*(3), 733–762.

Graham, J. R. (1999). Herding among investment newsletters: Theory and evidence. *Journal of Finance 54*(1), 237–268.

Harrison, G. W., J. Martínez-Correa, and J. T. Swarthout (2013). Inducing risk neutral preferences with binary lotteries: A reconsideration. *Journal of Economic Behavior and Organization 94*, 145–159.

Hartigan, J. A. (1975). *Clustering Algorithms*. Wiley series in probability and mathematical statistics. New York, NY: Wiley.

Hastie, T., R. Tibshirani, J. Friedman, and J. Franklin (2005). The elements of statistical learning: Data mining, inference and prediction. *The Mathematical Intelligencer 27*(2), 83–85.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies 66*(1), 169–182.

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

Hung, A. A. and C. R. Plott (2001). Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review 91*(5), 1508–1520.

Irlenbusch, B. and D. Sliwka (2006). Career concerns in a simple experimental labour market. *European Economic Review 50*(1), 147–170.

Keane, M. P. and D. Runkle (1998). Are financial analysts' forecasts of corporate profits rational? *Journal of Political Economy 106*(4), 768–805.

Khalmetski, K. (2016). Testing guilt aversion with an exogenous shift in beliefs. *Games and Economic Behavior 97*, 110–119.

Koch, A. K., A. Morgenstern, and P. Raab (2009). Career concerns incentives: An experimental test. *Journal of Economic Behavior and Organization 72*(1), 571–588.

Koch, C. and S. P. Penczynski (2018). The winner's curse: Conditional reasoning and belief formation. *Journal of Economic Theory 174*, 57–102.

Kübler, D. and G. Weizsäcker (2003). Information cascades in the labor market. *Journal of Economics 80*(3), 211–229.

Lamont, O. A. (2002). Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior and Organization 48*(3), 265–280.

Lauricella, G. (1893). Sulle funzioni ipergeometriche a più variabili. *Rendiconto del Circolo Matematico di Palermo 7*(1), 111–158.

Levy, G. (2007). Decicion making in committees: Transparency, reputation, and voting rules. *American Economic Review 97*(1), 150–168.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability 1*(14), 281–297.

Mattozzi, A. and M. Y. Nakaguma (2019). Public versus secret voting in committees.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective.* Adaptive computation and machine learning series. MIT Press.

Ottaviani, M. and P. N. Sørensen (2001). Information aggregation in debate: Who should speak first? *Journal of Public Economics 81*(3), 393–421.

Ottaviani, M. and P. N. Sørensen (2006a). Reputational cheap talk. *RAND Journal of Economics 37*(1), 155–175.

Ottaviani, M. and P. N. Sørensen (2006b). The strategy of professional forecasting. *Journal of Financial Economics 81*(2), 441–466.

Prat, A. (2005). The wrong kind of transparency. *American Economic Review 95*(3), 862–877.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American economic review*, 1281–1302.

Renes, S. and B. Visser (2018). Markets assessing decision makers and decision makers impressing markets: a lab experiment. Tinbergen Institute Discussion Papers 18-070/VII.

Roth, A. E. and J. K. Murnighan (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical psychology 17*(2), 189–198.

Scharfstein, D. S. and J. C. Stein (1990). Herd behavior and investment. *American Economic Review 80*(3), 465–479.

Schotter, A. (2003). Decision making with naive advice. *American Economic Review Papers & Proceedings 93*(2), 196–201.

Schotter, A. and B. Sopher (2007). Advice and behavior in intergenerational ultimatum games: An experimental approach. *Games and Economic Behavior 58*(2), 365–393.

Trueman, B. (1994). Analyst forecasts and herding behavior. *Review of Financial Studies 7*(1), 97–124.

Visser, B. and O. H. Swank (2007). On committees of experts. *Quarterly Journal of Economics 122*(1), 337–372.

Online Appendix

A Proofs

Proof of Proposition 1. Part (a) follows from Bayes updating by the evaluator as stated in equation (2). To compute evaluators' assessments after receiving a report $R = O$, recall that in this case, $\Pr(R|c, u, f) = \Pr(S|c, u)$, since the report can only be Orange if truthful. Noticing further that $\Pr(S = O|c = b, u = I) = 0$, $\Pr(S = O|c = o, u = I) = 1$, and $\Pr(S = O|c = o, u = U) = 1/2$, we have

$$p_{Ob} = \Pr(u = I|S = O, c = b) = \frac{\frac{1}{2}\Pr(S = O|c = b, u = I)}{\frac{1}{2}\left[\Pr(S = O|c = b, u = I) + \Pr(S = O|c = b, u = U)\right]} = 0.$$

$$p_{Oo} = \Pr(u = I|S = O, c = o) = \frac{\frac{1}{2}\Pr(S = O|c = o, u = I)}{\frac{1}{2}\left[\Pr(S = O|c = o, u = I) + \Pr(S = O|c = o, u = U)\right]} = \frac{2}{3}$$

If, instead, the report is $R = B$, the color of the Shell may be either B or O, and thus

$$\Pr(R|c, u, f) = \Pr(R|S = O, f)\Pr(S = O|c, u) + \Pr(R|S = B, f)\Pr(S = B|c, u) \Rightarrow$$
$$\Pr(R = B|c, u, f) = (1 - f)\Pr(S = O|c, u) + \Pr(S = B|c, u),$$

where we used the facts that (i) the reporter does not know the realizations of c and u when sending the report, (ii) the joint distribution of signal, state, and reporter type does not depend on the evaluator's beliefs, and (iii) after a blue signal the report is always $R = B$. Appropriately replacing the above in equation (2), we obtain

$$p_{Bb} = \Pr(u = I|R = B, c = b, f) = \frac{\frac{1}{2}}{\frac{1}{2}\left[1 + (1 - f)\frac{1}{2} + \frac{1}{2}\right]} = \frac{1}{\frac{3}{2} + (1 - f)\frac{1}{2}},$$

and

$$p_{Bo} = \Pr(u = I|R = B, c = o, f) = \frac{(1 - f)\frac{1}{2}}{\frac{1}{2}\left[(1 - f) + (1 - f)\frac{1}{2} + \frac{1}{2}\right]} = \frac{(1 - f)}{\frac{1}{2} + (1 - f)\frac{3}{2}}.$$

Note that p_{Bb} is continuous, strictly increasing, and convex in f for $f \in [0, 1]$, so it attains a minimum at $f = 0$ (where $p_{Bb} = 1/2$) and a maximum at $f = 1$ ($p_{Bb} = 2/3$). Similarly, p_{Bo} is strictly decreasing and concave for $f \in [0, 1]$; thus, it attains a minimum at $f = 1$ ($p_{Bo} = 0$) and a maximum at $f = 0$ ($p_{Bo} = 1/2$). This concludes the proof of part (a).

