



European
University
Institute



ECO 2016/07
Department of Economics

Can there be a market for cheap-talk information?
Some experimental evidence

Antonio Cabrales, Francesco Feri, Piero Gottardi, and
Miguel A. Meléndez-Jiménez

European University Institute
Department of Economics

**Can there be a market for cheap-talk information?
Some experimental evidence**

Antonio Cabrales, Francesco Feri, Piero Gottardi, and
Miguel A. Meléndez-Jiménez

This text may be downloaded for personal research purposes only. Any additional reproduction for other purposes, whether in hard copy or electronically, requires the consent of the author(s), editor(s). If cited or quoted, reference should be made to the full name of the author(s), editor(s), the title, the working paper or other series, the year, and the publisher.

ISSN 1725-6704

© Antonio Cabrales, Francesco Feri, Piero Gottardi, and Miguel A. Meléndez-Jiménez, 2016

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Can there be a market for cheap-talk information? Some experimental evidence

Antonio Cabrales, Francesco Feri, Piero Gottardi
and Miguel A. Meléndez-Jiménez*

April 6, 2016

Abstract: This paper reports on experiments testing the viability of markets for cheap talk information. We find that these markets are fragile. The reasons are surprising given the previous experimental results on cheap-talk games. Our subjects provide low-quality information even when doing so does not increase their monetary payoff. We show that this is not because subjects play a different (babbling) equilibrium. By analyzing subjects' behavior in another game, we find that those adopting deceptive strategies tend to have envious or non-pro-social traits. The poor quality of the information transmitted leads to a collapse of information markets.

Keywords: Experiment, Cheap talk, Auction, Information Acquisition, Information Sale

JEL Classification: D83, C72, G14

* Antonio Cabrales, Dept. of Economics, University College London; Francesco Feri, Dept. of Economics, Royal Holloway; Piero Gottardi, Dept. of Economics, European University Institute; Miguel A. Meléndez-Jiménez, Dept. Teoría e Historia Económica, Universidad de Málaga. We thank Andreas Blume, Marco Casari, Vince Crawford, Matthew Embrey, Peter Eso, Antonio Guarino, Navin Kartik, Meg Meyer, Marco Ottaviani, Joel Sobel, Matthias Sutter, seminar participants at the London Behavioral and Experimental Group, University of East Anglia, Università Bocconi, EUI, QMUL, Oxford University, Sussex University, and WZB Berlin for comments. Miguel A. Meléndez-Jiménez acknowledges financial support from the Spanish Ministry of Economy and Competitiveness through Project ECO2014-52345-P and from the Regional Government of Andalusia through Project SEJ2011-8065. All errors are our own.

1. Introduction

Information transmission is a common occurrence in economic life. Transmission occurs even when only soft/unverifiable information can be transmitted and the incentives of senders and receivers are not well aligned. Potential competitors sometimes share information, such as when business or academic recruiters discuss the characteristics and ability of individuals they would both be interested in hiring or when corporate raiders discuss potential takeover targets. The drawback of sharing information in this way is that the conflict of interest makes the source unreliable. However, when acquiring information is costly, market participants can benefit by sharing this information, possibly in return for some payment, thereby creating a market for information. Clearly, the lower the rivalry of information or the higher its cost, the more likely soft information is to be transmitted or sold. The literature on cheap-talk games initiated by Crawford and Sobel (1982) establishes the conditions under which soft information can be transmitted in strategic settings.

This paper reports on a series of experiments testing the viability of a market for information. There are several reasons that such an experimental analysis of these issues could be interesting. On the one hand, previous experimental (and empirical) evidence shows that real subjects often tell the truth even when this goes against their self-interest, suggesting that they derive utility from not breaching a truth-telling norm. This could strengthen the case for a market for information. On the other hand, any game with a cheap-talk element features multiple equilibria, and hence, it is conceivable that real players have some difficulty in coordinating their play. In our experiments, we find that markets for information are fragile. The reasons for this outcome are rather different from those we had anticipated. Our experimental subjects provide low-quality information even in situations where doing so does not increase their monetary payoff. We show that this is not because subjects play a different (babbling) equilibrium. In fact, through the analysis of subjects' behavior in another game, we find that the individuals who adopt deceptive or non-informative strategies tend to have envious or non-pro-social traits. The poor quality of the information transmitted leads to a collapse of information markets.

With the objective of understanding this problem, we study a stylized model in which information can be acquired and then transmitted via non-verifiable reports, prior to trades in a market. More specifically, we investigate a market in which a single unit of an object is sold via a second-price auction in which several potential buyers can

participate. The object has a number of possible varieties, and each buyer cares about only one of them, chosen randomly (and independently). The seller has no utility for the object. None of the agents knows the true variety of the object, but they can learn it by incurring a cost. In addition, there is a market in which information can be exchanged. Any agent who acquired information can post a price at which he is willing to sell a non-verifiable report to other potential buyers.

Assuming that agents have self-interested preferences and do not derive utility from truth-telling, Cabrales and Gottardi (2014), henceforth *CG*, characterize the equilibria in which agents always send truthful reports when they learn they are not interested in the object, and are thus indifferent between lying and telling the truth, and otherwise send an uninformative message. When information costs are not excessive there is usually only one trader who acquires information and then sells a report to other agents. The report is sold at a price that is positive but sufficiently low such that all but one of the buyers who did not acquire it directly purchase the message.

We ran a series of laboratory experiments based on this game. The baseline game considered has two possible object types and three potential buyers. The cost parameter is set at a level such that in the equilibrium characterized by *CG*, one agent acquires information directly, one agent purchases a report and the last agent remains uninformed.

Our results are quite conclusive, in that the market for information does not appear to work well and is more fragile than what the theory predicts. We find there are far fewer purchases of reports from informed players and, in the last repetitions of the game, these purchases decline further and practically disappear. The reasons turn out to be rather surprising. As expected from the previous experimental literature on cheap-talk games, we observe numerous truthful messages from sellers of information who are interested in the object, though this typically increases the demand for the object and thus the price paid in the auction. This clearly favors the emergence of a well-functioning information market. This effect is, however, counteracted by another one working in the opposite direction: many sellers of information who find that they are not interested in the object either lie or send uninformative messages.

The fact that some uninterested informed players are not sending informative messages is a novel finding, to the best of our knowledge. To understand this behavior, it is important to note that, when a seller of information is not interested in the object and truthfully reports its properties, the agent buying the information with positive probability likes the object and, in that case, gains it at a low price. As a consequence, the receiver's

expected payoff is higher than that of the sender. If the seller is envious or non-pro-social, he may thus prefer to lie and thereby lower the payoff gained by the buyer of information. Of course, alternative explanations are possible. For example, a babbling equilibrium could prevail in the message game, with agents simply randomizing in their reports or sending the same report regardless of the object's type, as this would also clearly lead to the collapse of the market for information.

It is not easy to find sufficient evidence directly from the *CG* game to properly compare these alternative explanations. This is because there is little, and declining, activity in the market for information, which implies that few reports are actually sent, and hence limited information is available on the reports sent by sellers. For this reason, we conducted additional experiments for a sender-receiver game that has a similar equilibrium structure as the message component of our full game. For those additional experiments, we also elicited risk attitudes and social preferences from the subjects.

