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

Debrah Meloso† Salvatore Nunnari‡ Marco Ottaviani§

October 6, 2018

Abstract

We experimentally study cheap talk by reporters motivated by their reputation for being well informed. Reputation is assessed by evaluators who see the report and the realized state of the world. In the laboratory, we manipulate the key drivers of misreporting incentives: uncertainty about the phenomenon to forecast and evaluators' beliefs. As predicted by theory, reporters are more likely to report truthfully when the state of the world is more uncertain and when evaluators conjecture that reporters always report truthfully. However, evaluators have difficulty learning reporters' strategies and tend to overreact to message accuracy, exacerbating reporters' incentives to misreport.

JEL classification: C91, D83

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 seminar audiences at the 'Political Economy: Theory Meets Experiments' Workshop at ETH Zurich, 2015 International Journal of Industrial Organization Conference, 2015 Conference of the Society for Experimental Finance, 2015 Conference of the Society for Economic Design, University of Alicante, Fondazione ENI Enrico Mattei, Toulouse Business School, International Symposium in Experimental Economics at NYU Abu Dhabi, and Fourth Women in Microstructure Meeting 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).

†Department of Finance and Economics, Toulouse Business School. E-mail: d.meloso@tbs-education.fr.

‡Department of Economics and IGIER, Bocconi University, Via Röntgen 1, 20136, Milan, Italy; and CEPR. E-mail: salvatore.nunnari@unibocconi.it.

§Department of Economics, BIDSA, and IGIER, Bocconi University, Via Röntgen 1, 20136, Milan, Italy; and CEPR. E-mail: marco.ottaviani@unibocconi.it.

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. 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.[1] In a different domain, after successfully calling the winner in forty nine U.S. states in the 2008 presidential election, Nate Silver was named one of The World's 100 Most Influential People by Time in 2009, he licensed his blog, FiveThirtyEight, for publication in the New York Times in 2010, and sold it to ESPN in 2013. The existence of such incentives might lead to believe that reputation motives and market forces 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 game proceeds as follows. 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.[2] Contrary to naïve intuition, this reputational incentive does not imply that the reporter always wants to truthfully report the private signal observed.

Variants of the reputational cheap talk game have been extensively used in the applied

[1]Alan Greenspan was chairman and president of Townsend-Greenspan & Co., Inc., and Lawrence Meyer was president of Lawrence H. Meyer and Associates.

[2]As explained in Section 2, this objective of the reporter can be derived as a reduced form payoff from a two-period model. While forecasters in the real-world examples above may have additional reasons to misreport (e.g., a stake in a decision a client takes following the forecaster's advice), the model we bring to the laboratory assumes away such reasons, thus isolating only the effect of reputation.

theory literature to model the strategic incentives for reputation management in settings including managerial decision-making (Holmström 1999; Scharfstein and Stein 1990; Prat 2005), recommendations by financial analysts (Trueman 1994; Graham 1999; Ottaviani and Sørensen 2006b), committee communication (Ottaviani and Sørensen 2001; Levy 2007; Visser and Swank 2007), strategic forecasting (Ehrbeck and Waldmann 1996; Lamont 2002; Ottaviani and Sørensen 2006a), professional election forecasting (Deb, Pai, and Said, forthcoming) and news reporting (Gentzkow and Shapiro 2006). This literature has focused on characterizing when we should expect the reputational concern for accuracy to induce the reporter to misrepresent the privately received signal. In turn, the evaluator, if aware of this incentive, should naturally discount the information reported. The lesson is that, in order to interpret forecasts, it is essential to understand forecasters' incentives.

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 this depend on uncertainty and evaluators' expectations? Are evaluators able to interpret forecasts or are they naïve? Our paper addresses these questions with controlled laboratory experiments. In the laboratory, we are able to measure and exogenously manipulate the information structure, the degree of uncertainty over the forecasted variable, and the expectations of evaluators.

A 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, and use its computerized implementation in our experiment. The scheme builds on the classic urn paradigm, which has been extensively used in the experimental literature since Anderson and Holt (1997) to test herding models à la Banerjee (1992) and Bikhchandani, Hirschleifer, and Welch (1992). 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 a crystal inner core inside the opaque 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.

The experiment proceeds as follows. A nested ball is drawn from an urn which is either informative or uninformative; the urn is covered so that neither the reporter nor the evaluator can observe from which urn the ball is drawn. The reporter privately observes the color of the shell and reports it (truthfully or not) to the evaluator. In addition to the report, the

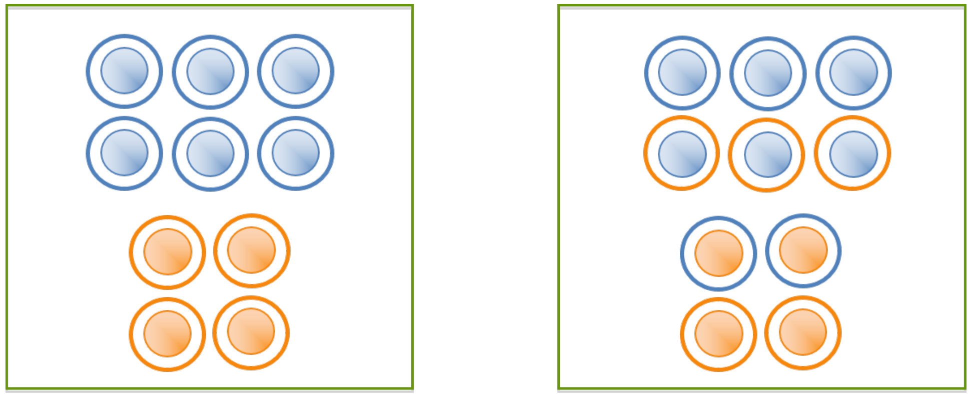

Figure 1: Urns with mildly unbalanced prior: $Q = 6$ balls have a blue core, giving $q = 6/10$. The left-hand panel represents the informative urn ($u = I$), the right-hand panel the uninformative urn ($u = U$).

evaluator observes the color of the core and, finally, assesses the probability that the ball was drawn from the informative urn. The payoff of the reporter is equal to the evaluator's assessment and the evaluator's payoff depends on the accuracy of his assessment.

Given our focus on understanding the behavior of reporters, we break down the game into its constituent components by controlling for strategic behavior and learning on the side of evaluators. 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, with the use of computerized evaluators, the evaluators' expectations.

We consider two values of q, which in our experiment corresponds to the fraction of balls with a blue core over the total number of balls in either urn. The values are chosen to generate different predictions about reporter behavior. Consider our *mildly unbalanced* prior of $q = 6/10$ (Figure 1). The informative urn in this case, is composed of six balls with a blue core and a blue shell and four balls with an orange core and an orange shell. 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 2), 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.

Theoretical predictions about reporters' incentives for truthfully reporting the observed shell depend on the prior belief q about the state of the world, and on evaluators' beliefs about the reporter's truthfulness. Hence, to further dissect reporters' strategic incentives, we vary how we control for evaluators' beliefs through four games: a game with *comput-*

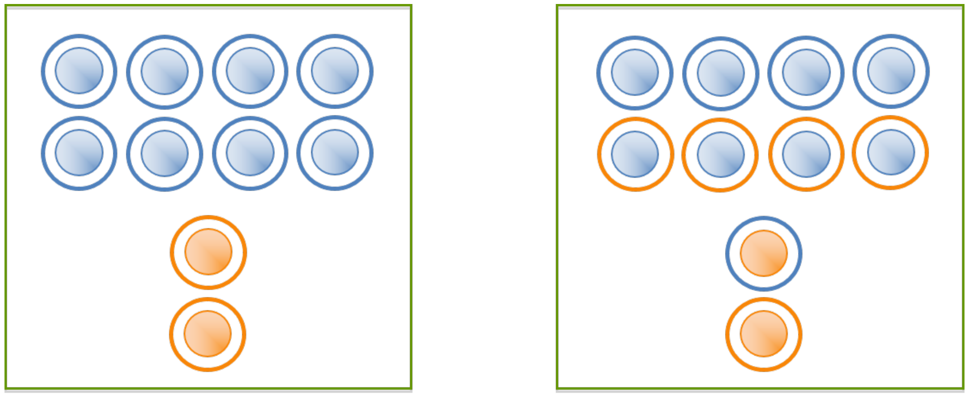

Figure 2: Urns with strongly unbalanced prior: $Q = 8$ balls have a blue core, giving $q = 8/10$. The left-hand panel represents the informative urn ($u = I$), the right-hand panel the uninformative urn ($u = U$).

erized evaluators programed to believe that all reporters always report *truthfully* (CT); a game with *computerized* evaluators programed to believe that the fraction of reporters who report truthfully is *uniformly* distributed (CU); a game with *computerized* evaluators whose beliefs are programed 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). According to the theory of reputational cheap talk, the 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.