To establish part (b), it is enough to note that $p_{Bb} \in [1/2, 2/3]$ and $p_{Bo} \in [0, 1/2]$. From part (a), this implies $p_{Bb} \geq p_{Bo}$ for all $f \in [0, 1]$, with $p_{Bb} = p_{Bo}$ whenever $f = 0$. Recalling that $p_{Oo} = 2/3$ we can also conclude that for all $f \in [0, 1]$, $p_{Oo} \geq p_{Bb}$, with $p_{Oo} > p_{Bb}$ for $f < 1$. Since $p_{Ob} = 0$ we can conclude that $p_{Bo} \geq p_{Ob}$ with $p_{Bo} > p_{Ob}$ if and only if $f < 1$. Summing up, $p_{Oo} \geq p_{Bb} \geq p_{Bo} \geq p_{Ob}$ for all $f \in [0, 1]$, with $p_{Oo} > p_{Bb} > p_{Bo} > p_{Ob}$ for

$f \in (0,1)$. For $f=0$, $p_{Oo} > p_{Bb} = p_{Bo} > p_{Ob}$. For $f=1$, $p_{Oo} = p_{Bb} > p_{Bo} = p_{Ob}$.

Proof of Proposition 2. For the first part, we must establish that the reporter's gain from misreporting is increasing in q. Since misreporting only matters when the observed signal is O, this gain is given by $\Pr(S=O)\,\Delta_{EU}(q,f)$, and, from equation (1),

$$\Delta_{EU}(q,f) = (1-q_O(q))\,[p_{Bo}(f) - p_{Oo}] + q_O(q)\,[p_{Bb}(f) - p_{Ob}]\,.$$

Note that the evaluator's assessments only depend on f. This is shown in the proof of proposition 1, and follows intuitively because evaluators make their assessments *after* observing the state of the world, and because their belief, f, is assumed exogenous. Therefore, the probability of the state of the world, q, gives the evaluator no additional information about the color of the shell or of the core. Instead, the posterior probability that the shell is B or O, does depend on q as follows:

$$\begin{aligned} q_O &= \Pr(c=b|S=O) = \tfrac{\Pr(S=O|c=b)q}{\Pr(S=O)} = \tfrac{\frac{1}{4}q}{\Pr(S=O)}, \text{ and} \\ 1-q_O &= \Pr(c=o|S=O) = \tfrac{\Pr(S=O|c=o)(1-q)}{\Pr(S=O)} = \tfrac{\frac{3}{4}(1-q)}{\Pr(S=O)} \end{aligned}$$

Therefore, the gain from misreporting is the following function of q and f:

$$\Pr(S=O)\,\Delta_{EU}(q,f) = \frac{3}{4}(1-q)\,[p_{Bo}(f) - p_{Oo}] + \frac{1}{4}q\,[p_{Bb}(f) - p_{Ob}]\,, \tag{9}$$

and, thus

$$\frac{\partial\,[\Pr(S=O)\Delta_{EU}(q,f)]}{\partial q} = \frac{3}{4}(p_{Oo} - p_{Bo}(f)) + \frac{1}{4}(p_{Bb} - p_{Bo}) > 0,$$

since, by proposition 1, $(p_{Oo} - p_{Bo}(f)) > 0$ and $(p_{Bb} - p_{Bo}) \geq 0$ for all values of f.

For the second part of the proposition, replace the values of the evaluator's assessments in equation (9) (the values are given in Section 3.1 and again in the proof of proposition 1), to obtain

$$\Pr(S=O)\,\Delta_{EU}(q,f) = \frac{3}{4}(1-q)\left[\frac{(1-f)}{\frac{1}{2}+(1-f)\frac{3}{2}} - \frac{2}{3}\right] + \frac{1}{4}q\left[\frac{1}{\frac{3}{2}+(1-f)\frac{1}{2}}\right].$$

If the evaluator holds point beliefs about the reporter's truthfulness, the reporter should misreport only if the above expression—gain from misreporting—is positive, which boils down to $q > (4-f)\,/\,[4(2-f)]$. This concludes the proof of Proposition 2.

Proof of Proposition 3. Consider Δ_{EU} from equation (10). We have:

$$\frac{\partial\,[\Pr(S=O)\Delta_{EU}(q,f)]}{\partial f} = \frac{(24q-6)f^2 + (-96q+48)f + 128q - 96}{(3f^2 - 16f + 16)^2}.$$

The denominator is always positive. Thus, the derivative is negative if and only if the numerator is negative. The numerator is negative if and only if

$$q < \frac{3f^2 - 24f + 48}{12f^2 - 48f + 64}.$$

This means that, if $q \leq \frac{3}{4}$, the derivative is negative for any $f \in [0,1]$, and hence, it follows that the expected gain from misreporting is minimized at $f = 1$. When $q > \frac{3}{4}$, the numerator is negative if and only if

$$f > \frac{q - (1-q) - 2\sqrt{\frac{1}{3}q(1-q)}}{\frac{3}{4}q - \frac{1}{4}(1-q)} = \overline{f}(q).$$

As long as $q < 27/28$, there is $f \in [0,1]$ such that this condition is satisfied. Thus, for prior beliefs $q \in [3/4, 27/28]$, $\Pr(S = O)\,\Delta_{EU}(q, f)$ is a concave function of f, minimized either at $f = 0$ or $f = 1$. It is easy to see that, as long as $q < 9/10$, $\Pr(S = O)\Delta_{EU}(q, 1) = \frac{8q-6}{9-6q} < \Pr(S = O)\Delta_{EU}(q, 0) = \frac{16q-8}{48-32q}$, thus completing the proof of the statement.

Additionally, notice that for $q \in [9/10, 27/28]$, the concave incentives for misreporting are minimized when $f = 0$. Moreover, when $q > 27/28$, expected gains from misreporting are a strictly increasing function of f, immediately implying that they are minimized at $f = 0$ (and maximized at $f = 1$).

Proof of Proposition 4. First consider a strongly unbalanced prior belief about the state, $q \in (3/4, 1]$, and notice that the condition for the reporter to prefer misreporting, $q > (4-f)/[4(2-f)]$, is satisfied for all values of f: the right-hand side is strictly increasing in f, and equals 3/4 when $f = 1$. This means that there can only be equilibria where the reporter misreports. Since the condition is also satisfied when $f = 0$—when the evaluator believes the reporter misreports for sure—misreporting and misreporting beliefs indeed constitute a perfect Bayesian Nash equilibrium. Hence, with a strongly unbalanced prior, the unique equilibrium of the game is a pooling equilibrium.

Now consider a mildly unbalanced prior, q, and suppose the evaluator holds beliefs $f^*(q) = \frac{8q-4}{4q-1}$. Simple transformations give a simplified expression for the reporter's expected gain from misreporting,

$$\Pr(S = O)\,\Delta_{EU}(q, f) = \frac{4q(2-f) + (f-4)}{2(4-f)(4-3f)},$$

whose denominator is always positive. Replacing $f = \frac{8q-4}{4q-1}$, the numerator of the above expression equals 0, meaning that the reporter is indifferent between misreporting and truth-telling. Hence, the evaluator's beliefs that the reporter picks a mixed strategy of truth-telling with probability $\frac{8q-4}{4q-1}$ can indeed be sustained with reporter's best-replying behavior for those beliefs. This shows that the MSE exists.