The results of this second set of experiments support the conjecture that agents' social preferences play an important role in explaining the surprising behavior we observed in agents' reporting. Individuals who are envious and non-pro-social are considerably less likely to tell the truth when they do not like the object. Even more surprisingly, these senders also lie when they like the object and the receiver does not. This reduces not only the sender's payoff but also that of the receiver. Thus, given selfish preferences, the behavior of these senders is not a best response to the observed behavior of the receivers. Moreover, we find no evidence that subjects are playing a babbling equilibrium.

Notably, the novelty of our findings can arise in part because of a subtle but important difference between the game we consider and the class of standard sender-receiver games examined in the experimental cheap-talk literature, following Crawford and Sobel (1982). In the usual experimental implementation of those games,¹ truth-telling (or, more precisely, separating) equilibria exist when the interests of the sender and the receiver are aligned and, more important from our perspective, their monetary payoffs coincide. In this case, both sender and receiver strictly gain from the sender's truth-telling behavior. In fact, the experimental evidence shows that in those situations, truth-telling behavior prevails. When the payoffs of the sender and the receiver conflict, truth-telling behavior is not consistent with equilibrium, although it is sometimes observed. In

¹ See, e.g., Dickhaut, McCabe and Mukherji (1995), Blume et al. (1998), or Cai and Wang (2006).

contrast, in our setup, truth-telling behavior is consistent with equilibrium when the sender's monetary payoff is lower than the receiver's, and the sender's payoff is not affected by his truth-telling behavior. In this case, we observe significant deviations from truth-telling.

We show that our results concerning the collapse of the information market are robust to variations in the design of the experiment, in particular to the consideration of the case in which the sender cannot participate in the auction and hence there is no conflict of interest between the sender and the receiver of reports.

1.1. Literature

First, we should mention the seminal work of Crawford and Sobel (1982) on strategic information transmission, which studies how the alignment of preferences between sender and receiver affects information transmission (Sobel (2013) reviews the vast theoretical literature following that paper). As noted above, with respect to that paper (and the subsequent literature), we consider a different and richer game structure that allows for some novel results. In particular, the amount of information available to agents is endogenously determined, and we allow payments to be required for the transmission of messages. Crucially, the alignment of interests between senders and receivers is not commonly known, as it depends on the preferences of the sender and the realized type of the object.²

The experimental literature on information transmission has concentrated primarily on analyzing sender-receiver games à la Crawford and Sobel (1982). A first series of papers (e.g., Dickhaut, McCabe and Mukherji (1995), Blume et al. (1998, 2001), and Kawagoe and Takizawa (1999)) demonstrates that when the interests of the sender and receiver are well aligned (the underlying game is one of common interest), play tends to converge to informative/separating equilibria, although other equilibria (babbling/pooling) exist. A more recent strand of the literature (see Sánchez-Pagés and Vorsatz (2007), Kawagoe and Takizawa (2005), Cai and Wang (2006), and Wang, Spezio and Camerer (2010)) finds more evidence of truth-telling than the most informative equilibrium in Crawford-Sobel would predict in games in which interests do not align

² There is also a relevant theoretical literature that studies information transmission when agents may have a preference for telling the truth (see Kartik, Ottaviani and Squintani (2007) and Bolton, Freixas and Shapiro (2007)).

well, which can be explained by a truth-telling norm. While in our experiments we also find some evidence of aversion to lying, we also observe a substantial amount of deception/misinformation even when lying does not increase the senders' payoff but reduces that of the receivers. This tends to be the case for subjects who display (independently measured) non-pro-social or envious preferences.

Gneezy (2005) also explores deception for message senders who sometimes do not (substantially) benefit materially from the deception, while the receiver is significantly harmed. We will discuss this work at greater length in the main body of the paper when we comment on our results. Here, it is worth noting that although there are relevant differences in the design of the experiment between his and our paper (for example, in Gneezy (2005), the experimenter does not inform the receivers of the payoffs of the sender, who is then unaware of the potential conflict of interest),³ the proportions of lying and truth-telling behavior are similar between the two papers.⁴ An important difference is that the senders in Gneezy (2005) always have a strict incentive to deceive others if they are self-interested. In our setup, in some situations, senders have no benefit from lying if they are self-interested.⁵

The sender's expectations regarding the receiver's response play an important role in assessing truthful behavior by the sender. Sutter (2009) examines this issue in the setup of the same games considered by Gneezy (2005), eliciting beliefs from senders regarding whether their recommendation will be followed. He shows that senders sometimes expect their recommendation not to be followed and then tell the truth (which in that case leads to the sender's preferred outcome). He then makes the argument that those "truths" should be called deceptive, which means that some behavior that appears to be altruistic when one does not consider senders' beliefs might actually be self-interested. This issue is of much less concern in our setup because the game is repeated a number of times. Additionally, we build on the methodology of Costa-Gomes, Crawford and Broseta (2001) to find the most likely strategy for every subject. We show that for no subject is

³ Moreover, Gneezy (2005), like us, explains his findings using individuals' social preferences, but he does not measure social preferences for the same individuals who play the information transmission game.

⁴ Charness and Dufwenberg (2006) explore theoretically and experimentally whether the results in Gneezy (2005) can be explained by a norm that induces guilt in senders if they "let-down" receivers (i.e., senders believe that they harm receivers relative to what the latter believe they will receive). They construct a different game from Gneezy (2005) that allows them to separate the role of social preferences from that of an aversion to disappointing receivers.

⁵ Hurkens and Kartik (2009), using the setup of Gneezy (2005), examine the relationship between lying aversion and social preferences. Their evidence supports the conjecture that the two traits are independent.

this strategy consistent with a *reverse* equilibrium, in which senders announce the strategy that is opposite to the truth and that is commonly understood.

Brandts and Charness (2003) establish experimentally that deception by a sender concerning her intended action leads the receiver to impose a costly punishment. Punishments would be lower when the sender takes the same action (which leads to an unequal payoff benefitting the sender) but she does not lie about it. That is, a violation of trust induces more negative feelings regarding the sender than simple envy. We also find in our games that misinformation is punished by receivers, by stopping to buy information from uninformative senders.

The paper is organized as follows. The game and its equilibria are presented in Section 2. Section 3 describes the experimental design and the results from the basic treatment. Section 4 presents the experimental evidence from the simplified sender-receiver game. Section 5 discusses the robustness of the results to the case in which the sender does not have a conflict of interest. Section 6 concludes.

2. The game and equilibria

There is one object for sale. The object can be any of K possible varieties, assumed to be equally likely *ex ante*. Let $v \in \{1, 2, \dots, K\}$ be the true variety of the object. There are N potential buyers. Each buyer $i \in \{1, \dots, N\}$ has positive utility for only one, randomly and independently chosen, variety θ_i that is his private information and denotes his *type*. The object is allocated to buyers via a second-price auction.

We assume that no trader knows the variety of the object for sale. Before the auction takes place, any buyer can learn the true variety of the object by paying a cost c . Any agent who paid this cost can then sell a cheap-talk message about the information he learned. The utility of buyer i is denoted by $\pi_i = I_{vi} - cI_{ci} - t_i$, where I_{vi} is an indicator variable that takes value 1 if buyer i gains the object and its true variety equals i 's type and 0 otherwise, I_{ci} is another indicator that takes value 1 if i acquires information directly and 0 otherwise. Finally t_i is the sum of the net monetary payments made by buyer i in the auction, to gain the object, and to the other traders, to sell/purchase information to/from them. Clearly, an important feature of this environment is that the different agents' preferences for the object are not always in conflict, but it is also not common knowledge whether this is the case.