According to the theory, the reporter should choose to be truthful only when the second force prevails. 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 contrarian information 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.

However, we find human evaluators' behavior to be incompatible with beliefs based on correct inference or Bayesian learning about reporters' strategies in different environments. In particular, we find that human evaluators incorrectly weigh reporter accuracy—an accurate report predicts the realized state by matching the color of the core—and inaccuracy—an inaccurate report does not match 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 behavior may reflect a higher willingness of evaluators to reward accuracy or punish inaccuracy at a cost, 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.

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 have negative consequences on forecasters' performance and the transmission of information.[3]

[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.

While there is a large experimental literature on the statistical herding models of Banerjee (1992) and Bikhchandani, Hirschleifer, and Welch (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 2017) focus on situations with multiple senders. Fehrler and Hughes (2018) and Mattozzi and Nakaguma (2017) 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, Morgenstern, and Raab (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 works is the information available to the evaluator before making his assessment.[5] 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.[6]

Our work relates to a broader experimental literature testing models of cheap talk (see Blume, Lai, and Lim 2017 for a comprehensive review). Differently from our setup, in these experiments, the sender cares about a decision taken by the receiver—rather than about his reputation—and the key driver of information transmission is the preference alignment between the sender and the receiver. Our work also relates to the experimental literature on "naïve advice", which explores the determinants and consequences of advice transmitted from senders who have limited information and internalize, at least partially, the receiver's well being (Schotter 2003; Schotter and Sopher 2007; Chaudhuri, Schotter, and Sopher 2009; Çelen, Kariv, and Schotter 2010). Finally, this paper relates to the experimental study of psychological game theory (Geanakoplos, Pearce, and Stacchetti 1989; Battigalli and

[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, Rabin, and Weizsäcker (2015). Anderson and Holt (2008) provide an excellent review.

[5]Similarly to these laboratory experiments, Meade and Stasavage (2008) and Hansen, McMahon, and Prat (2017) empirically explore the effect of transparency on deliberations within committees composed of career concerned experts using a natural experiment in the Federal Open Market Committee.

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

Dufwenberg 2009), since reporter utility depends on evaluator beliefs. In this literature, the hypothesized effect of beliefs on utility is mediated by emotions, which reduces 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 theoretical model and the testable hypotheses we take to the laboratory. Section 3 describes the experimental design and Section 4 presents the experimental results. Section 5 concludes. Appendix A derives in detail our testable hypotheses by analyzing theoretically the model; the material presented streamlines in a self-contained way results that have already appeared in the literature. Appendix B 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 our experimental design and analysis, but is also of independent interest. Proofs, supplementary empirical results, and full experimental instructions are relegated to a Supplementary Appendix.

2 Model and Testable Hypotheses

2.1 Model

We consider a simple 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 weakly 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 with conditional distribution

$$\Pr(S = B|c = b, u = I) = 1 - \Pr(S = B|c = o, u = I) = 1$$

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

and reporters with $u = U$ receive perfectly *uninformative* signals with

$$\Pr(S = B|c = o, u = U) = \Pr(S = B|c = b, u = U) = \frac{1}{2}.$$

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 $S = O$.

These strategies only differ when the reporter receives an Orange signal; thus, our theoretical analysis focuses on this event.

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, $E[p_{Rc}]$. 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.

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

[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.

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 the evaluators, the reporter is paid the expected value of her information. This justifies reporters in our setting being paid their expected future informativeness (the evaluator's assessment).

2.2 Testable Hypotheses

We now outline the testable hypotheses of the model that our laboratory experiment is specifically designed to investigate. Here we focus on the intuitive logic underlying the predictions; we present a detailed theoretical derivation in Appendix A.

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).*

As the fraction of balls with a blue core increases, the fraction of balls with a Blue shell in the informative urn increases while it remains unchanged in the uninformative urn (Figures 1 and 2 display such changes in urn composition). Thus, an Orange shell becomes a stronger indication that the urn is uninformative, giving the reporter a stronger incentive to misreport.

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

After receiving an Orange report, the evaluator is sure the shell of the ball is Orange, because reporters are not allowed to misreport after observing a Blue shell. This means that the evaluator's belief that the reporter is truthful, f, only affects p_{Bb} and p_{Bo} and, thus, the reporter's expected payoff from misreporting. On one hand, as f increases, the gain from misreporting when the evaluator observes a blue core, $p_{Bb} - p_{Ob}$, increases, because then 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 loss from misreporting when the evaluator observes an orange core, $p_{Oo} - p_{Bo}$, increases, because now a higher belief that the shell was truly Blue more strongly indicates that the urn was uninformative. 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, $p_{Oo} \geq p_{Bb} \geq p_{Bo} \geq p_{Ob}$).*
First, notice that matching shell and core are the strongest indication of a ball coming from an informative urn. $p_{Oo} \geq p_{Bb}$ because a matching orange report and core can only result if the shell is orange (given that misreporting is precluded when $S = B$); 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 if the report accurately predicts the state than if the report does not match 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).*
As remarked in the explanation of HP2, an Orange report perfectly reveals an observed Orange shell, 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. The assessment after a Blue report is thus moved further away from 1/2, the prior probability that the signal is informative.[9]

[9]When f is low, the evaluator is expecting to observe mostly blue reports—a blue report thus carries very little information, leaving the evaluator's belief about the informativeness of the signal mostly unchanged at 1/2, the prior belief that the signal is informative. A larger f makes a blue report more informative about the urn and thus spreads both possible assessments following a blue report further away from 1/2.

3 Experimental Design

The experiments 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. We ran 12 sessions. Subjects participated in only one session and maintained their role of reporter or evaluator throughout the whole session. Each session lasted around two hours and the average earnings (including a €5 show-up fee) were €36.05 for reporters and €39.66 for evaluators. The experiments were programmed using z-tree (Fischbacher 2007). Full experimental instructions for all treatments are reported in Supplementary Appendix E.

Information Structure. A first challenge in bringing our theoretical model to the laboratory is the complex information structure of a reputational cheap talk game. The solution we adopt is an innovative urn scheme, building on the setup introduced by Anderson and Holt (1997). At the beginning of each period, a ball is drawn from one of two urns, $u \in \{I, U\}$, with equal probability.[10] Each urn contains 10 nested balls. Each nested ball consists of an opaque outer *shell* and of a crystal inner *core*. The color of the outer shell is either Blue or Orange, $S \in \{B, O\}$, and corresponds to the signal privately observed by the reporter. The color of the inner core is either blue or orange, $c \in \{b, o\}$, and corresponds to the state of the world. The number of balls with a blue core is the same in both urns, $Q = 10q$. The urn informativeness determines the chance that the color of the shell matches the color of the core: the informative urn (I) contains only balls whose shell is the same color as the core; in the uninformative urn (U), instead, the color of the shell is the same as the color of the core only for half of the balls, and differs for the remaining half.

Counting Heuristic. The urn scheme with nested balls makes it possible to carry out the Bayesian updating discussed in Appendix A through a simple counting heuristic (Anderson and Holt 1996). To illustrate, consider an example with $q = 6/10$, that is, $Q = 6$. Figure 1 shows the two urns for this case. Suppose the reporter observes an orange shell. First, the reporter can compute the probability the evaluator observes a blue core. This is $q_O = \Pr(c = b|S = O) = 3/9 = 1/3$, since 3 out of the 9 balls with an orange shell have a blue core in the two urns combined. Second, the reporter can compute the evaluator's expected assessment if she reports $R = O$. Suppose the evaluator believes the reporter always reports truthfully: after observing an orange core, the evaluator should

[10] Whether the ball is drawn from urn I or urn U is determined by the computerized toss of a fair coin. Similarly, the random selection of a ball from the selected urn is done by the computer.

assess $p_{Oo} = \Pr(u = I|R = O, c = o) = 4/6 = 2/3$ since 4 of the 6 balls with an orange core and an orange shell are in urn I. When instead the evaluator observes a blue core, $p_{Ob} = \Pr(u = I|R = O, c = b) = 0$, since urn I only contains balls whose shells match the core. Thus, by reporting truthfully, the reporter obtains $(1/3)(0) + (2/3)(2/3) = 4/9$.