If the evaluator holds beliefs $f > \frac{8q-4}{4q-1}$, $\Delta_{EU} < 0$, meaning that the reporter prefers to be truthful. Therefore, the only belief $f > \frac{8q-4}{4q-1}$ that can be sustained by the reporter's behavior, is $f = 1$, where the reporter is truthful and the evaluator believes this: a separating equilibrium exists.

If the evaluator holds beliefs $f < \frac{8q-4}{4q-1}$, $\Delta_{EU} > 0$, meaning that the reporter prefers misreporting. Therefore, the only belief $f < \frac{8q-4}{4q-1}$ that can be sustained by the reporter's best-replying behavior is $f = 0$, where the reporter misreports and the evaluator believes so: a pooling equilibrium exists.

$$\Delta_{EU}(q, f) = EU(M) - EU(T) = \frac{\frac{3}{4}(1-q)}{\frac{1}{4}q+\frac{3}{4}(1-q)} \left(\frac{(1-f)}{\frac{1}{2}+(1-f)\frac{3}{2}} - \frac{2}{3} \right) + \frac{\frac{1}{4}q}{\frac{1}{4}q+\frac{3}{4}(1-q)} \left(\frac{1}{\frac{3}{2}+(1-f)\frac{1}{2}} \right). \quad (10)$$

B The Effect of Experience

The game faced by our experimental subjects is complicated and it might require some time for subjects to understand the underlying incentives. This is the reason why we focused the analyses on experienced subjects, as is customary in experimental economics. To explore the possibility that behavior adapted to accumulated experience, we compare reporters' and evaluators' behavior in the first block of each treatment (decisions 1–16), when subjects were relatively inexperienced, to the second block (decisions 17–32), after subjects had been exposed to feedback and a chance to learn. Tables 13, 14 and 15 report summary statistics on reporters' and evaluators' behavior by block. Figures 10 and 11 show reporters' and evaluators' behavior by period and treatment. Table 16 reports estimates of the effect of experience on reporters' behavior as a function of the game and the treatment. Table 17 reports estimates of the effect of experience on human evaluators' behavior as a function of the treatment and the observed report-core pair.

With the exception of CT with $q = 6/10$ (where there is significantly less misreporting with experience) and of CU with $q = 8/10$ (where the effect of experience is negligible), experienced reporters are significantly less likely to report truthfully than inexperienced reporters (significance at the 1% level, except for HF with $q = 6/10$ and CL with $q = 8/10$, which are significant at, respectively, the 10% and the 5% level). Notice that in game CT with $q = 6/10$ reporting truthfully is the best response to the beliefs of computerized evaluators so experience leads reporters to make better choices, as is the case for HF with $q = 8/10$ and CT with $q = 8/10$. Regarding CU with $q = 6/10$, learning is away from the best response: experienced reporters are more likely to misreport than inexperienced reporters but misreporting gives a lower EU than truth-telling. At the same time, we must note that the differences in EU between misreporting and reporting is minimal: €0.01 in each period. The behavior of evaluators is only marginally affected by experience: experienced evaluators punish significantly less severely an inaccurate blue report with $q = 6/10$ (possibly as a consequence of increased misreporting by experienced reporters, which dampens the evaluators' ability to infer the informativeness of the urn) and punish more severely (significant at the

5% level) an inaccurate orange report (in the direction of what a Bayesian evaluator would do). Finding 7 summarizes this discussion.

FINDING 7: Except when reporting truthfully is a best response to computerized evaluators' beliefs (game CT with $q = 6/10$), experienced reporters are more likely to misreport than inexperienced reporters. Human evaluators' assessments are mostly unaffected by experience.

	$q = 6/10$			$q = 8/10$		
	1st Block	2nd Block	Theory	1st Block	2nd Block	Theory
HF	0.53	0.49	[0,1]	0.44	0.37	0
	(752)	(745)		(736)	(736)	
CT	0.59	0.65	1	0.43	0.34	0
	(752)	(752)		(752)	(752)	
CU	0.52	0.44	1	0.38	0.35	0
	(725)	(752)		(752)	(752)	
CL	0.54	0.47	[0,1]	0.41	0.36	0
	(768)	(768)		(768)	(768)	

Table 13: Fraction of periods the reporters choose the truthful plan of action, by treatment and experience. The number of observations is in parentheses. Each reporter makes 16 decisions in each treatment and block. There are 46 reporters in HF; 47 in CU and CT; and 48 in CL. In HF, there is one additional reporter making 16 decisions in the 1st Block with $q = 6/10$ and 9 decisions in the 1st Block with $q = 8/10$.

	Assessments		Empirical Frequency		Theory
	1st Block	2nd Block	1st Block	2nd Block	
Blue Report & Blue Core	0.60	0.59	0.56	0.55	[0.50, 0.67]
	(385)	(383)	(385)	(383)	
Blue Report & Orange Core	0.30	0.40	0.37	0.37	[0, 0.50]
	(196)	(194)	(196)	(194)	
Orange Report & Blue Core	0.23	0.01	0.00	0.00	0.00
	(60)	(54)	(60)	(54)	
Orange Report & Orange Core	0.66	0.65	0.64	0.68	0.67
	(111)	(105)	(111)	(105)	

Table 14: Median assessment by the evaluators, by observed Report-Core and experience, $q = 6/10$. The number of observations is in parentheses. There are 47 the evaluators, each making 16 assessments in each treatment and each block. Because of missing observations by the corresponding reporter, there are 7 missing observations for the 1st Block with $q = 8/10$ and 16 missing observations for the 2nd Block with either q.

	Assessments		Empirical Frequency		Theory
	1st Block	2nd Block	1st Block	2nd Block	
Blue Report & Blue Core	0.60	0.55	0.59	0.56	[0.50, 0.67]
	(536)	(522)	(536)	(522)	
Blue Report & Orange Core	0.36	0.30	0.37	0.46	[0, 0.50]
	(113)	(117)	(113)	(117)	
Orange Report & Blue Core	0.10	0.01	0.10	0.01	0
	(56)	(59)	(56)	(59)	
Orange Report & Orange Core	0.48	0.66	0.65	0.89	0.67
	(40)	(38)	(40)	(38)	

Table 15: Median assessment by the evaluators, by observed Report-Core and experience, $q = 8/10$. The number of observations is in parentheses. There are 47 the evaluators, each making 16 assessments in each treatment and each block. Because of missing observations by the corresponding reporter, there are 7 missing observations for the 1st Block with $q = 8/10$ and 16 missing observations for the 2nd Block with either q.