To be more precise, the timing of the game is as follows:

1. Each buyer decides whether to pay c to acquire information about the object. This decision to acquire information, but not the content of the information itself, is observable by all agents.
2. Any buyer who acquired information can post a price p at which he is willing to sell a message about the type.⁶
3. Any buyer who did not acquire information in stage 1 decides whether to purchase a message from one of the agents selling information.
4. Every buyer who paid the cost c learns the true variety of the object and then sends a (common) report to all the buyers who purchased information from him.
5. A second-price auction takes place among all buyers to allocate the object.

The set of messages available to a seller of information is given by the set of possible varieties of the object plus one additional message. We will refer to this last message as the empty message, denoted by 0. Thus, the set of messages is as follows:

$$M = \{0, 1, 2, \dots, K\}.$$

Because the information sold in stage 3 of the game is unverifiable (a cheap-talk report), and babbling equilibria always exist in cheap-talk games (see Crawford and Sobel (1982)), the game considered has several equilibria. *CG* characterize the properties of equilibria in which a seller of information always sends a truthful message whenever telling the truth belongs to the agent's best response to the other players' strategies and beliefs. This appears to be a reasonable refinement, which can be formalized by assuming that agents have a very small cost of lying (as in Kartik (2009)), either from an intrinsic disutility or from fear of being caught and punished. As we mentioned in the introduction, this is also consistent with the existing experimental evidence.

In particular, *CG* show that the game described above has always an equilibrium in which any agent, when he is a seller of information, adopts the following message strategy:

$$m_i = \begin{cases} v & \text{if } v \neq \theta_i \\ 0 & \text{if } v = \theta_i \end{cases} \quad (1)$$

where m_i is the report issued by agent i . Therefore, buyer i is truthful about the variety of the object when this is different from his own type. However, when he likes the object, he sends the empty message 0. To understand why this strategy may be optimal, notice first that when buyer i learns that he is not interested in the object, he will not bid in the

⁶ The fact that the price of information is posted before the content of the information is learnt implies that the price posted has no signalling value.

auction; hence, his payoff (as specified in π_i) is independent of the message he sends, and he is thus willing to tell the truth. Consider next the situation in which buyer i learns that he likes the object. In this case, buyer i would gain by sending a message that deceives buyers and induces them to make the lowest bid. Because message strategy (1) conveys some information and hence the bids of other buyers depend on the message sent, this is achieved by sending the empty message. *CG* also focus on equilibria in which players use undominated strategies in the auction (we refer to them as truthful bidding strategies because each buyer makes a bid equal to his expected valuation for the object, conditional on his information) and on pure strategy equilibria.⁷

Proposition 1. For all $c \geq 0$, when $K \geq N - 2$ there exists a perfect Bayesian equilibrium of the game described above in which the sellers of information adopt the reporting strategy in equation (1) while buyers choose a truthful bidding strategy in the auction. Furthermore,

1. If $c > c^I \equiv \frac{1}{K} \frac{K-1}{K} + (N-2) \frac{1}{K} \left(\frac{K-1}{K}\right)^{N-1}$, no buyer acquires information; the object goes to a randomly chosen buyer, at a price $1/K$.
2. If $c \leq c^I$, one buyer acquires information and sells a report about it at a price $p = \min \left\{ \frac{1}{K} \left(\frac{K-1}{K}\right)^{N-1}, c \right\}$, at which all the other buyers except one purchase information; the object goes to a buyer who likes it, if such a buyer exists, at a price equal to $1/K$ (when either the seller of information or only one buyer of information likes the object), 0 (when neither the seller nor any buyer of information likes the object) and 1 otherwise.⁸

Thus, when information costs are low, information is acquired in equilibrium and transmitted via a message that is sometimes informative. Information is sold for a low enough price that all buyers except one purchase it. The market for information is typically a monopoly.⁹ Furthermore, the seller of information always obtains the object when he likes it; when he does not like it, the object goes to one of the buyers of

⁷ Both the experimental (see, e.g., Erev and Rapoport (1998) and references therein) and field evidence (Berry (1992), Ericson and Pakes (1995)) for entry games, which share a similar strategic structure, tend to favor pure strategy equilibria.

⁸ See *CG* for the proof of Proposition 1.

⁹ As shown in *CG*, when c is sufficiently low ($c \leq \frac{1}{K} \left(\frac{K-1}{K}\right)^{N-1}$) another equilibrium exists, with the same reporting and bidding strategy, with two sellers of information posting a zero price for information. This situation, however, does not arise for the parameter values considered in our experimental design.

information who likes it, if such a buyer exists, and otherwise to the buyer not purchasing information.

3. The experiment and the results

3.1 Design of the experiment

At the beginning of the experiment, subjects are divided into groups of three individuals. The subjects in any given group interact for 20 iterations of the game, and this feature is common information. Additionally, within each group of three subjects, each individual is randomly assigned a player position (1, 2 or 3) that remains fixed throughout the experiment.

Having fixed groups, and even fixed positions within group, was a design choice made for two main reasons. First, this is a fairly complicated experimental design from a cognitive perspective, and we wanted to maximize the probability that the players learned the best strategies to play the game. This becomes easier against a single group of players in a fixed position than against changing opponents and/or changing roles. Second, it increases substantially the number of independent observations on which to base our statistical analysis.

Obviously, the procedure also has disadvantages. The most important is that the repetition of the game creates new equilibria, and thus the theoretical benchmark is less clear. However, the main new equilibria of the game are those in which the amount of truth-telling increases because of reputational concerns. As we will see, the amount of truth-telling in our results is smaller even than that in the equilibrium described in Proposition 1. Additionally, the dynamic trends apparent in the data are easy to explain using simple learning heuristics, without resorting to complicated strategies in the repeated game.

In our main treatment (which we label *Base*), we implement the game described in Section 2, with parameters $N = 3$, $K = 2$ and $c = 20$. Players are informed that, in each round, they will have the opportunity to buy an object by bidding in an auction. The object can be either green or orange (its color is randomly drawn at the beginning of the round with equal probability). Similarly, each player has a randomly assigned color for the round (also green or orange, with equal probability). The object has a value of 200 ECUs (experimental currency units) for a player if it is of his assigned color and of 100 ECUs

otherwise. At the beginning of every round, each player is endowed with 250 ECUs and is informed of his assigned color but not of others' colors nor of the color of the object.

In every round, there are three stages. In the first stage, each player decides whether to pay 20 ECUs to learn the color of the object. This decision is made in sequence by the three players in any group, with an order randomly drawn at the beginning and then held fixed through the experiment, and with each player knowing the decisions of his predecessors,

In the second stage, we have the market for reports. Each player who paid to acquire the information in the previous stage (but does not yet know the color of the object) sets a price for his (future) report on the color of the object. The price can be any integer number of ECUs less than or equal to 20. Prices are set simultaneously by all sellers of information. Then, those players who remain uninformed decide which of the reports, if any, they want to buy at the indicated price (each one of those players can buy at most one report). These decisions are again made in sequence, with each player knowing the choices of his predecessors.

In the third and final stage, the reports are issued, and then the auction takes place. The content of the report can be "the object is orange", "the object is green" or "the object is orange or green". The three players simultaneously make their bids (from 0 to 250 ECUs) for the object. The bid can be any number of ECUs less than or equal to 250. The players know that the highest bidder will obtain the object, earning 200 ECUs if it is of his assigned color and 100 ECUs if it is not, and paying a price equal to the second-highest bid.¹⁰ The remaining bidders neither earn nor pay anything.