Reporters' Choices. Reporters submit their reports through the strategy method. Given the focus of the theory is on the incentives to misreport signal $S = O$, at the beginning of each period, reporters 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.

Evaluators' Choices. After observing the report and the core, evaluators are asked to assess the probability that the ball was drawn from the informative urn. We manipulate experimentally the nature of evaluators, who are human in one treatment and are impersonated by different computer algorithms in the other treatments.

Experimental Treatments. In the experiment, we exogenously manipulate q and, with the use of computerized evaluators, f. We consider two values of q: a mildly unbalanced prior, $q = 6/10$, and a strongly unbalanced prior, $q = 8/10$. The urns corresponding to these two priors are illustrated in Figures 1 and 2, respectively. We employ four different games varying the experimenters' degree of control on f:

- **CT** (*Computerized Trusting*). Evaluators are computerized and believe that all reporters always report truthfully: $f = 1$.

- **CU** (*Computerized Uniform*). Evaluators are computerized and believe that the fraction of reporters who report truthfully is uniformly distributed: $f \sim U(0, 1)$.[11]

[11] The theoretical analysis presented in Appendix A, treats f as the probability that the reporter truthfully report $R = O$ when observing $S = O$. Here, we treat it alternatively as a characteristic or type of the reporter, with a distribution over the entire population of reporters. The evaluator faces a sample of this population during a session and may use it to refine his belief about the true population distribution. The evaluators in CT and CU do not learn.

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.

- **CL** (*Computerized Learning*). Evaluators are computerized and initially believe that $f \sim U(0,1)$; the belief then evolves according to Bayes' rule depending on the outcome of past individual interactions with reporters, resulting in *generalized Beta learning.* Details of this novel learning model are given in Appendix B.

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

In each session, subjects played the same game in all periods—HF, CU or CT (that is, we used a *between-subject design* for f), but were confronted with both prior beliefs on the state—$q = 6/10$ and $q = 8/10$ (that is, we used a *within-subject design* for q). We run at least two sessions with the same game. Table 1 reports the details of the experimental design.

Periods and Blocks. Each session consisted of 4 blocks of 16 periods (for a total of 64 periods) to allow learning. Each reporter was randomly rematched with an evaluator at the beginning of each period. The value of q was fixed during a block but it changed from one block to the next, so that each value occurred in two non-consecutive blocks. We ran 6 sessions for each of the two orders: in order 6868, we started with $q = 6/10$; in order 8686 we started with $q = 8/10$. We use the term *first block* to indicate the block in which a given value of q was encountered for the first time, and *second block* refers to the block in which this same q was encountered for the second time.

Payoffs. At the beginning of each block of periods, the reporter received a budget of €4. In each period, the reporter paid an operating fee of €0.25 and obtained 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—that is, p_{Rc}. We compensated human evaluators in HF using a binarized scoring rule (Harrison, Martínez-Correa, and Swarthout 2013; Hossain and

Okui 2013).[12] Once the scoring rule is binarized, it is optimal (that is, incentive compatible) 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 Supplementary Appendix E).

Feedback. At the end of each period, each reporter received individual feedback about the outcome of the period that just elapsed. This feedback consisted of: 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, each human evaluator in HF received the following feedback at the end of a period: the type of urn from which the ball was drawn and the evaluator's payoff. We gave this feedback to reporters and evaluators to allow them to gain experience and learn how to play the game. Given that in each period's play evaluators observed neither the strategy of reporters nor the color of the shell, feedback given to evaluators was also meant to help them learn the strategy used by reporters.

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 2.2.[13] Throughout the Results section, whenever we state a result is significant, unless otherwise indicated, we refer to significance at the 1% level.

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

[12]To assist evaluators in their assessment, the software provided a slider and computed the payoff for different provisional assessments that were provided by the subjects.

[13]For games CT and CU, with exogenous evaluators' beliefs, HP1 is based on Proposition 2. For games HF and CL, HP1 is based on Proposition 4: the set of equilibria admits both truthful reporting and misreporting with $q = 6/10$, while it admits only misreporting when $q = 8/10$. HP2 is based on Proposition 3. All propositions are stated in Appendix A.

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.

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). In this and in the following tables, we focus on experienced subjects, that is, on decisions belonging to the second block for each treatment.[14] 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,

[14] Our findings are unchanged if we use all decisions. We discuss the effect of experience in Section 4.3 and we present summary statistics for first-block decisions in Supplementary Appendix D.

	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.

	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.

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 behavior observed in game HF to test hypotheses HP3 and HP4 from Section 2.2.[15]

4.2.1 Effect of Observing Different Reports and Cores

Tables 4 and 5 show summary statistics for human evaluators' assessments of the probability the urn is informative. Each row is for a different pair of observed report and observed core.

[15]Both HP3 and HP4 are based on Lemma 1. Our use of q as a proxy for f when we test HP4, is partly based on Proposition 4 and on the empirical findings on reporters' behavior discussed in Section 4.1.1. All formal theoretical results are in Appendix A.

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.[16] 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

[16] 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.

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 an accurate report and punishing to a larger extent an inaccurate report) 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 $q = 6/10$.

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

	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.

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.

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 Appendix B, 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

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.

neutral signal bears no information; a muddy signal assigns positive but distinct likelihood to 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 Appendix B 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 which 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)

	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 Bayesianly 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.

and an inaccurate report as a negative signal (that is a signal that the reporter he met this 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

	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.

evaluators believing (incorrectly, and similarly to human evaluators) that human reporters 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 which 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).

4.3 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. Table 13 reports estimates of the effect of experience on reporters' behavior as a function of the game and the treatment. Table 14 reports estimates of the effect of experience on human evaluators' behavior as a function of the treatment and the observed report-core pair.

	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	0.6	0.8	0.6	0.8	0.6	0.8	0.6	0.8
N	1488	1481	1504	1504	1504	1504	1536	1536

Table 13: Random effects GLS regressions. Each subject is a panel and periods are times within a panel. Standard errors in parentheses. Constant is omitted.

	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	0.6	0.6	0.6	0.6	0.8	0.8	0.8	0.8
N	768	390	114	216	1058	230	115	78

Table 14: 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.

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.

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).

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 design consists in the introduction of nested crystal 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, and more methodological, innovation in the experimental design consists in controlling for strategic behavior and learning on the side of the evaluator in a number of intermediate experiments in which we computerize the 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 Appendix D. 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 provide answers to the open questions about models of reputational cheap talk that motivated our study: *Do these models accurately predict behavior? Do reporters misreport available information to appear competent? How does this depend on uncertainty and evaluators' expectations? Are evaluators able to interpret forecasts or are they naïve?*

Empirically, 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, exacerbates the reporters' incentives for misreporting. 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.

Appendix A: Theoretical Analysis

Here we derive the theoretical foundations for the hypotheses stated in Section 2.2. We initially focus on the individual perspectives of evaluator and reporter, then obtain comparative statics predictions, and finally characterize the equilibria depending on the parameters. All proofs are relegated to the Supplementary Appendix C.

A.1 Evaluator's Assessments

Consider the evaluator, who acts as receiver in this sequential game of strategic communication. As long as the evaluator believes there is a positive chance the reporter is truthful, the evaluator learns from observing the report and the state of the world: if the report accurately predicts the state of the world, the evaluator should become more confident that the reporter is informed; instead, a report that does not match the state indicates that the reporter is uninformed. When the evaluator believes that the reporter is more likely to be truthful, the observed combination of R and c is more informative, resulting in a posterior further away from the balanced prior about informativeness.

Given report R, state of the world c, and belief f, the evaluator's assessment can be determined through Bayesian updating as

$$p_{Rc} = \Pr\left(u = I | R, c, f\right) = \frac{\frac{1}{2}\Pr\left(R|c, u = I, f\right)}{\frac{1}{2}\left[\Pr\left(R|c, u = I, f\right) + \Pr\left(R|c, u = U, f\right)\right]}. \tag{1}$$

Assessments p_{Oo} and p_{Ob} follow an Orange report, $R = O$, which can only occur when $S = O$, given that we do not allow the reporter to misreport $R = O$ when $S = B$. Therefore, in this case, deducing the signals informativeness from the report and the state is equivalent to deducing it from the signal and the state, $\Pr\left(u = I | R = O, c, f\right) = \Pr\left(u = I | S = O, c\right)$, regardless of the held belief about the reporter's truthfulness, f. An Orange report followed by the observation of a blue state implies a mismatch of signal and state. Since an informative signal is never different from the state, we have $p_{Ob} = 0$. On the other hand, observing an orange state implies that the signal matches the state. Since an informative signal is twice as likely to coincide with the state than an uninformative signal, and both signal types are ex-ante equally likely, $p_{Oo} = 2/3$.