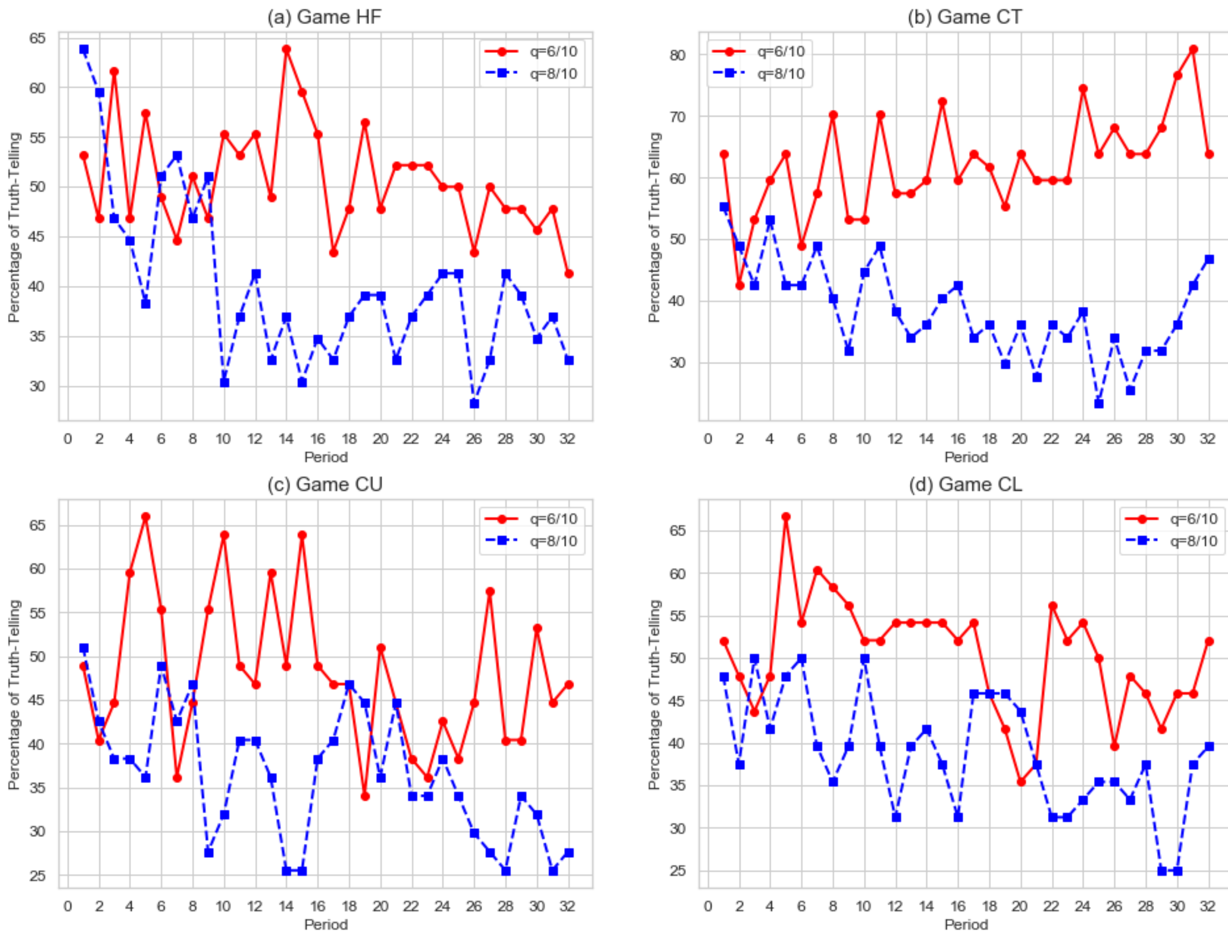

Figure 10: Time Series of Reporters' Behavior by Game and Prior Beliefs on State (q).

	Pr[Reporter Chooses Truthful Plan of Action]							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2^{nd} Block	-0.04	-0.07	0.07	-0.09	-0.08	-0.03	-0.07	-0.05
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Game	HF	HF	CT	CT	CU	CU	CL	CL
q	6/10	8/10	6/10	8/10	6/10	8/10	6/10	8/10
N	1488	1481	1504	1504	1504	1504	1536	1536

Table 16: Random effects GLS regressions. Each subject is a panel and periods are times within a panel. Standard errors in parentheses. Constant is omitted.

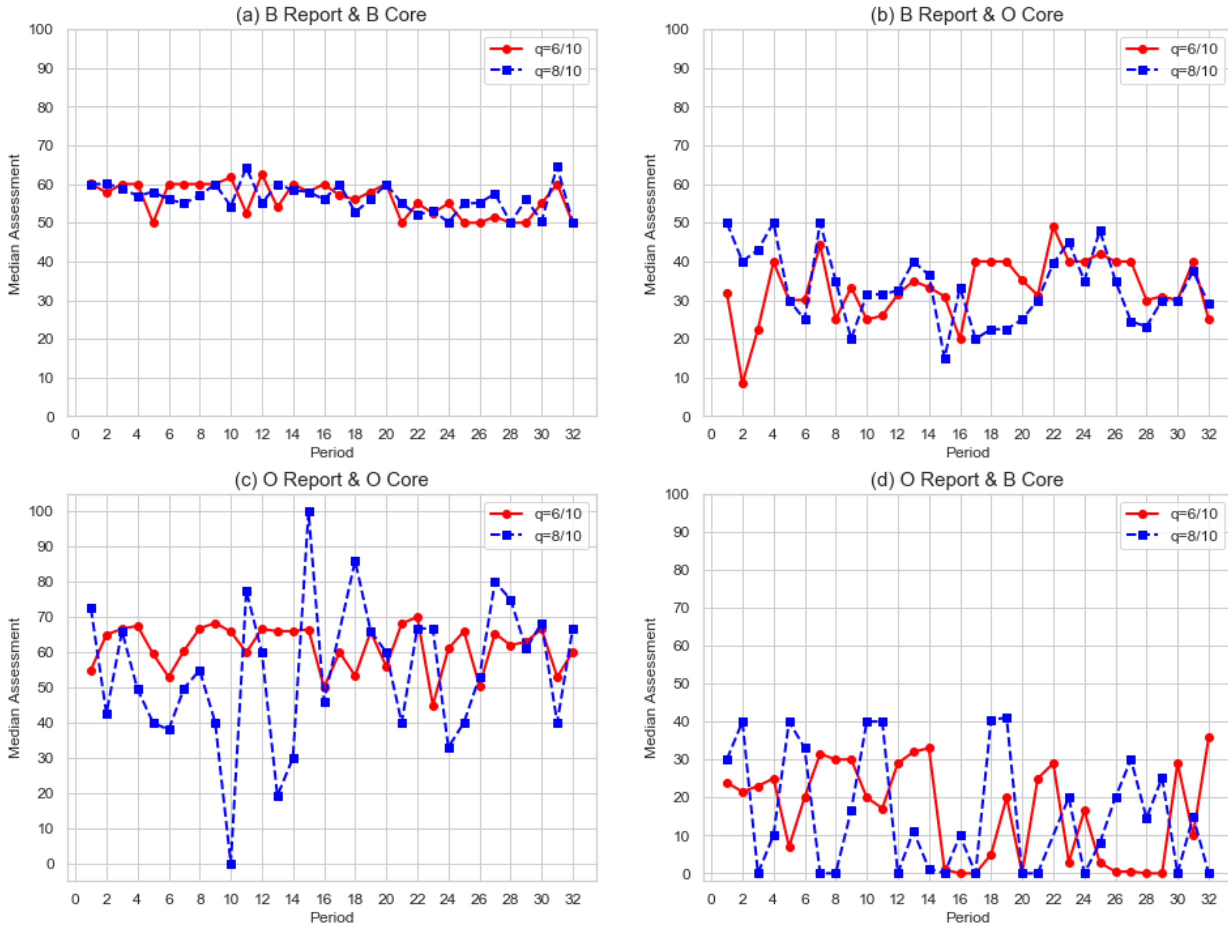

Figure 11: Time Series of Average Human Evaluators' Assessments by Observed Report-Core Pair and Prior Beliefs on State (q).