In the auction, we follow the strategy method: players who bought the information in the first stage are asked to choose their bid both in the event that the object is green and in the event that the object is orange. Players who bought a report are asked to choose their bid in the event that the report says that the object is green, in the event that it says the object is orange and in the event that it is uninformative. Players who neither bought the information nor a report are simply asked to make their bid. As a consequence, after players have made their strategy choices in the auction, those players that directly acquired the information in the first stage learn the color of the object and determine the content of their report. The realized color of the object and the chosen content of the

¹⁰ In the event of a tie, the acquirer of the object is randomly selected among the highest bidders. Note that in this case, the highest and second-highest bids coincide.

reports determine the bids of each player (according to their choices in the auction stage). Then, payoffs are realized.

At the end of each round, each player is informed of his payoff, the true color of the object, the bids made by each player, and the player who won the object and the price he paid.

Given the result stated in Proposition 1, the equilibrium prediction of the game considered in our base treatment is as follows. Regarding the information acquisition stage, one player acquires information directly, a second player buys a report at a price equal to 12.5 ECUs, and the remaining player stays uninformed. The informed player reports the true color of the object when he is not interested in the object and sends an empty message otherwise. In the auction, the informed player bids 200 ECUs if he is interested in the object and 100 ECUs otherwise; the acquirer of the report bids 200 (resp. 100) ECUs if he receives a report that indicates that he is interested (resp. not interested) in the object and 150 ECUs if he receives an uninformative report. The uninformed player bids 150 ECUs. In equilibrium, the expected payoff of the informed player is 17.5 ECUs, the expected payoff of the buyer of the report is 0, and the expected payoff of the uninformed player is 12.5 ECUs. At the end of the experiment, subjects were paid their payoffs from 4 randomly selected rounds at a conversion rate of 100 ECUs = 1 euro.

We ran three sessions of treatment *Base* at the laboratory of experimental economics of the University of Siena (LabSi) in November 2013 and March 2014. A total of 33 subjects participated in these sessions, providing a total of 11 groups. The subjects were recruited from the LabSi pool of human subjects, primarily consisting of undergraduate students from the University of Siena. No subject was allowed to participate in more than one session. After subjects had read the instructions, the instructions were read aloud by an experimental administrator. Throughout the experiment, we ensured anonymity and effective isolation of subjects to minimize any interpersonal influences that could stimulate cooperation. The average duration of sessions was 70 minutes (including the reading of instruction, excluding payment procedures). The experiment was computerized and conducted using the experimental software z-Tree (Fischbacher (2007)). The experimental instructions, translated into English, are reported in the online Appendix.

To understand the agents' behavior in the market for reports, we ran an additional treatment, denoted *Simplified*, that consists of a simple two-player sender-receiver game that resembles the structure of this market. We also ran further treatments, denoted

Uninterested, *Option*, and *Unint-Opt*, to test the robustness of the observed results with respect to some variations on the base game (focusing again primarily on the market for reports). Table 1 provides a summary of all our treatments.

Table 1. Experimental treatments¹¹

<i>Treatment</i>	<i>#sessions</i>	<i># groups</i>	<i># subjects</i>
<i>Base</i>	3	11 (4+4+3)	33
<i>Simplified</i>	3	10 (3+4+3)	40
<i>Uninterested</i>	3	12 (4+4+4)	48
<i>Option</i>	2	6 (3+3)	18
<i>Unint-Opt</i>	2	6 (3+3)	24

3.2 Results

In this section, we present the experimental results for the *Base* treatment and compare them with the theoretical predictions of the analysis in Section 2.¹²

Table 2 presents the results concerning the behavior of subjects in the auction (the third stage of the game). In the columns, we report the bids made in the first half (rounds 1 to 10) and second half (rounds 11 to 20) of the experiment. In the rows, subjects are differentiated according to their available information. We also report the bids made by subject in the equilibrium characterized in Proposition 1, referring to these as [Predictions].

In the first two rows, we report the behavior of the informed players, i.e., those who acquired information directly, specifically their average bid when the color of the object coincided with their assigned color (Color – Yes) and when it differed (Color – No). We observe a fairly clear learning pattern: when we move from the first to the second half of the experiment, the average bid when the color of the object does not coincide with the player’s assigned color increases from 73.15 to 102.77, very close to the theoretical prediction (100). A similar learning pattern occurs when the color of the object coincides with the player’s assigned color: the average bid increases from 125.37 to 176.97 (the theoretical prediction is 200).

¹¹ As in the case of the *Base* treatment, all sessions of the robustness treatments were run at LabSi (University of Siena): treatment *Simplified* in December 2014, treatment *Uninterested* in November 2013 and March 2014, treatment *Option* in March 2014 and treatment *Unint-Option* in March 2014 and June 2014.

¹² The results of treatment *Simplified* are reported in Section 4, those of treatment *Uninterested* in Section 5 and those of treatments *Option* and *Unint-Opt* in Section 6.