Turning to assessments p_{Bo} and p_{Bb} following a Blue report, $R = B$, the evaluator should take into account that the reporter might have misreported Blue after observing an Orange signal, $S = O$. The evaluator's assessments after a Blue report thus crucially depend on the likelihood that the reporter is truthful, f. A belief that the reporter may misreport, $1 - f > 0$, should therefore dampen 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}. According to (1),

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

Intuitively, when the reporter is believed to always misreport, $f = 0$, the posterior assessments must be equal to the prior, $p_{Bb} = p_{Bo} = 1/2$. The evaluator's assessments based on comparing the match between the state and the report are most extreme and informative when the reporter is always truthful, $f = 1$. Assessment $p_{Bb} \in [1/2, 2/3]$ is convex and strictly increasing in f; $p_{Bo} \in [0, 1/2]$ is concave and strictly decreasing in f. Given that the quality of the signal is independent of the prior probability of the blue state, q, notice also that the evaluator's assessment is independent of q.

Lemma 1. *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 $f = 1$, $p_{Oo} > p_{Bb} = p_{Bo} > p_{Ob}$. If $f = 0$, $p_{Oo} = p_{Bb} > p_{Bo} = p_{Ob}$.*

A.2 Reporter's Expected Utility and Best Reply

The probabilities of the two possible assessments, p_{Rb} or p_{Ro}, following each report, R, are given by the reporter's posterior belief about the state of the world after observing signal $S = O$,

$$q_O \equiv \Pr(c = b | S = O) = \frac{\Pr(S{=}O|c{=}b)q}{\Pr(S{=}O|c{=}b)q + \Pr(S{=}O|c{=}o)(1-q)} = \frac{\left(\frac{1}{4}\right)q}{\left(\frac{1}{4}\right)q + \left(\frac{3}{4}\right)(1-q)}. \tag{2}$$

Note that $q_O \in [1/4, 1]$ is convex and strictly increasing in the prior belief, q. Thus, the reporter's expected utility when misreporting and sending $R = B$ is

$$EU(M) = (1-q_O)\, p_{Bo} + q_O p_{Bb} = \left(\frac{\frac{3}{4}(1-q)}{\frac{1}{4}q + \frac{3}{4}(1-q)}\right)\left(\frac{(1-f)}{\frac{1}{2} + (1-f)\frac{3}{2}}\right) + \left(\frac{\frac{1}{4}q}{\frac{1}{4}q + \frac{3}{4}(1-q)}\right)\left(\frac{1}{\frac{3}{2} + (1-f)\frac{1}{2}}\right).$$

Similarly, truthfully reporting $R = O$ gives expected utility

$$EU(T) = (1-q_O)\, p_{Oo} + q_O p_{Ob} = \left(\frac{\frac{3}{4}(1-q)}{\frac{1}{4}q + \frac{3}{4}(1-q)}\right)\frac{2}{3} + \left(\frac{\frac{1}{4}q}{\frac{1}{4}q + \frac{3}{4}(1-q)}\right)0.$$

Hence, the expected gain from misreporting rather than truth-telling after observing an Orange signal, is a function of the prior belief on the state, q, and of the evaluator's belief that the reporter is truthful, f:

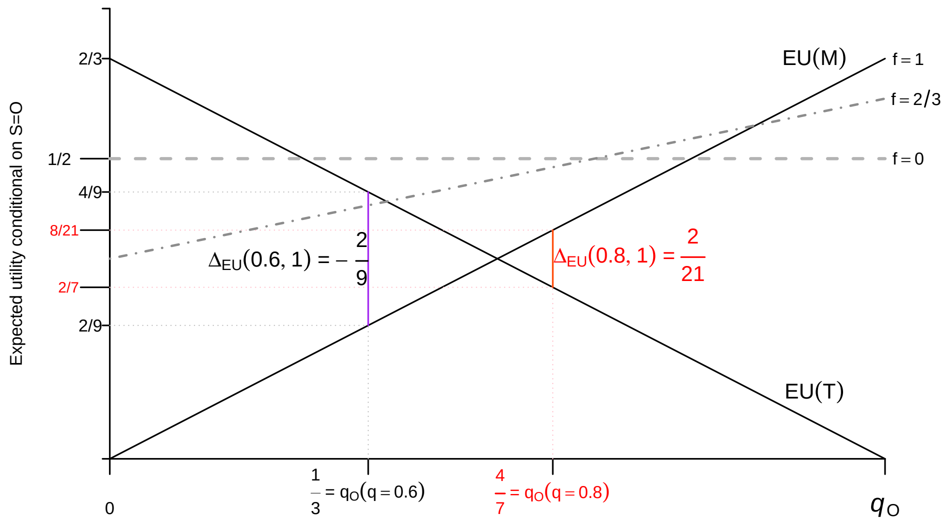

Figure 3: Reporter's expected utility from from *Misreporting* or *Truth-telling* after observing signal $S = O$, for a *mildly unbalanced* prior, $q = 6/10$, and a *strongly unbalanced* prior, $q = 8/10$. $EU(T)$ is equal for all values of f while $EU(M)$ varies with f. Three values of f are represented with the solid ($f = 1$), the dot-dashed ($f = 2/3$), and the dashed ($f = 0$) lines.

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

Since strategies M and T are identical when the observed signal is B, the reporter's choice is only affected by the expected gain from misreporting rather than reporting truthfully when the signal is O, weighted by the probability of this signal: $\Pr(S = O)\,\Delta_{EU}(q, f)$. The reporter is strictly better off misreporting if and only if $\Delta_{EU}(q, f)$ is strictly positive.

Proposition 1 (Reporter's Best Reply to Exogenous Beliefs)**.** *Consider an evaluator whose exogenous belief about the reporters' strategy is a random variable drawn from the distribution $G(f)$ with density $g(f)$. The reporter strictly prefers to misreport rather than to report truthfully if and only if $\int_0^1 \Delta_{EU}(q, f)g(f)df > 0$. If $G(f)$ is a degenerate distribution, the reporter strictly prefers to misreport rather than to report truthfully if and only if $q > \frac{4-f}{4(2-f)}$.*

Figure 3 illustrates the interaction between the reporter's posterior belief, q_O, and the evaluator's assessments, p_{Rc}, in the reporter's expected utility (equation (3)), and provides the intuition for Proposition 1. Consider the extreme case of a *perfectly trusting* evaluator, who is certain that $f = 1$. In Figure 3, the downward-sloping solid line represents the reporter's expected utility from a truthful report when her signal is O, $EU(T)$, as a function of her posterior belief, q_O, when the evaluator is perfectly trusting. This line crosses the vertical axis at 2/3, since when $q_O = 0$, $EU(T) = p_{Oo} = 2/3$ and for $q_O = 1$ we have $EU(T) = p_{Ob} = 0$. The upward-sloping solid line represents the reporter's expected utility from misreporting when her signal is O, $EU(M)$, as a function of her posterior belief, q_O. It crosses the vertical axis at 0, since when the evaluator is perfectly trusting and $q_O = 0$, $EU(M) = p_{Bo} = 0$ and for $q_O = 1$ we have $EU(M) = p_{Bb} = 2/3$. The two solid lines cross when the reporter's posterior belief about the state is $q_O = 1/2$. Therefore, the reporter maximizes her expected utility by truth-telling (resp. misreporting) as long as the posterior probability she assigns to state b is smaller (resp. larger) than 1/2. From equation (2) we know that a posterior belief $q_O = 1/2$ corresponds to a prior belief of $q = 3/4$. Application of Proposition 1 to the case of a trusting evaluator ($f = 1$), yields the same conclusion: the reporter prefers to misreport rather than to report truthfully if and only if $q > 3/4$.