	Human Evaluator's Assessment							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2^{nd} Block	0.04	4.86	-2.14	-0.62	-0.66	-3.13	-5.80	1.67
	(1.00)	(1.36)	(2.68)	(1.94)	(0.70)	(2.19)	(2.89)	(5.02)
Report	Blue	Blue	Orange	Orange	Blue	Blue	Orange	Orange
Core	Blue	Orange	Blue	Orange	Blue	Orange	Blue	Orange
q	6/10	6/10	6/10	6/10	8/10	8/10	8/10	8/10
N	768	390	114	216	1058	230	115	78

Table 17: Random effects GLS regressions. Game HF. Each subject is a panel and periods are times within a panel. Standard errors in parentheses. Constant is omitted.

C Experimental Instructions

Experimental instructions were delivered in print and using a video of power point slides with explanations of the situation and decisions to be made. The videos for each game can be found at the following web addresses:

- For game CT, `http://bocconicohortstudy.org/t11a.mp4`
- For game CU, `http://bocconicohortstudy.org/t11u.mp4`
- For game CL, `http://bocconicohortstudy.org/t22.mp4`
- For game HF, `http://bocconicohortstudy.org/t33.mp4`

Here we reproduce the words and some of the figures contained in the slides handed out to the subjects. Information outside boxes is relevant for all games; information inside boxes is relevant for a specific game or set of games only, as indicated in the title of the box. Some wording is slightly different between games CT, CU, and CL, on one side, and game HF on the other. These alternate wordings are indicated inside square brackets, with the wording used in game HF indicated in italics. We use square brackets and small caps to insert comments about the graphical interface of the delivered instructions.

An Experiment With Balls. Instructions
<u>Welcome</u>

- In this experiment your earnings will depend on your decisions, so that different participants may earn different amounts
- Your earnings will be paid in cash at the end of the session in a separate room to preserve the confidentiality of your scores
- Please be aware that your participation is voluntary and can be withdrawn at any time without giving any reasons, but in that case your earnings will be nil

<u>Informed Consent Form</u>

- Please read carefully the <u>Information for Data Subjects and Consent Request</u> document handed out along with these instructions. Please tick, date, and sign the Informed Consent Form at the end of that document
- The data will be collected in an anonymous way by associating a code with your identity
- The users of the data will associate the data with the code, but they will never be able to associate the data with your individual identities
- The anonymized data will be stored and analyzed by the Principal Investigator for the purpose of a research project on reporting and evaluating

- The anonymized data will be kept indefinitely by the Principal Investigator and will be made available to other researchers if and when the project leads to a publication in a scientific journal

Practicalities

- Please remember to turn off your cell phones
- Once the experiment starts, please do not talk or in any way communicate with other participants
- If you have any question or problem at any point, please raise your hand
- Participants intentionally violating rules may be asked to leave the experiment and may not be paid
- You can contact Marco Ottaviani (`marco.ottaviani@unibocconi.it`), the project's Principal Investigator, to ask for corrections, updates, or cancellation of your data at any time
- In case of ethical concerns related to the experiment, you can contact Bocconi's Ethical Committee (`comitatoeticoricerca@unibocconi.it`)

The Experiment

- This experiment consists of **four (4) blocks** of periods
- Each block consists of **sixteen (16) periods**

Games CT, CU, and CL only

- In each period you will play the role of reporter and you will interact with a **computerized evaluator**

Game HF only

- In each period you interact with another participant
- Half of you are assigned the role of reporter, the other half the role of evaluator
- You maintain the role assigned in the first period for the entire experiment

[Screenshots are shown to illustrate the initial message which assigns the role of reporter or evaluator to each subject]

- In each period a reporter is randomly paired with an evaluator
- If you are a reporter, in each period you are equally likely to be paired with any of the evaluators, regardless of the evaluator you were paired with in the previous period

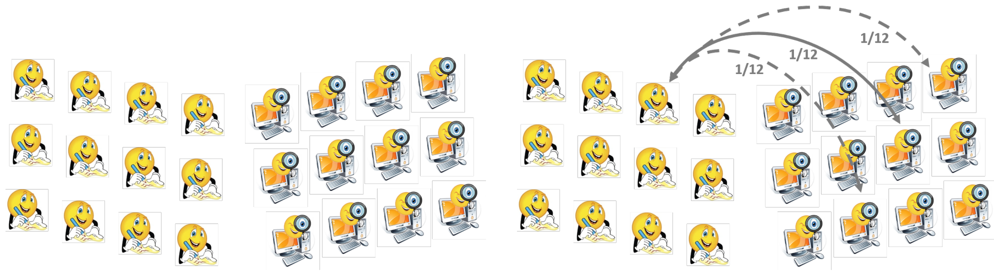

Figure 12: Game CL only: diagram to illustrate reporters' matching with computerized evaluators.

- You will never know the identity of the evaluators you are paired with
- If you are an evaluator, in each period the same mechanism randomly pairs you with a reporter, whose identity you will never know

GAME CL ONLY

[ACCOMPANIED BY THE DIAGRAM IN FIGURE 12.]

Reporters & Computerized Evaluators

The number of computerized evaluators is the same as the number of reporters, which in turn is equal to the number of experimental subjects in this room

Random Pairing

In each period you will be randomly paired with one of the computerized evaluators
Regardless of the computerized evaluator you were paired with in the previous period, in each period you are equally likely to be paired with any of the computerized evaluators

Balls

- In each period the software draws a **ball**
- Each ball is made of two parts: a crystal **inner core** and an opaque **outer shell**
- The inner core is either **blue** or **orange**; similarly, the outer shell that covers the core is either **blue** or **orange**
- Overall, there are four kinds of balls:
    1. Balls with blue core and blue shell
    2. Balls with blue core and orange shell
    3. Balls with orange core and blue shell
    4. Balls with orange core and orange shell

Urns

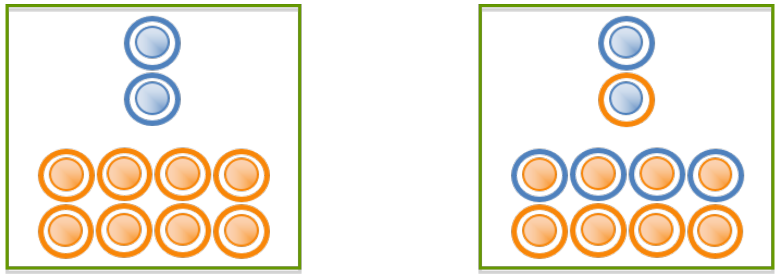

Figure 13: Informative (left) and uninformative (right) urn used as an example in the experimental instructions.