Table 2. Average bids by type of player and block of 10 rounds

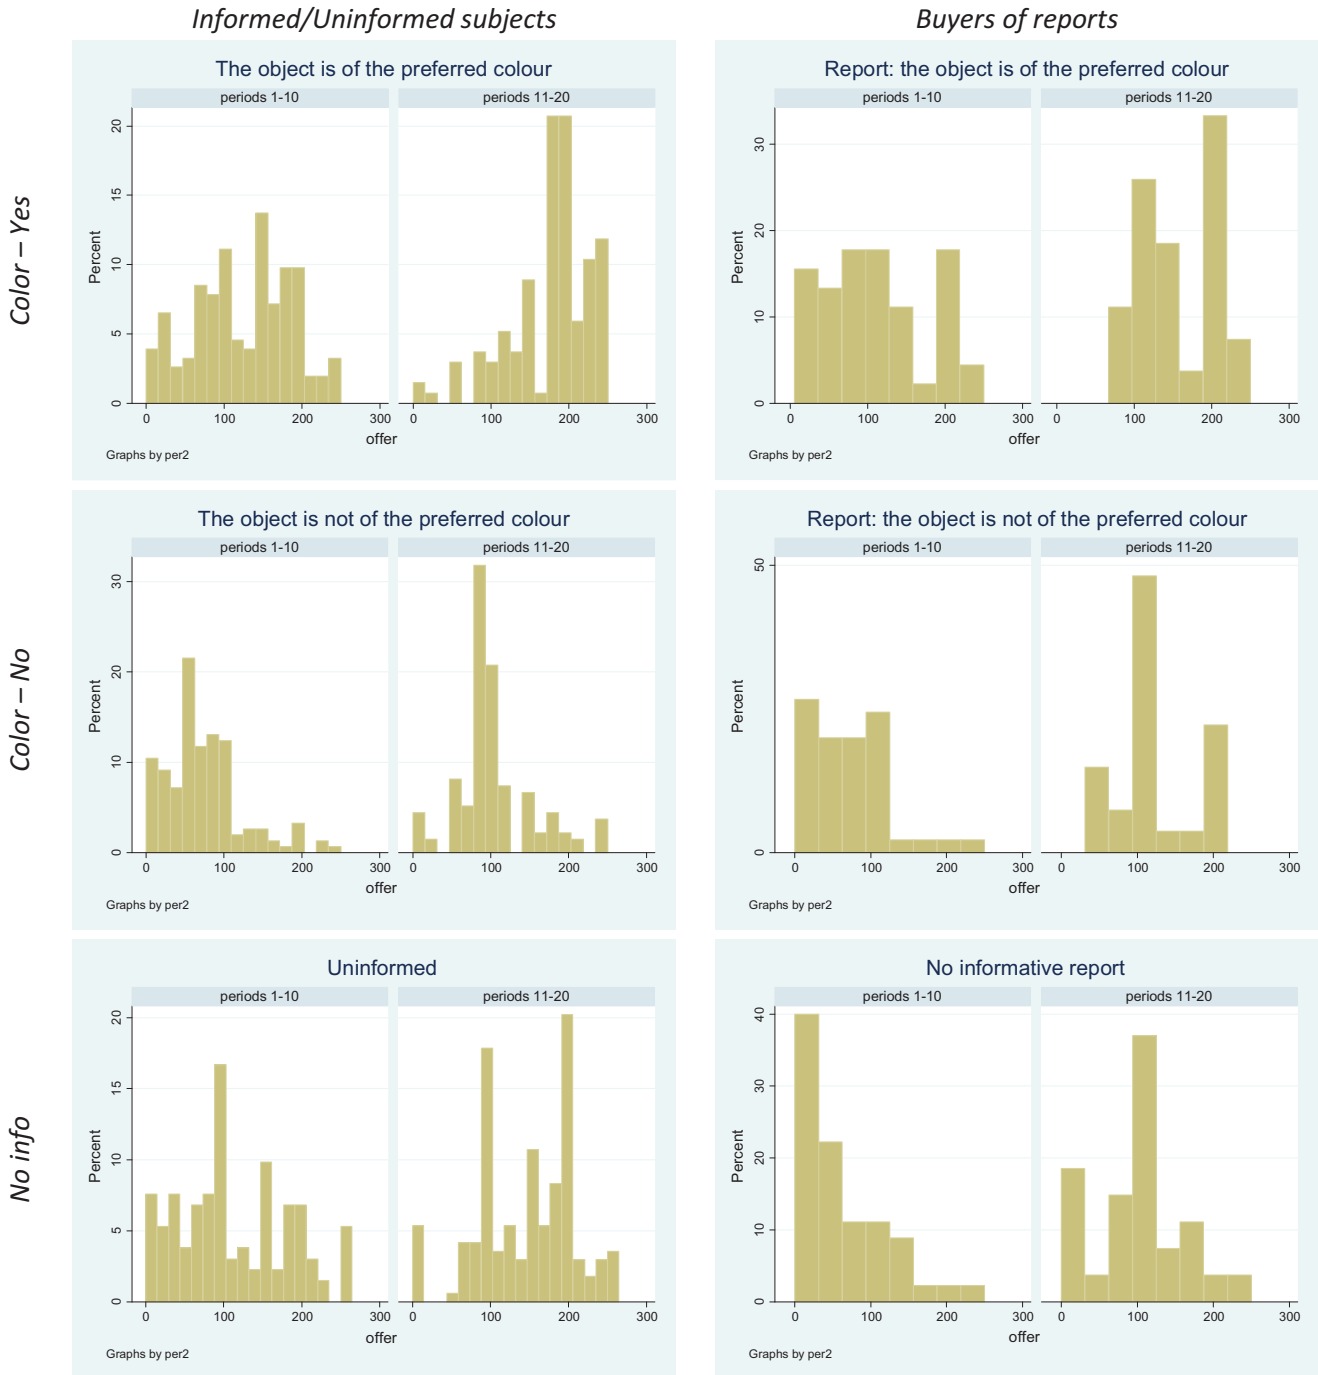
		<i>Rounds 1-10</i>	<i>Rounds 11-20</i>
<i>Informed players</i>	<i>Color - Yes</i> [Prediction: 200]	125.37 (62.52)	176.97 (53.41)
	<i>Color - No</i> [Prediction: 100]	73.15 (50.16)	102.77 (50.51)
<i>Uninformed players</i>	[Prediction: 150]	112.12 (70.35)	145.87 (62.17)
<i>Buyers of report</i>	Content: <i>Color - Yes</i> [Prediction: 200]	108.04 (67.37)	155.18 (49.95)
	Content: <i>Color - No</i> [Prediction: 100]	72.06 (50.91)	119.88 (50.91)
	Content: 0 [Prediction: 150]	61.86 (61.07)	97.81 (61.23)

(Std. Dev. in brackets)

The next row displays the average bid of the uninformed players (i.e., those who neither acquired information directly nor purchased a report): we observe an increase in the average bid from 112.12 to 145.87, quite close to the theoretical prediction (150).

The last rows display the average bid of indirectly informed players (buyers of reports) when the report states that the color of the object coincided with their assigned color (Color – Yes), when it did not coincide (Color – No) and when the report said “the object is orange or green” (we refer to this as the 0 report). It is interesting to compare the behavior in the first two cases with that of the informed players: we observe that the average bid responds to the content of the message received but considerably less than how the average bid of the informed players responds to the observation of the true color of the object. It equals 119.88 when “Color – No” and 155.15 when “Color – Yes”. Note that the theoretical predictions, when sellers of information adopt the message strategy in (1), truthfully reporting the information in those cases, are that bids should be 100 and 200, identical to those of the informed players. Hence, we can interpret this evidence as suggesting that receivers believe the report in those cases is only imperfectly informative. Finally, the average bid when the 0 report is received is (slightly) less than 100, the value of the object for a player in the worst-case scenario. The theoretical prediction is that in this case the report is not informative of the true color of the object but reveals that the seller of information is interested in gaining the object. This may explain why buyers of information are very conservative in their bidding (their expected value should be 150).

Figure 1. Distribution of bids by type of player and block of 10 rounds



To better understand the data, it is useful to also consider the distribution of bids across players. In Figure 1, we present the histograms of the bids in each of the cases considered. They show that in the second half of the experiment, the bids of informed players are indeed concentrated around the theoretical predictions. For the uninformed players, the evidence indicates that the average bid is close to 150, which is the result of the presence of two peaks (at 100 and 200), revealing rather curious behavior by bidders whose true value of the object is 100 and 200 with equal probability.

For buyers of reports, in the event that the content of the report is “Color – No”, the modal play is 100, the theoretical prediction if the message was fully believed, but there is also a significant peak at 200. Similarly, in the event “Color – Yes”, the modal play is at 200 but with a significant peak at 100.¹³ This suggests that a fraction of the players, when acting as buyers of information, do not trust that the message received is truthful in those cases. They still believe, however, that the message has some informational content because bids respond to the content of the message.

This evidence already demonstrates the importance of assessing the informational content of the reports sent by the sellers of information. In Table 3, we present the report sent, distinguishing the case in which the seller is interested in the object (the object is of his assigned color) and the case in which he is not interested.