In Figure 3 we determine the reporter's best reply for the values of q used in our experiment: a *mildly unbalanced* prior belief of $q = 6/10 < 3/4$, and a *strongly unbalanced* prior belief of $q = 8/10 > 3/4$. When the prior is mildly unbalanced, the posterior belief that the state is b is $q_O = 1/3 < 1/2$ and the reporter's expected gain from misreporting conditional on seeing signal O is $\Delta_{EU}(6/10, 1) = -2/9 < 0$ (unconditionally, taking into account the probability that the signal is $S = O$, the expected gain is $-1/10$), so the reporter is better off *Truth-telling.* When the prior is strongly unbalanced, the reporter's posterior belief after observing signal O equals $q_O = 4/7 > 1/2$, and her conditional gain from misreporting instead of truth-telling is $\Delta_{EU}(8/10, 1) = 2/21 > 0$ (unconditionally, 1/30), so the reporter prefers *Misreporting.*

Lighter, dashed lines in Figure 3 illustrate the role of the evaluator's belief, f. The expected utility from truth-telling when the signal is $S = O$, $EU(T)$, is clearly unaffected by f (only a truthful reporter reports $R = O$). Thus this expected utility is always represented by the same, downward-sloping, solid line. The expected utility from misreporting, $EU(M)$, changes as f changes. The larger the value of f, the larger the slope of the relation between expected utility and posterior belief. This is so because when the evaluator is more trusting, a blue report is more informative, and the difference between p_{Bb} and p_{Bo} is larger. Therefore, a change in the relative probabilities of these payoffs, q_O, has a stronger effect on the reporter's expected utility when the evaluator is more trusting.

In the extreme case where the evaluator is perfectly skeptical (certain that $f = 0$), a blue report carries no information, so the evaluator maintains her prior belief about the signal informativeness regardless of the report. Thus $p_{Bb} = p_{Bo} = 1/2$, and the reporter's posterior belief is irrelevant for her expected utility from misreporting. This case is represented by the flat dashed line with intercept $1/2$. As long as $q_O > 1/4$, this flat line lies above the expected utility from truth-telling, $EU(T)$ and, thus, the reporter prefers to misreport. Since $q \geq 1/2 \Rightarrow q_O \geq 1/4$, when the evaluator is perfectly skeptical, the reporter is always better off by misreporting (recall our assumption that $q > 1/2$). If the evaluator holds other, intermediate, degenerate beliefs about the reporter's truthfulness, $f \in (0,1)$, the line representing the reporter's expected utility from misreporting as a function of her posterior belief, has a slope between those of the perfectly trusting and perfectly skeptical cases considered so far. An example of such an f is given by the dot-dashed line in Figure 3: it represents $EU(M)$ when the evaluator believes $f = 2/3$ for sure.

The examples in Figure 3 were limited to evaluators with degenerate beliefs. We allow evaluators to hold distributional beliefs, for example because they are unsure about the reporter's behavior or they learn from experience. An example of such distributional beliefs, that corresponds to game CU in our experiment, is that of an *agnostic* evaluator who believes f is equally likely to take on any value on the interval $[0,1]$, $f \sim U[0,1]$. Applying Proposition 1, we obtain that when faced with an agnostic evaluator, the reporter strictly prefers to misreport rather than to report truthfully if and only if $q > \frac{2\ln(4)}{8\ln(4)-6\ln(3)} \approx 0.6163$.

A.3 Comparative Statics for Reporter's Incentive to Misreport

Equation (3) highlights the two main drivers of reporters' behavior in this simple game of reputational cheap talk: the common prior belief on the state of the world, q, and the evaluator's belief on the reporter's strategy, f. In this section, we analyze how each factor affects the reporter's propensity to misreport. While q is an exogenous variable which can vary in different environments, f is an endogenous variable and is determined in equilibrium. We derive the equilibrium f as a function of q in Section A.4. In this Section, we investigate how the reporter's best reply changes with an exogenous change to f. This is instructive of the reporter's incentive and directly relates to experimental treatments where we employ computerized evaluators to exogenously manipulate f.

Proposition 2 (Comparative Statics with respect to q)**.** *The reporter's incentive to misreport, i.e., the expected gain from misreporting with respect to reporting truthfully, is strictly increasing in q for any $f \in [0,1]$ and, thus, also for probabilistic evaluator's beliefs.*

Intuitively, when the reporter is more confident that the state of the world is b, she is more confident that a report $R = B$ will match the state and that a report $R = O$ will not, regardless of the signal she receives. If she misreports, the evaluator is more likely to assess

p_{Bb} and less likely to assess p_{Bo}. Since $p_{Bb} \geq p_{Bo}$ for any $f \in [0,1]$, this weakly increases her expected utility from misreporting. If she reports truthfully, the evaluator is more likely to assess p_{Ob} and less likely to assess p_{Oo}. Since $p_{Oo} > p_{Ob}$ for any $f \in [0,1]$, this strictly decreases her expected utility from reporting truthfully.

The effect of f on the reporter's expected gain from misreporting is ambiguous. While the gain that accrues if the state is o decreases with f, the gain that accrues if the state is b increases with f. Depending on q and on the evaluator's initial belief, increased trust (an increase in f) may either lower or increase the reporter's incentives to misreport. Nonetheless, for both values of q used in our experimental setup, incentives to misreport are lowest when the evaluator is strictly trusting.

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

Let us discuss the intuition behind Proposition 3 and, more generally, behind the effect of changes in evaluator beliefs on reporters' incentives. The evaluator's belief about the reporter's honesty only affects p_{Bb} and p_{Bo} and, thus, the expected assessment from misreporting. A larger f makes a Blue report more informative and thus, it brings both possible assessments following a Blue report further from 1/2, the prior belief that the signal is informative. On one hand, this increases the assessment when the core is blue and the report is accurate, implying that the gain from misreporting, $p_{Bb} - p_{Ob}$, increases with f. On the other, it decreases the assessment when the core is orange and the report is inaccurate, implying that also the loss from misreporting, $p_{Oo} - p_{Bo}$, increases with f. The loss from misreporting increases faster in f than the gain, which suggests that the net gain from misreporting decreases steadily in f and reaches a minimum at $f = 1$. However, when the reporter considers the expected gain from misreporting, she weighs gains and losses by their probability, q_O and $1 - q_O$, respectively. Thus, the speed at which expected gains and losses from misreporting increase with f, also depend on q_O, and thus, on q. For small values of q, expected losses indeed grow faster than expected gains, making $\Delta_{EU}(q, f)$ strictly decreasing in f. Above a first threshold for the value of q, $\Delta_{EU}(q, f)$ becomes concave in f, but the loss from misreporting grows sufficiently fast for f near 1, that $\Delta_{EU}(q, f)$ is still minimized at $f = 1$. Above a second threshold, $q = 9/10$, $\Delta_{EU}(q, f)$, still concave in f, is no longer minimized at $f = 1$, but at $f = 0$. A third threshold for q, makes $\Delta_{EU}(q, f)$ a strictly increasing function of f, since after weighting by (a very large) q_O, gains from misreporting grow faster with f than losses.[17]

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

A.4 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 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.[18]

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:*

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

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 4 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 4, 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 4 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

[18]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.

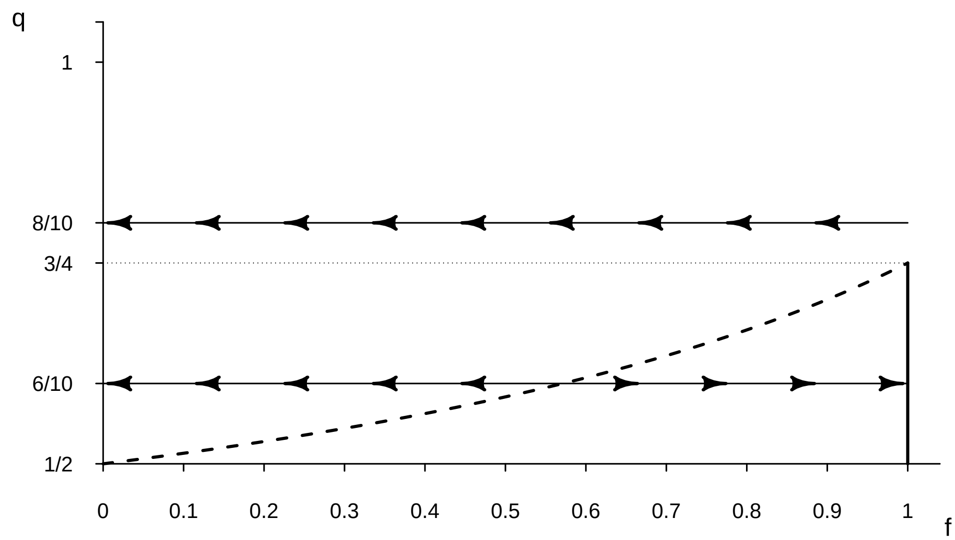

Figure 4: Equilibrium Set and Basin of Attraction.