- The ball is drawn from one of **two urns** [Figure 13 is shown]
- The number of balls in each of the two urns is always equal to 10
- In each urn the number of balls with a **blue core** is equal to Q
- At the beginning of a block of periods you are told the number of balls with a blue core, Q, contained in each urn in every period of that block; the remaining 10 - Q balls in each urn have an orange core
- In the example above, in both urns Q=2 balls have a blue core, so that the remaining 10 - Q=8 balls have an orange core

<u>The Informative Urn</u>

In the informative urn, the core of each and every ball is covered by a shell of the same color
EXAMPLE [The left panel of Figure 13 is shown]: The informative urn contains:

- Two (Q=2) balls with a blue core and a blue shell
- Eight (10 - Q=8) balls with an orange core and an orange shell

<u>The Uninformative Urn</u>

In the uninformative urn, for half of the balls the core is covered by a shell of the same color, and for the remaining half of the balls the core is covered by a shell of the other color
EXAMPLE [The right panel of Figure 13 is shown]

- Out of the two (Q=2) balls with blue core, one (2/2=1) is covered by an orange shell and one by a blue shell
- Out of the eight (10 - Q=8) balls with orange core, four (8/2=4) are covered by an orange shell and four by a blue shell

Notice that in the uninformative urn five (5) balls always have a blue shell and five (5) balls always have an orange shell

<u>Draw</u>

- At the beginning of each period, the computer will simulate the toss of a fair coin to determine from which of the two urns the ball is drawn
- If the coin lands Heads, the ball will be drawn from the informative urn
- If the coin lands Tails, the ball will be drawn from the uninformative urn
- When the ball is drawn, neither you (the reporter) nor the [computerized] evaluator know the outcome of the coin toss

Thus nobody knows from which of the two urns the ball is drawn

[Your Task] [*Task of the Reporter*]

[Your task as reporter] [*The task of the reporter*] is to make a report about the color of the shell

- The report has to be made through a plan to which [you] [*the reporter*] must commit before seeing the color of the shell

You [*The reporter*] must choose one of the following two plans:

(1) If I see a BLUE shell, I will report: "The shell is BLUE". If I see an ORANGE shell, I will report: "The shell is ORANGE".
(2) If I see a BLUE shell, I will report: "The shell is BLUE". If I see an ORANGE shell, I will report: "The shell is BLUE".

EXAMPLE: [Caricature of a reporter who thinks the following sentence.] If I see an ORANGE shell, I will report "The shell is BLUE".

[Screenshots are given to illustrate how this choice can be made using the computer interface of the experiment. See Figure 14.]

Implementation of Plan of Action

- After submitting the plan, [you] [*the reporter*] see the color of the shell of the ball that was actually drawn
- At this point, a report is automatically sent to the computerized evaluator according to [the plan you have previously chosen] [*the plan previously chosen by the reporter*]
- Recall that the report sent to the computerized evaluator is determined both by [your] [*the reporter's*] plan and by the color of the shell of the ball that was actually drawn
- Notice that the plan is made before [you] [*the reporter*] see the actual color of the shell

EXAMPLE: If I see an ORANGE shell, I will report "The shell is BLUE".
The following ball is drawn [Graphical display of a ball with an orange shell and an orange core. The shell is then isolated for the reporter to see. A

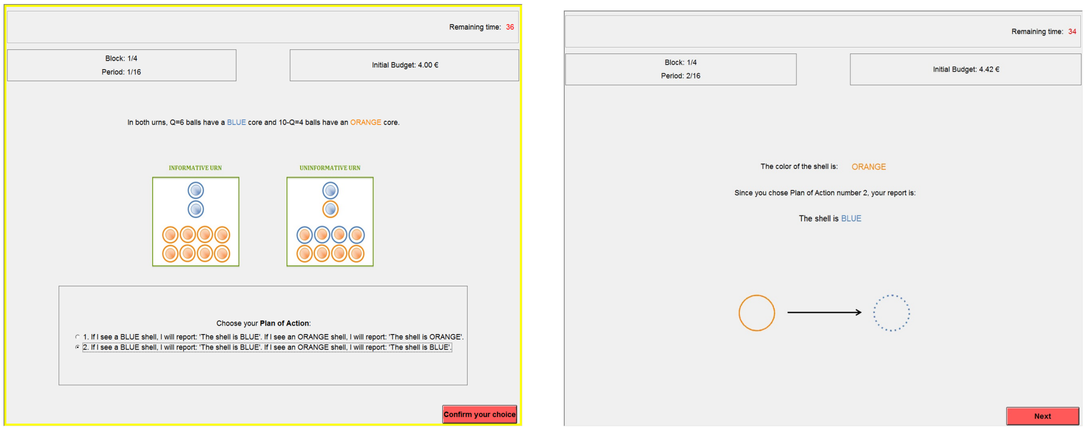

Figure 14: The reporter chooses plan of action (2) (left), a ball is drawn that has an Orange shell, so the reporter's report is Blue (right).

DASHED BLUE SHELL (INDICATING THE REPORT) IS THEN SENT TO THE EVALUATOR.]

[Your goal as reporter] [*The goal of the reporter*] is to be perceived as having seen a ball drawn from the informative urn

[SCREENSHOTS ARE GIVEN TO ILLUSTRATE HOW THE SHELL IS SHOWN TO THE REPORTER AND A REPORT IS AUTOMATICALLY SENT USING THE RULE GIVEN BY THE REPORTER'S CHOSEN PLAN OF ACTION. SEE FIGURE 14.]

Task of the [Computerized] Evaluator

The task of the [computerized] evaluator is to assess how likely it is that the ball was drawn from the informative urn

The [computerized] evaluator makes the assessment after receiving two pieces of information:

- The **report** sent by the reporter about the color of the shell
- The color of the **core** of the ball that has been drawn

GAME CT ONLY

- Throughout all the periods of this experiment, the computerized evaluator is programmed to believe that you always use plan (1)
  - Thus, in each period, the evaluator you face believes that the color of the shell you report is equal to the color of the shell you see

Game CU only

Throughout all the periods of this experiment, the computerized evaluator is programmed to interpret the report based on the belief that:

- A fraction f of the reporters uses plan (1) and a fraction $1 - f$ uses plan (2)
- All values of f between 0 and 1 are equally likely

This means, for example, that the computerized evaluator believes that the probability that a fraction $f = 2/10$ of reporters use plan (1) is the same as the probability that a fraction $f = 9/10$ of reporters use plan (1), and so on for all possible values of the fraction f

Game CL only

- In order to interpret the report and assess whether the ball was drawn from the informative urn, the computerized evaluator is programmed to believe that a fraction f of the reporters uses plan (1) and a fraction $1 - f$ uses plan (2)
- However, computerized evaluators do not know the value of f