Table 3. Content of the report by type of seller (interested in the object or uninterested)

<i>Seller</i>	<i>0 report</i>	<i>False report</i>	<i>Truthful report</i>	<i>Total</i>
<i>Uninterested</i>	8 (21.05)	8 (21.05)	22 (57.89)	38 (100)
<i>Interested</i>	10 (35.71)	5 (17.85)	13 (46.42)	28 (100)

(% over total row), [Pearson $\chi^2(2) = 1.7539$ Pr = 0.416], predictions in bold

In the equilibrium characterized in Proposition 1, sellers of information are always truthful when they cannot benefit by lying. Hence, uninterested sellers send a truthful report (i.e., reveal the color of the object) while interested sellers send a 0 report. In the experimental evidence, we observe a significant departure from such behavior: while the modal choice of uninterested sellers (57.98%) is indeed to reveal the true color, there is also a significant fraction of 0 (21.05%) and even false reports (21.05%). At the same time, interested sellers send truthful reports with a high frequency (46.42%). Importantly, note that the number of reports sent is only 66, considerably lower than the number of times the games were played (220, as there were 20 iterations per group and a total of 11 groups), and in the equilibrium considered, a report should be sent every time the game is played. This already clearly suggests that the market for information functioned far worse than the equilibrium predicted. Moreover, the relatively limited information concerning the reporting strategy followed by sellers, because of the small number of

¹³ When the 0 report is received, there is instead a single peak at 100.

reports sent, is another reason for the need to consider the additional treatment *Simplified*, which focuses on the behavior of message senders and receivers. (See Section 4 below.)

In Table 4, we present the frequencies of the different choices players could make regarding the acquisition of information: (i) directly acquire information in the first stage (*Inform*), (ii) buy a report in the second stage (*Buy rep*) and (iii) remain uninformed (*Uninf*). In presenting the data, we distinguish by mover (that is, by the order assigned to the player at the beginning of the session) and by block of 10 rounds. In the equilibrium considered, as explained at the end of Section 3.1, we should observe an equal number of each of these choices: in each round, one player in the group should directly acquire information, one player should buy a report, and one player should remain uninformed. However, the observed data follow a rather different pattern. In the first half of the experiment, the modal choice was to acquire the information directly (46.36%), followed by the choice of remaining uninformed (40%) and with a low frequency of purchases of reports (13.63%).

Table 4. Behavior in information markets – Absolute number of observations

	Rounds 1-10			Rounds 11-20		
	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>
<i>Mover 1</i>	41 (37.27) [31.06]	48 (43.63) [31.37]	21 (19.09) [46.66]	47 (42.72) [27.97]	47 (42.72) [34.81]	16 (14.54) [59.25]
<i>Mover 2</i>	46 (41.81) [34.84]	52 (47.27) [33.98]	12 (10.90) [26.66]	60 (54.54) [35.71]	42 (38.18) [31.11]	8 (7.27) [29.62]
<i>Mover 3</i>	45 (40.90) [34.09]	53 (48.18) [34.64]	12 (10.90) [26.66]	61 (55.45) [36.30]	46 (41.81) [34.07]	3 (2.72) [11.11]
<i>Total</i>	132 (40.00) [100]	153 (46.36) [100]	45 (13.63) [100]	168 (50.90) [100]	135 (40.90) [100]	27 (8.18) [100]

(% over total row), [% over total column]

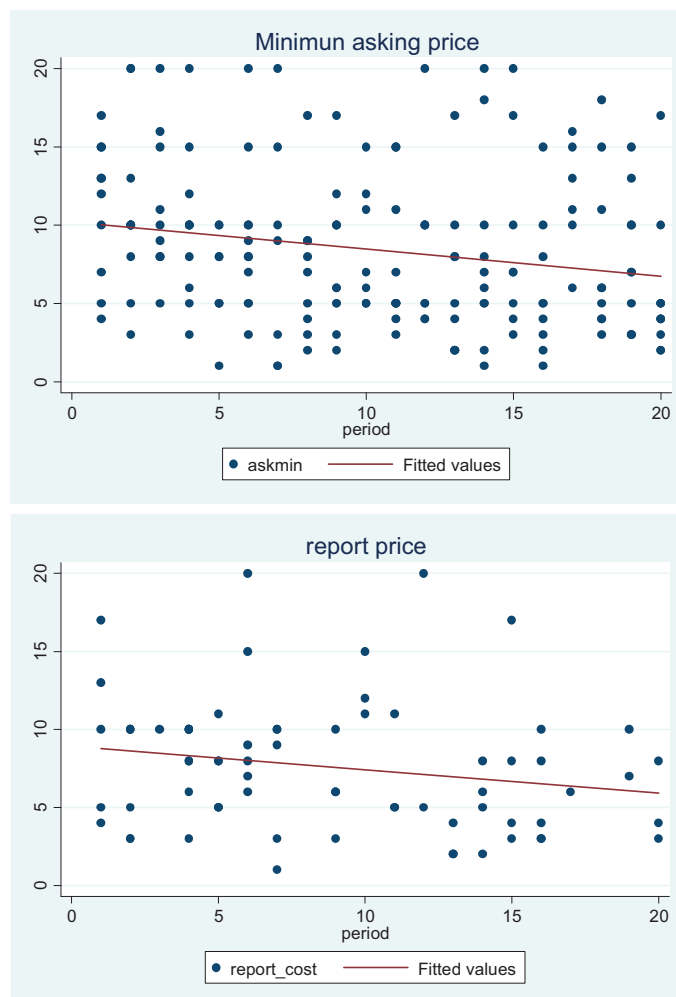
The frequency with which reports are purchased is further reduced in the second half of the experiment, where we find that the market for reports progressively collapses (the average frequency is 8.18%, with almost no activity in the last rounds). In this second half of the experiment, the modal choice is still to remain uninformed (50.9%), and players chose to acquire information directly with a frequency of 40.9%.

These results show that the market for reports is never very large and ultimately collapses. To understand the reasons for this collapse of the market for reports, one is

naturally led to relate it to the informational content of the reports sent. From Table 3, we observe that the number of truthful reports is close to 50%, in line with the theoretical prediction; hence, we can say that the informational content is not worse than predicted. However, the distribution of truthful reports between the case in which the sender is uninterested and that in which he is interested is quite different from the predicted distribution (the prediction is that all the truthful messages occur when the sender is uninterested). This has a significant effect on the benefit a receiver obtains from the purchase of information (a truthful message received when the sender is interested in the object means that the receiver will face aggressive bidding from the sender in the auction).

To understand the declining pattern of the activity in the market for information, we investigate the evolution of the prices posted for the sale of reports. In Figure 2, we plot in the top panel the prices asked by senders and in the bottom panel the prices accepted by receivers.

Figure 2. Market for reports: Minimum asked and accepted (report) prices



We see that the average posted price is lower than its level in the equilibrium considered (12.5) and declines over time. Recall, however, that, as noted in the previous paragraph, the benefit from purchasing information is also lower than predicted. We also find that the price is lower when there are two competing sellers of information¹⁴ and that more people buy information when prices decline.

On this basis, we proceed in Table 5 to estimate the determinants of the decision to buy a report. Because to be able to participate as a buyer in the market for reports, a subject must have chosen not to acquire information directly in the first stage, we jointly estimate the *selection equation* (the probability of not acquiring information in the first stage) and the *report equation* (the probability of buying a report in the second stage given that the player did not acquire information directly and that at least one of the other two players acquired information in the first stage).