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).

Appendix B: 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.[19]

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

$$g(x; v_1, v_2, d_1, \ldots, d_H, \delta_1, \ldots, \delta_H) = \frac{x^{v_1-1}(1-x)^{v_2-1}}{\mathrm{B}(v_1, v_2)} \times \\ \frac{(1-\delta_1 x)^{d_1-1} \ldots (1-\delta_H x)^{d_H-1}}{F_D^{(H)}(v_1; 1-d_1, \ldots, 1-d_H; v_1+v_2; \delta_1, \ldots, \delta_H)}, \tag{5}$$

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

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

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

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

[19]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).

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

$$l\left(\sigma_0, \ldots, \sigma_J|\,\theta\right) = \left[\pi_{0|0}^{\sigma_0} \times \ldots \pi_{J-1|0}^{\sigma_{J-1}} \times \pi_{J|1}^{\sigma_J}\right] \theta^{\sigma_J} \left(1-\theta\right)^{\sigma_0} \prod_{j=1}^{J-1} \left[1 - \left(1 - \frac{\pi_{j|1}}{\pi_{j|0}}\right)\theta\right]^{\sigma_j}. \qquad (6)$$

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 $h = 1, \ldots, J-1$.[20] The numerator of the above expression is given by

[20] 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):

$$
F_D^{(H)}\left(a; \mu_1, \ldots, \mu_H; c; \gamma_1, \ldots, \gamma_H\right) = \tfrac{1}{\mathrm{B}(a, c-a)} \int_0^1 t^{a-1}\left(1-t\right)^{c-a-1} \prod_{h=1}^{H}\left(1-\gamma_h t\right)^{-\mu_h} \mathrm{d}t. \tag{7}
$$

□

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*

$$
\varphi\left(x; k, \zeta_0, \zeta_1, \ldots, \zeta_{S,1}, \ldots, z_S\right) = \zeta_0 x^k \prod_{s=1}^{S}\left(1-\zeta_s x\right)^{z_s-1} \tag{8}
$$

is

$$
E\left[\varphi\left(x\right)\right] = \zeta_0 \tfrac{\Gamma(v_1+v_2)\Gamma(v_1+k)}{\Gamma(v_1)\Gamma(v_1+v_2+k)} \tfrac{F_D^{(H+S)}(v_1+k, 1-d_1, \ldots, 1-d_H, 1-z_1, \ldots, 1-z_S; v_1+v_2+k; \delta_1, \ldots, \delta_H, \zeta_1, \ldots, \zeta_H)}{F_D^{(H)}(v_1, 1-d_1, \ldots, 1-d_H; v_1+v_2; \delta_1, \ldots, \delta_H)}. \tag{9}
$$

Proof. Using (5) and (8) and collecting terms, we have

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

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 7, we replace the numerator and conclude that

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

□

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

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

which means these functions satisfy the assumptions of Proposition 6.

References

Anderson, Lisa and Charles Holt (1996). "Classroom Games: Understanding Bayes' Rule." *Journal of Economic Perspectives* 10.2, pp. 179–187.

— (1997). "Information Cascades in the Laboratory." *American Economic Review* 87.5, pp. 847–862.

— (2008). "Information Cascade Experiments." *Handbook of Experimental Economics Results*. Ed. by Charles Plott. Vol. 1. Elsevier, pp. 335–343.

Banerjee, Abhijit V. (1992). "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107.3, pp. 797–817.

Battigalli, Pierpaolo and Martin Dufwenberg (2009). "Dynamic Psychological Games." *Journal of Economic Theory* 144.1, pp. 1–35.

Bikhchandani, Sushil, David Hirschleifer, and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100.5, pp. 992–1026.

Blume, Andreas, Ernest K Lai, and Wooyoung Lim (2017). "Strategic Information Transmission: A Survey of Experiments and Theoretical Foundations." Working Paper. The University of Arizona.

Çelen, Boğaçhan and Shachar Kariv (2004). "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review* 94.3, pp. 484–498.

Çelen, Boğaçhan, Shachar Kariv, and Andrew Schotter (2010). "An Experimental Test of Advice and Social Learning." *Management Science* 56.10, pp. 1687–1701.

Charness, Gary and Martin Dufwenberg (2006). "Promises and Partnership." *Econometrica* 74.6, pp. 1579–1601.

Chaudhuri, Ananish, Andrew Schotter, and Barry 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, pp. 91–122.

Deb, Rahul, Mallesh M. Pai, and Maher Said (forthcoming). "Evaluating Strategic Forecasters." *American Economic Review.*

Dufwenberg, Martin and Uri Gneezy (2000). "Measuring Beliefs in an Experimental Lost Wallet Game." *Games and Economic Behavior* 30, pp. 163–182.

Ederer, Florian and Alexander Stremitzer (2017). "Promises and Expectations." *Games and Economic Behavior* 106, pp. 161–178.

Ehrbeck, Tilman and Robert Waldmann (1996). "Why Are Professional Forecasters Biased? Agency Versus Behavioral Explanations." *Quarterly Journal of Economics* 111.1, pp. 21–40.

Ellingsen, Tore, Magnus Johannesson, Sigve Tjøtta, and Gaute Torsvik (2010). "Testing Guilt Aversion." *Games and Economic Behavior* 68, pp. 95–107.

Eyster, Erik, Matthew Rabin, and Georg Weizsäcker (2015). "An Experiment on Social Mislearning." CEPR Discussion Paper 11020.

Fehrler, Sebastian and Niall Hughes (2018). "How Transparency Kills Information Aggregation: Theory and Experiments." *American Economic Journal: Microeconomics* 10.1, pp. 181–209.

Fischbacher, Urs (2007). "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10, pp. 171–178.

Geanakoplos, John, David Pearce, and Ennio Stacchetti (1989). "Psychological Games and Sequential Rationality." *Games and Economic Behavior* 1.1, pp. 60–79.

Gentzkow, Matthew and Jesse M. Shapiro (2006). "Media Bias and Reputation." *Journal of Political Economy* 114.2, pp. 280–316.

Goeree, Jacob K, Thomas R Palfrey, Brian W Rogers, and Richard D McKelvey (2007). "Self-Correcting Information Cascades." *Review of Economic Studies* 74.3, pp. 733–762.

Graham, John R. (1999). "Herding Among Investment Newsletters: Theory and Evidence." *Journal of Finance* 54.1, pp. 237–268.

Hansen, Stephen, Michael McMahon, and Andrea Prat (2017). "Transparency and Deliberation within the FOMC: A Computational Linguistics Approach." *Quarterly Journal of Economics* 133.2, pp. 801–870.

Harrison, Glenn W., Jimmy Martínez-Correa, and J. Todd Swarthout (2013). "Inducing Risk Neutral Preferences with Binary Lotteries: A Reconsideration." *Journal of Economic Behavior and Organization* 94, pp. 145–159.

Holmström, Bengt (1999). "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* 66.1, pp. 169–182.

Hossain, Tanjim and Ryo Okui (2013). "The Binarized Scoring Rule." *Review of Economic Studies* 80.3, pp. 984–1001.

Hung, Angela A and Charles R Plott (2001). "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review* 91.5, pp. 1508–1520.

Irlenbusch, Bernd and Dirk Sliwka (2006). "Career Concerns in a Simple Experimental Labour Market." *European Economic Review* 50.1, pp. 147–170.

Keane, Michael P. and David Runkle (1998). "Are Financial Analysts' Forecasts of Corporate Profits Rational?" *Journal of Political Economy* 106.4, pp. 768–805.

Khalmetski, Kiryl (2016). "Testing Guilt Aversion with an Exogenous Shift in Beliefs." *Games and Economic Behavior* 97, pp. 110–119.

Koch, Alexander K., Albrecht Morgenstern, and Philippe Raab (2009). “Career Concerns Incentives: An Experimental Test.” *Journal of Economic Behavior and Organization* 72.1, pp. 571–588.

Kübler, Dorothea and Georg Weizsäcker (2003). “Information Cascades in the Labor Market.” *Journal of Economics* 80.3, pp. 211–229.