Experience and Dynamics of Beliefs

Computerized evaluators accumulate experience across periods with the same value of Q, so that their belief about f evolves depending on their experience

- In each period, the **experience** of each computerized evaluator consists of the outcome of the interaction with the reporters with whom this computerized evaluator was paired in all previous periods with the same value of Q
- In the first period of a block of periods with a value of Q that has never been encountered before, all evaluators believe that any value of f between 0 and 1 is equally likely
- Thus, in the first period, the computerized evaluator believes, for example, that the probability that a fraction $f = 2/10$ of reporters use plan (1) is the same as the probability that a fraction $f = 9/10$ of reporters use plan (1), and so on for all possible values of the fraction f
- Each computerized evaluator updates its belief about the fraction f on the basis of the experience accumulated in each of the previous **individual** interactions with reporters. This experience consists of:
  - The **reports** received by that specific computerized evaluator
  - The color of the **cores** observed by that specific computerized evaluator

- Whether each ball was drawn from the informative or the uninformative urn

This experience allows the computerized evaluator to make an inference about the plan used by the reporters it encountered in all previous periods

- Note that the first time a block of periods with a certain Q starts, learning from experience starts anew
- The "memory" of the computerized evaluator is then reset to believe that all values of f between 0 and 1 are equally likely
- However, if a block starts a second time with the same Q as in an earlier block, the computerized evaluator carries over the experience from the earlier block with that same Q

Task of the [Computerized] Evaluator, continued

- The assessment of the [computerized] evaluator takes the following form:

"Given the core that I see and the reported shell, how likely is it that the ball was drawn from the informative urn? My assessment is P% =__%."

- The number P is between 0 and 100

The goal of the [computerized] evaluator is to make an accurate assessment
EXAMPLE: The following ball is drawn [GRAPHICAL DISPLAY OF A BALL WITH AN ORANGE SHELL AND AN ORANGE CORE. THE CORE IS SEPARATED FROM THE SHELL. THE CORE IS DIRECTLY GIVEN TO THE EVALUATOR TO SEE. THE SHELL IS GIVEN TO THE REPORTER WHO SENDS A DASHED BLUE SHELL (REPORT) TO THE EVALUATOR. THE GRAPHIC EVALUATOR PONDERS:] Given the core that I see and the reported shell, how likely is it that the ball was drawn from the informative urn? My assessment is P% = __%.

Notice that the [computerized] evaluator sees [your] [*the reporter's*] report, but sees neither the reporter's plan nor the actual color of the shell

GAME HF ONLY

[SCREENSHOTS ARE GIVEN TO ILLUSTRATE THE INFORMATION THE EVALUATOR WILL HAVE AT THE TIME WHEN SHE/HE WILL MAKE HER/HIS CHOICE. SEE FIGURE 15.]

[Your Payoff] [*Payoff of the Reporter*]

- At the beginning of each block of periods [you] [*the reporter*] receive a budget of 4 euros

Figure 15: The evaluator is reminded that the report she/he will see is the choice of the reporter (up, left), and is given the opportunity to make a choice after receiving the report and observing the core of the drawn ball (up, right). The evaluator can make her/his choice using a slider (down, left), or by typing in a number (down, right).

- In each period [you] [*the reporter*] pay an operating fee of 25 euro cents and obtain a payoff equal to P euro cents
- P% represents the [computerized] evaluator's assessment of the probability that the ball was drawn from the informative urn

[A screenshot is shown to illustrate how feedback is given to the reporter about her choice, her payoff, and the truth about the core of the drawn ball and the urn informativeness. See Figure 16.]

Game HF only

Payoff of the Evaluator

The payoff structure of the evaluator is designed to give the evaluator an incentive to make and report an accurate assessment of the probability that the ball was drawn from the informative urn

- Depending on the evaluator's assessment, P%, the evaluator receives the following numbers of lottery tickets:
  - $N_I = \left[1 - (1 - P/100)^2\right] \times 10000$ tickets that are marked by I and numbered consecutively from 1 to N_I
  - $N_U = \left[1 - (P/100)^2\right] \times 10000$ tickets that are marked by U and numbered consecutively from 1 to N_U
- When the evaluator assesses P, the software displays the numbers N_I and N_U corresponding to every value of P in a friendly format
- The payoff of the evaluator depends on the outcome of the lottery as follows:

  i. Selection of the letter:
     * If the ball was drawn from the informative urn, letter I is selected
     * If the ball was drawn from the uninformative urn, letter U is selected

  ii. Selection of the number: The software extracts a random number between 1 and 10000 (in each period all numbers are equally likely to be extracted and extractions are independent across periods)

  iii. If the evaluator owns the ticket with the selected letter and the selected number, the evaluator wins 75 euro cents; otherwise, the evaluator wins 0 euro cents

[Screenshots are given to show how the evaluator can make an assessment using either the keyboard or the slider in the experiment's computer interface. See Figure 15.]

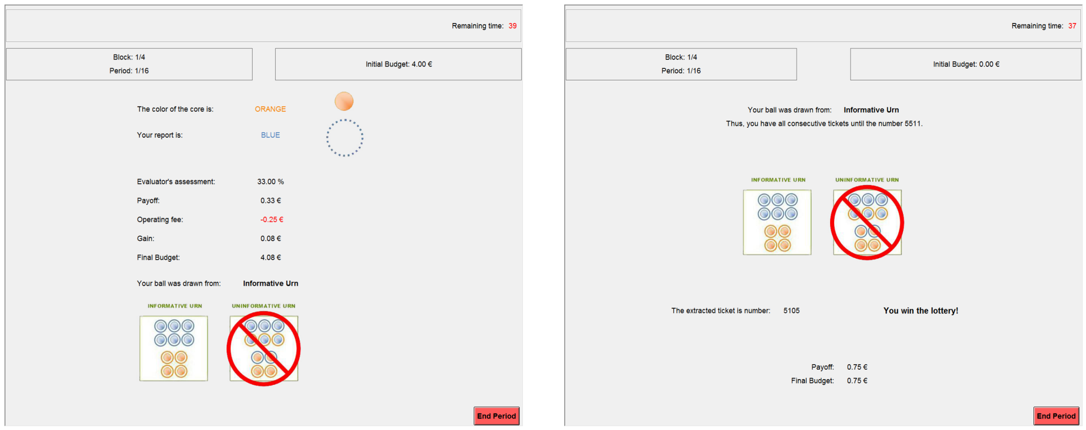

Figure 16: Feedback is given to the reporter(left) and to the evaluator (right) at the end of each period.

<u>Payoff of the Evaluator</u>

- Suppose that the ball was actually drawn from the informative urn
- Suppose that the software randomly extracts number 5105
- Given that the evaluator owns the winning ticket (number $I - 5105$), the evaluator wins 75 cents!
- Note that if the number extracted had been greater than 5511, the evaluator would have lost the lottery

[A SCREENSHOT IS GIVEN WITH THE EVALUATOR'S FEEDBACK ON PAYOFF. SEE FIGURE 16.]