Our model is quite similar to a Heckman probit estimation (that allows for the possibility of correlation between the selection and the report equations, measured by parameter *Rho* in Table 5). In addition, we need to account for the fact that, to be able to participate in the market for reports, another condition is required: at least one player in the group must have acquired the information in the first stage (i.e., there has to be a seller).¹⁵

The variables included in the selection equation (where we estimate the probability that a subject does not directly acquire the information) are *round*, a variable that represents the iteration of the game (from 1 to 20); *info_1*, a dummy that takes value 1 if the subject acts as mover 2 and the predecessor (mover 1) has directly acquired the

¹⁴ In particular, in 109 cases, we had only one seller (and two potential buyers). Of these cases, in 29 observations, one report was sold and in 6 cases two reports were sold. The average price of the report in these cases was 8.3 ECUs. Hence, we observed that, in 35 out of 109 cases (i.e., 32.11%), at least one report was sold. However, in 76 cases, we had two sellers (and one potential buyer), and in these cases, we observed 31 reports sold (i.e., 40.78%) with an average price of 6.6 ECUs.

¹⁵ Note that, to observe whether a player buys a report in the second stage, it is necessary that this player did not acquire the information in the first stage (otherwise, we could not observe the variable “buy a report”). If this were the only condition for the player to be a potential buyer in the market for reports, we could directly use the Heckman method (i.e., in the *selection equation*, we would estimate the probability of not acquiring information in the first stage, and in the *report equation*, we would estimate the probability of buying a report, correcting for self-selection by incorporating a transformation of the predicted individual probabilities of not acquiring information in the first stage as an additional explanatory variable). However, in our case, we need an additional condition for the player to be a potential buyer in the market for reports: at least one of the remaining players must have bought the information in the first stage (i.e., there has to be a seller). Hence, if we want to use all of the observations in our estimation (i.e., retain those in which no one buys the information in the first stage), we need to use a (maximum likelihood) modification of the Heckman method that includes a second condition for the *selection*: that at least one player bought the information in the first stage. The Stata program to perform this estimation is available from the authors upon request.

information and value 0 otherwise; *info_12*, a variable that takes value 0 if the subject is not mover 3 or if he is mover 3 and no predecessor has directly acquired the information, and value 1 (resp. 2) if he is mover 3 and 1 (2) predecessor(s) has (have) acquired directly the information; and *last*, a variable that takes value 1 (-1) if the last time the subject bought a report, this contained a true message (a false or uninformative message) and value 0 if the subject has not bought any report yet. The last variable is intended to capture the effects of previous experience with purchasing reports on subsequent information acquisitions.

Table 5. Market for reports

A) Model Estimation - Selection equation		B) Marginal effects - Selection equation	
<i>Round</i>	0.0121 (0.0086)	<i>round_m1</i>	0.0048 (0.0034)
<i>info_1</i>	0.6022*** (0.1481)	<i>round_m2</i>	0.0046 (0.0033)
<i>info_12</i>	0.1486* (0.0810)	<i>round_m3</i>	0.0048 (0.0034)
<i>Last</i>	-0.2468*** (0.0660)	<i>info_1</i>	0.2265*** (0.0509)
<i>Cons</i>	-0.1233 (0.1091)	<i>info_12</i>	0.0588* (0.0317)
	- Report equation	<i>last_m1</i>	-0.0984*** (0.0263)
<i>Round</i>	-0.0287** (0.0125)	<i>last_m2</i>	-0.0958*** (0.0259)
<i>Last</i>	0.1786** (0.0877)	<i>last_m3</i>	-0.0979*** (0.0264)
<i>Askmin</i>	-0.0473*** (0.0155)		- Report equation
<i>inf_tot</i>	0.3519*** (0.1356)	<i>Round</i>	-0.0065** (0.0028)
<i>Cons</i>	-0.8663*** (0.3093)	<i>Last</i>	0.0409** (0.0198)
<i>Rho</i>	0.9998*** (0.1783)	<i>Askmin</i>	-0.0108*** (0.0037)
		<i>inf_tot</i>	0.0807*** (0.0308)

[Number of obs = 660], [***, **, * significant at the 1%, 5%, and 10% level, respectively]

The variables included in the report equation (in which we estimate the probability that a subject buys a report) are *round*, *last*, *askmin*, which is the minimum price asked by a seller for a report, and *inf_tot*, which is the number of sellers of information (which can be 1 or 2).

In panel *A*) of Table 5, we present the model estimation, and in panel *B*), we report the marginal effects. Beginning with the latter, for the report equation, we provide the marginal effects of the variables *round*, *last*, *askmin* and *inf_tot*. We find that the marginal effect of *round* is negative and significant, showing the negative trend in the purchase of reports. The marginal effect of *last* is positive and significant, showing that the informational content of the previous report purchased matters: a negative experience in the market for reports (that is, receiving a false or uninformative report) decreases the future probability of buying a report. We also find that the marginal effect of *askmin* is negative and significant, indicating that the higher the (minimum) price asked for a report is, the lower the probability that the report is bought. Finally, the marginal effect of *inf_tot* is positive and significant: having two (rather than one) sellers of reports increases the probability that an individual buys a report.¹⁶

For the selection equation, we present the marginal effects of *round* and *last* for movers 1, 2 and 3 (*round_m1*, *round_m2*, *round_m3* and *last_m1*, *last_m2*, *last_m3*, respectively), and the marginal effects of *info_1* (for mover 2) and *info_12* (for mover 3). We find that the marginal effects of *info_1* and *info_12* are positive and significant (the first one only at the 10% level). This indicates that observing that mover 1 (or 1 and 2) already acquired information increases the probability that mover 2 (or mover 3) does not directly acquire information. There is now no significant time trend (*round*) for any mover. Note that the effect of *last* is negative and significant for all three movers, indicating that the higher the quality of the past reports, the lower the probability of not acquiring information directly (that is, the higher the probability of acquiring information directly). To understand this finding, note that in this case the variable *last* has two opposite effects on the decision to directly acquire information: on the one hand, previous experience of high quality reports provides an incentive to purchase a report instead of acquiring information directly. On the other hand, in groups where a significant share of high-quality reports has been observed in the past, it is more likely that reports are bought in the future (as shown by the marginal effect of *last* in the report equation), thus increasing the benefits of acquiring information directly in the first stage to then sell reports in the second stage. Our results indicate that the second effect dominates the first.

¹⁶ In this game, the only reason for buying a report is to be better informed than some other agent. If two agents have acquired the information directly, it is impossible to be better informed than anyone else, and hence there is no reason for buying a report. Thus, the fact that the marginal effect of *inf_tot* is positive and significant is in contrast to the theory.

Finally, in Table 6 we report the average payoff obtained by subjects, differentiating between the first and the second half of the experiment and with respect to the information available to subjects.

Table 6. Average payoffs

	<i>Rounds 1-10</i>	<i>Rounds 11-20</i>
<i>Informed players</i>	257.72 (51.41) [153]	244.62 (37.06) [135]
<i>Buyers of report</i>	245.17 (28.83) [45]	249.62 (39.91) [27]
<i>Uninformed players</i>	264.43 (47.26) [132]	246.83 (36.62) [168]

(Std. Dev.), [observations]

Although the performance of buyers of reports is fairly poor in the first half of the experiment (potentially due to the lower benefits gained from the purchase of information, relative to the price paid), towards the end of the experiment, the average payoffs of all types of players are very similar. This may be because the price of reports sold declines, as shown in Figure 2.