Lamont, Owen A. (2002). “Macroeconomic Forecasts and Microeconomic Forecasters.” *Journal of Economic Behavior and Organization* 48.3, pp. 265–280.

Lauricella, Giuseppe (1893). “Sulle Funzioni Ipergeometriche a Più Variabili.” *Rendiconto del Circolo Matematico di Palermo* 7.1, pp. 111–158.

Levy, Gilat (2007). “Decicion Making in Committees: Transparency, Reputation, and Voting Rules.” *American Economic Review* 97.1, pp. 150–168.

Mattozzi, Andrea and Marcos Y. Nakaguma (2017). “Public Versus Secret Voting in Committees.” Working Paper. European University Institute.

Meade, Ellen E and David Stasavage (2008). “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve.” *Economic Journal* 118.528, pp. 695–717.

Ottaviani, Marco and Peter Norman Sørensen (2001). “Information Aggregation in Debate: Who Should Speak First?” *Journal of Public Economics* 81.3, pp. 393–421.

— (2006a). “Reputational Cheap Talk.” *RAND Journal of Economics* 37.1, pp. 155–175.

— (2006b). “The Strategy of Professional Forecasting.” *Journal of Financial Economics* 81.2, pp. 441–466.

Prat, Andrea (2005). “The Wrong Kind of Transparency.” *American Economic Review* 95.3, pp. 862–877.

Renes, Sanders and Bauke Visser (2018). “Committees of Experts in the Lab.” Working Paper. Erasmus University Rotterdam.

Scharfstein, David S. and Jeremy C. Stein (1990). “Herd Behavior and Investment.” *American Economic Review* 80.3, pp. 465–479.

Schotter, Andrew (2003). “Decision Making with Naive Advice.” *American Economic Review Papers & Proceedings* 93.2, pp. 196–201.

Schotter, Andrew and Barry Sopher (2007). “Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach.” *Games and Economic Behavior* 58.2, pp. 365–393.

Trueman, Brett (1994). “Analyst Forecasts and Herding Behavior.” *Review of Financial Studies* 7.1, pp. 97–124.

Visser, Bauke and Otto H. Swank (2007). “On Committees of Experts.” *Quarterly Journal of Economics* 122.1, pp. 337–372.

Supplementary Appendices

Supplementary Appendix C: Proofs

Proof of Proposition 1. The proposition follows directly from equation 3. We give further details of the derivation of this equation. First, recall that the reporter's strategies differ only if the signal is O, in which case the expected payoffs from misreporting or truth-telling, are, respectively

$$\begin{aligned} EU(M) &= (1-q_O)\, p_{Bo} + q_O p_{Bb}, \text{ and} \\ EU(T) &= (1-q_O)\, p_{Oo} + q_O p_{Ob}, \end{aligned}$$

where, by Bayes' rule

$$q_O = \Pr(c=b|S=O) = \frac{\Pr(S=O|c=b)q}{\Pr(S=O|c=b)q + \Pr(S=O|c=o)(1-q)} = \frac{\frac{1}{4}q}{\frac{1}{4}q+\frac{3}{4}(1-q)}.$$

Recall that $p_{Oo} = 2/3$ and $p_{Ob} = 0$, independent of the evaluator's beliefs, $G(f)$. In the case of a Blue report, the evaluator assessments are derived from

$$p_{Rc} = \Pr(u=I|R,c,f) = \frac{\frac{1}{2}\Pr(R|c,u=I,f)}{\frac{1}{2}[\Pr(R|c,u=I,f)+\Pr(R|c,u=U,f)]}.$$

with

$$\begin{aligned} & \Pr(R|c,u,f) \\ = \; & \Pr(R|S=O,c,u,f)\Pr(S=O|c,u,f) + \Pr(R|S=B,c,u,f)\Pr(S=B|c,u,f) \\ = \; & \Pr(R|S=O,f)\Pr(S=O|c,u) + \Pr(R|S=B,f)\Pr(S=B|c,u), \end{aligned}$$

where in the last equality we used the fact that the reporter does not know the realizations of c and u when sending the report and that the joint distribution of signal, state, and reporter type does not depend on the evaluator's beliefs.

Given the evaluator's beliefs and the reporter's available strategies, we know that $\Pr(R = B|S=B,f) = 1$, and $\Pr(R=B|S=O,f) = 1-f$, and therefore

$$\begin{aligned} \Pr(R=B|c=o,u=I,f) &= 1-f, \\ \Pr(R=B|c=b,u=I,f) &= 1, \\ \Pr(R=B|c=o,u=U,f) = \Pr(R=B|c=b,u=U,f) &= 1-\frac{1}{2}f, \end{aligned}$$

so that the evaluator's assessments after a Blue report are

$$\begin{aligned} p_{Bo} &= \tfrac{(1-f)}{\frac{1}{2}+\frac{3}{2}(1-f)} \\ p_{Bb} &= \tfrac{1}{\frac{3}{2}+\frac{1}{2}(1-f)}. \end{aligned}$$

Using the expressions for $p_{Oo}, p_{Ob}, p_{Bo}, p_{Bb}$, and q_O we obtain the gain from misreporting conditional on observing an Orange signal, $\Delta_{EU}(q,f)$, reported in equation (3). Ex-ante, before knowing the received signal, the reporter's expected gain from misreporting equals $\Pr(S=O)\Delta_{EU}(q,f)$, which preserves its sign.

If the evaluator holds distributional beliefs, $G(f)$, with density $g(f)$, the conditional expected gain from misreporting becomes

$$\tilde{\Delta}_{EU}(q, G(f)) = \int_0^1 \Delta_{EU}(q,f)\, g(f) df;$$

the reporter should misreport only if this expression is positive. If the evaluator holds point beliefs with $G(f)$ degenerate, the reporter should misreport if $\Pr(S=O)\Delta_{EU}(q,f) > 0$, i.e.,

$$\tfrac{3}{4}(1-q)\left(\tfrac{(1-f)}{\frac{1}{2}+(1-f)\frac{3}{2}} - \tfrac{2}{3}\right) + \tfrac{1}{4}q\left(\tfrac{1}{\frac{3}{2}+(1-f)\frac{1}{2}}\right) > 0,$$

which boils down to $q > \frac{4-f}{4(2-f)}$.

Proof of Proposition 2. Consider Δ_{EU} from equation (3) and recall that the reporter's ex-ante expected gain from misreporting is

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

The derivative $\frac{\partial[\Pr(S=O)\Delta_{EU}(q,f)]}{\partial q} = \frac{4-2f}{(4-f)(4-3f)} > 0$ given that $4-2f$, $4-f$, and $4-3f$ are positive for all $f \in [0,1]$.

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

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

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

Additionally, notice that for $q \in \left[\frac{9}{10}, \frac{27}{28}\right]$, the concave incentives for misreporting are minimized when $f=0$. Moreover, when $q > \frac{27}{28}$, 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, $q > \frac{4-f}{4(2-f)}$, 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.

Supplementary Appendix D: Summary Statistics By Experience

	$q = 0.6$			$q = 0.8$		
	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 15: 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 $q = 0.6$ and 9 decisions in the 1st Block with $q = 0.8$.

	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 16: Median assessment by the evaluators, by observed Report-Core and experience, $q = 0.6$. 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 $q = 0.8$ 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 17: Median assessment by the evaluators, by observed Report-Core and experience, $q = 0.8$. 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 $q = 0.8$ and 16 missing observations for the 2nd Block with either q.

Supplementary Appendix E: 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

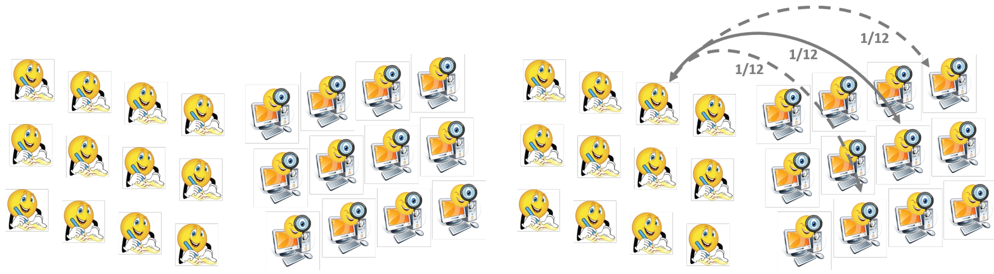

Figure 5: Game CL only: diagram to illustrate reporters' matching with computerized evaluators.