<u>Evaluator Feedback</u>

Evaluator

At the end of each period, the evaluator receives the following feedback about the outcome of that period:

- The urn (informative or uninformative) from which the ball was drawn
- The evaluator's own payoff

Recall that the evaluator sees neither the reporter's plan nor the color of the shell

[A SCREENSHOT IS GIVEN TO SHOW THE HISTORICAL FEEDBACK GIVEN TO EVALUATORS IN BETWEEN EXPERIMENTAL PERIODS. SEE FIGURE 17.]

<u>[Your] [*Reporter*] Feedback</u>

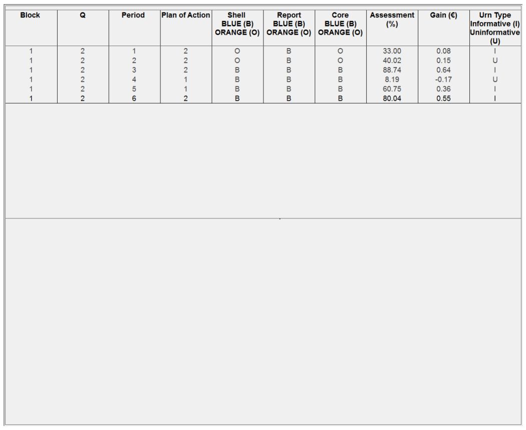

Block	Q	Period	Plan of Action	Shell BLUE (B) ORANGE (O)	Report BLUE (B) ORANGE (O)	Core BLUE (B) ORANGE (O)	Assessment (%)	Gain (€)	Urn Type Informative (I) Uninformative (U)
1	2	1	2	O	B	O	33.00	0.08	I
1	2	2	2	O	B	O	40.02	0.15	U
1	2	3	2	B	B	B	88.74	0.64	I
1	2	4	1	B	B	B	8.19	-0.17	U
1	2	5	1	B	B	B	60.75	0.36	I
1	2	6	2	B	B	B	80.04	0.55	I

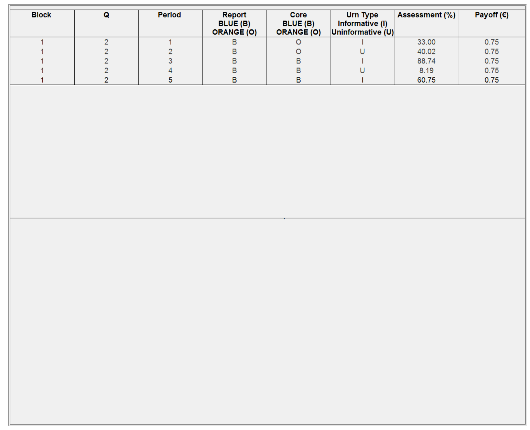

Block	Q	Period	Report BLUE (B) ORANGE (O)	Core BLUE (B) ORANGE (O)	Urn Type Informative (I) Uninformative (U)	Assessment (%)	Payoff (€)
1	2	1	B	O	I	33.00	0.75
1	2	2	B	O	U	40.02	0.75
1	2	3	B	B	I	88.74	0.75
1	2	4	B	B	U	8.19	0.75
1	2	5	B	B	I	60.75	0.75

Figure 17: In between periods, the reporter (left) and the evaluator (right) are reminded of important outcome variables for all past periods.

At the end of each period, [you] [*the reporter*] receive the following feedback about the outcome of that period:

- The color of the core of the drawn ball
- [Your] [*The reporter's*] own payoff
- The urn (informative or uninformative) from which the ball was drawn

[A SCREENSHOT IS SHOWN ILLUSTRATING THE HISTORICAL FEEDBACK GIVEN TO THE REPORTER IN BETWEEN EXPERIMENTAL PERIODS. SEE FIGURE 17.]

GAME CL ONLY

Evaluator Feedback

Computerized Evaluator

At the end of each period, the computerized evaluator receives feedback about the urn (informative or uninformative) from which the ball was actually drawn

Recall that the computerized evaluator sees neither your reporting plan nor the color of the shell

Transition Across Periods & Blocks

- At the end of each period the ball is returned to the urn from which it was drawn
- At the beginning of the following period a new coin flip is simulated and a new ball is drawn from the urn selected by the coin flip
- Urn selections and ball draws are therefore **independent** across periods
- You are allowed to take notes on scrap paper throughout the experiment

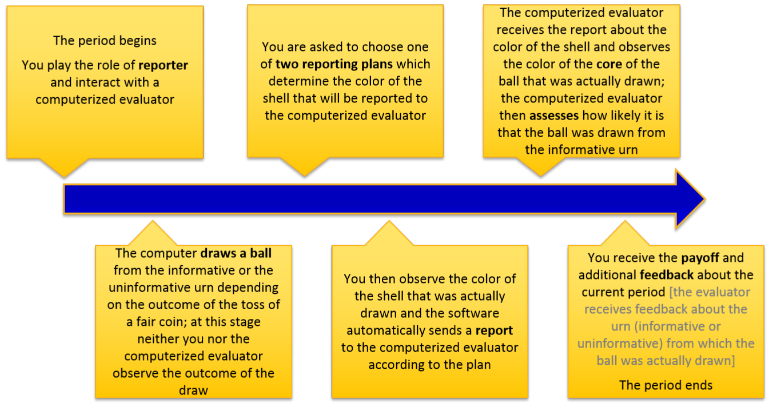

(a) Games CT, CU, and CL. Additional text for CL in gray.

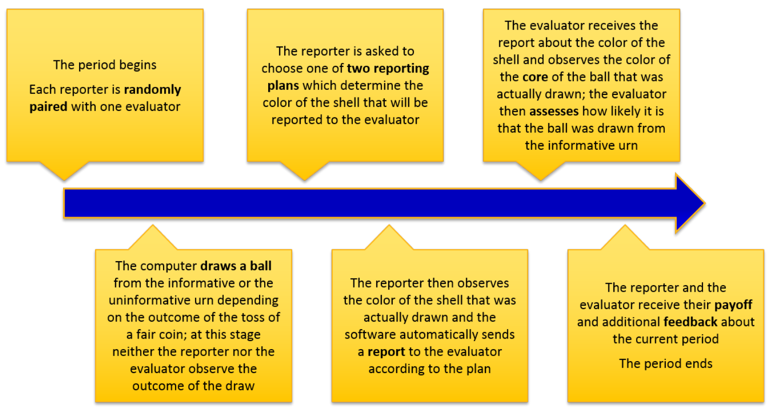

(b) Game HF.

Figure 18: Graphical summary of the experiment used in the experimental instructions.

- At the end of each block of periods you will have time to take notes about your experience during that block
- You are advised to go over your notes whenever you happen to play again a block of periods with the same number (Q) of balls with a **blue core**

Summary

At the beginning of each of the four (4) blocks of periods you are told the value of Q, the number of balls with **blue core** out of the total ten (10) balls that are contained in each of the two urns

For each of the sixteen (16) periods within each block, the timing is as follows: [see Figure 18.]