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
- 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 5.]

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**

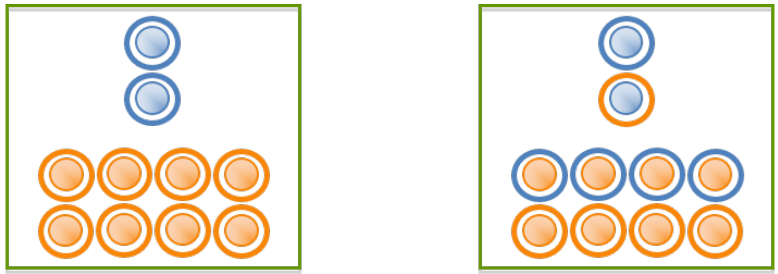

Figure 6: Informative (left) and uninformative (right) urn used as an example in the experimental instructions.

- 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

- The ball is drawn from one of **two urns** [Figure 6 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

The Informative Urn

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 6 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

The Uninformative Urn

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 6 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

<u>[Your Task] [*Task of the Reporter*]</u>

[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 7.]

<u>Implementation of Plan of Action</u>

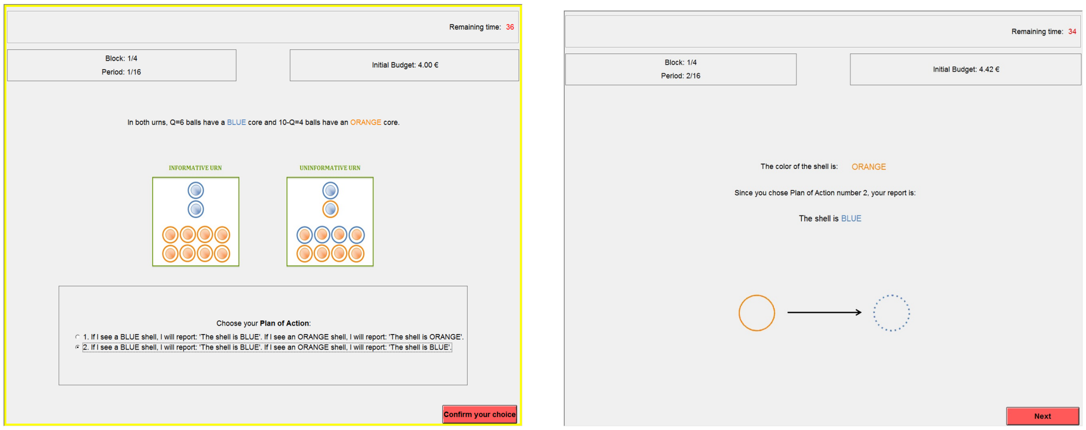

Figure 7: The reporter chooses plan of action (2) (left), a ball is drawn that has an Orange shell, so the reporter's report is Blue (right).

- 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 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 7.]

<u>Task of the [Computerized] Evaluator</u>

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

Figure 8: 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).

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 8.]

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

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

- 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 9.]

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 8.]

Payoff of the Evaluator

- 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 9.]

Evaluator Feedback

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 10.]

[Your] [*Reporter*] Feedback

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

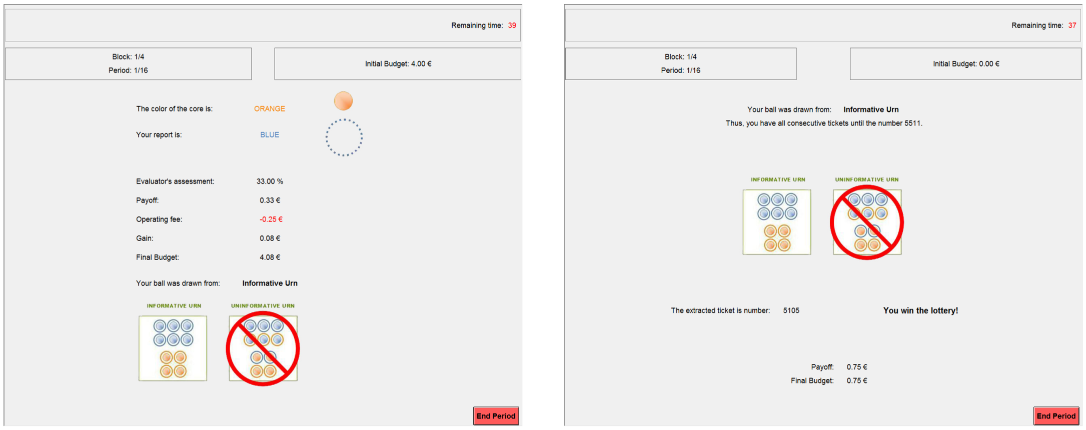

Figure 9: Feedback is given to the reporter(left) and to the evaluator (right) at the end of each period.

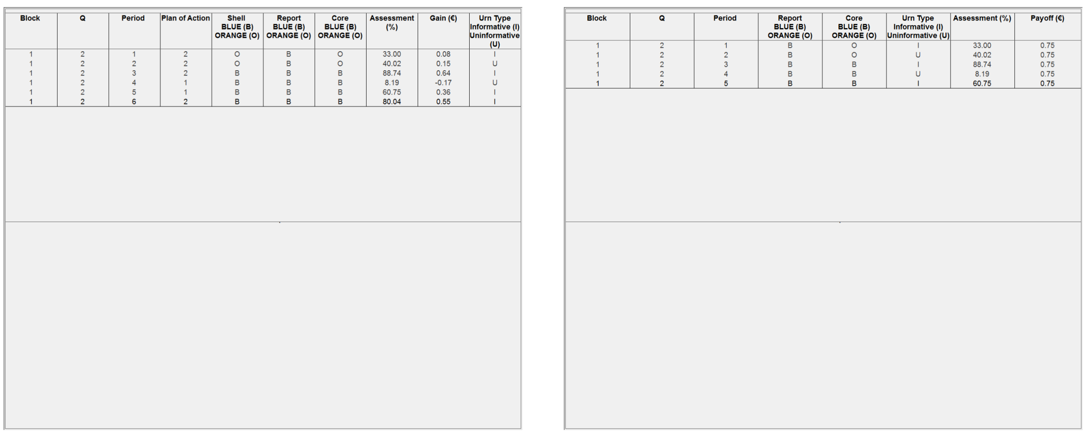

Block	Q	Period	Plan of Action	Shell BLUE (B) ORANGE (O)	Report BLUE (B) ORANGE (O)	Core BLUE (B) ORANGE (O)	Assessment (%)	Gain (€)	Urn Type Informative (I) Uninformative (U)
1	2	1	2	O	B	O	33.00	0.08	I
1	2	2	2	O	B	O	40.02	0.15	U
1	2	3	2	B	B	B	88.74	0.64	I
1	2	4	1	B	B	B	8.19	-0.17	U
1	2	5	1	B	B	B	60.75	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 10: In between periods, the reporter (left) and the evaluator (right) are reminded of important outcome variables for all past periods.

[A SCREENSHOT IS SHOWN ILLUSTRATING THE HISTORICAL FEEDBACK GIVEN TO THE REPORTER IN BETWEEN EXPERIMENTAL PERIODS. SEE FIGURE 10.]

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
- 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 11.]

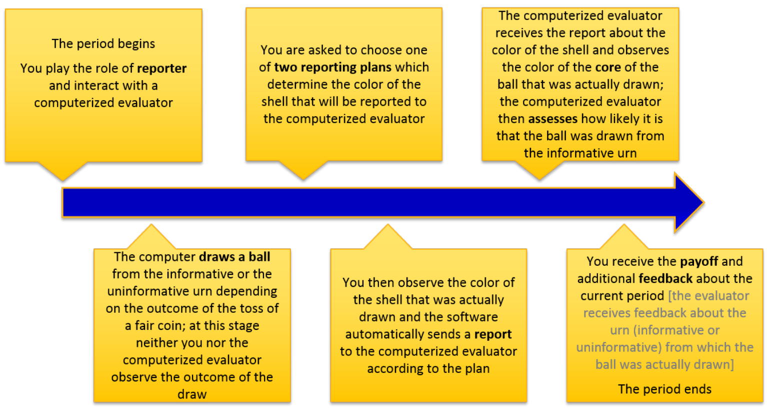

(a) Games CT, CU, and CL. Additional text for CL in gray.

(b) Game HF.

Figure 11: Graphical summary of the experiment used in the experimental instructions.