4. The simplified sender-receiver game

The market for reports and agents' decisions in this market constitute the more novel elements in the game considered. A proper understanding of subjects' behavior in such a market plays a key role in the analysis of the data obtained from the *Base* treatment. To this end, we run an additional treatment, denoted *Simplified*, which consists of a simple two-player sender-receiver game focused on these decisions. This game allows us to focus on a simpler environment in which the only decisions are the content of the report sent by the informed agent and the action chosen by the receiver of the message in response to it. We can thus obtain a much larger set of observations of behavior in this game (in the previous game, the observations on this were limited because of the small size of the information market). We can also characterize precisely the complete set of equilibria of the game (in the previous analysis, we concentrated on the most informative equilibrium). In the next section, we describe the design and report the results for this treatment.

4.1 The design

We design a two-player game that resembles the market for reports in the *Base treatment*, abstracting from decisions regarding direct acquisition of the information and from behavior in the auction. In this treatment, subjects were assigned to groups of four players and played 40 rounds of a sender-receiver game. In each round, subjects were randomly matched within their group, and in each pair, one subject was randomly assigned the role of player 1 (sender) with the other subject acting as player 2 (receiver).

We now describe the game.¹⁷ In each round, each player is randomly assigned a color (either black or white), with the assigned colors being i.i.d. There is also an i.i.d. random draw for the color of an object, which can be either black or white. Player 2 is only informed of his assigned color, whereas player 1 is informed of the colors assigned to both agents and of the color of the object. In the event that the assigned color of a player coincides with the color of the object, we say that the player is *interested* (otherwise, *uninterested*).

Next, player 1 sends a message (or report) regarding the color of the object to player 2. The message can be either “the color is white” or “the color is black”, i.e., player 1 can either send a truthful message or a false message. Player 2 observes the message sent by player 1 and makes a choice that can be either *left*, *center*, or *right*. The players’ payoffs depend on whether they are interested in the object and on the choice of player 2, as indicated in Table 7.

Table 7. Payoffs in the simplified game (in each cell: player 1’s payoff, player 2’s payoff)

Player 2’s choice	Left	Center	Right	Left	Center	Right
<i>Player 1 not interested</i>	20, 20	20, 60	20, 100	20, 160	20, 120	20, 50
<i>Player 1 interested</i>	20, 30	70, 90	120, 120	20, 60	70, 50	120, 40
	<i>Player 2 not interested</i>			<i>Player 2 interested</i>		

The payoffs resemble a situation in which player 1 (the informed player) makes a high bid in a hypothetical auction when he is interested in the object and a low bid otherwise. Player 2’s action left could be interpreted as a high bid in the hypothetical auction, action center as a medium bid, and action right as a low bid. Player 2 makes his

¹⁷ The experimental instructions, translated into English, are reported in the online Appendix.

choice not knowing whether he is interested in the object but observing player 1's message.

In the event that player 1 is not interested, his payoffs do not depend on player 2's choice, whereas if he is interested, he prefers that player 2 chooses *right* rather than *center* and *center* rather than *left*. The payoff for player 2 is related to her action as follows: if she is interested, then action *left* is preferred to *center*, which in turn is better than *right*. If she is not interested, the order is reversed.¹⁸

As we mentioned earlier, the specification of the payoffs in this game aims to capture the features of the market for reports in the *Base* treatment, but the game is sufficiently rich to allow for a variety of possible situations. We have, in fact, the case where player 1's payoffs are independent of the action of player 2, that is, when player 1 is not interested. We also have the case where preferences of player 1 are strictly aligned with those of player 2, when player 1 is interested and player 2 is not interested as well as the case where preferences are strictly misaligned, when players 1 and 2 are both interested.

The equilibrium analysis of the *Simplified* game is developed in the online Appendix, where we characterize the (weak) perfect Bayesian equilibria. We label *informative equilibrium* the equilibrium that is analogous to that of the *Base* treatment described in Proposition 1. In the *informative equilibrium*, player 1 sends a message containing the true color of the object if he is not interested in the object; however, when player 1 is interested, the message contains the opposite color of the color assigned to player 2. The equilibrium response of player 2 is to choose action *left* if the report says that she is interested (i.e., if it contains her assigned color) and action *right* otherwise. As in our *Base* treatment, in addition to the informative equilibrium, there is also a *babbling equilibrium* in which player 1's message is uninformative (i.e., the sender uses the same rule to determine the content of the message regardless of what he learned about his type and the type of the object) and player 2 chooses action center regardless of the content of the message received.¹⁹ Finally, there is another equilibrium (which we call for simplicity *extra equilibrium*) in which player 1's message contains the color of the object when the assigned colors of players 1 and 2 differ, and the message contains the opposite color of

¹⁸ Moreover, player 2's payoff is strictly higher when her assigned color differs from the assigned color of player 1 than when they coincide.

¹⁹ The case in which player 2 responds using action *left* on the out-of-equilibrium path (and action *center* on the equilibrium path) also constitutes a babbling (weak) perfect Bayesian equilibrium that is outcome equivalent to the babbling equilibrium described in the text. See the online Appendix for details.

the true color of the object when the assigned colors of players 1 and 2 coincide (i.e., player 1's message says that player 2 is interested in the object when player 1 is not interested, and it says that player 2 is not interested in the object when player 1 is interested). In this case, the equilibrium response of player 2 is to choose either action *left* or action *center* if the report says that she is interested (both actions provide the same expected payoff to player 2) and to choose action *right* otherwise.²⁰

In this *Simplified* treatment we do not have a 0 message, and hence the sender can only announce a type of the object, and the only alternative to the truthful message is to send a message that indicates the opposite of the true type of the object, which we define as deceptive behavior by the sender. This allows a cleaner identification of this type of behavior, as the analysis of deception and its determinants is one of the main purposes of this treatment. Note that in the informative equilibrium of the game in the *Simplified* treatment, the sender interested in the object is deceptive when the receiver is also interested in it. This behavior is possible and profitable because the sender knows the receiver's type. In the *Base* treatment, the sender does not know the receiver's type, and in the *informative* equilibrium characterized in Proposition 1, a sender interested in the object sends a message saying "the object is orange or green", which is not deceptive but uninformative. It is also useful to note that, in comparing the *informative* and the *extra* equilibrium in the game of the *Simplified* treatment, there is a higher level of deception in the latter, as the sender sometimes lies even when he is not interested in the object.

The equilibrium (expected) payoffs to players 1 and 2 are, respectively, 70 and 105 (in the *informative* equilibrium), 45 and 80 (in the *babbling* equilibrium), and 70 and 85 (in the *extra* equilibrium).

After the 40 rounds of play, we elicited the subjects' attitudes towards risk and social preferences. We used the risk test proposed by Charness and Gneezy (2010).²¹ Regarding social preferences, we used the approach proposed by Bartling et al. (2009), which allows us to identify pro-social and envious attitudes (see Table 10). This information allows us to investigate possible rationales for the behavior in the market for reports in terms of these variables.

²⁰ Of course, associated with each one of the three equilibria, there is a reverse equilibrium, in which the colors –white and black– are used inversely by player 1. See the online Appendix for details.

²¹ The subjects decide how much of their endowment (5 euros) to invest in a risky asset and how much to keep. They earn 2.5 times the amount invested if the asset is successful (prob. 0.5) and lose the amount invested otherwise.

4.2 Results

In Table 8, we report the behavior of senders. In particular, we report the frequency of true messages, distinguishing the cases in which the senders and the receivers are, respectively, interested or not interested in the object. According to the informative equilibrium, the frequency should be 1 except for the cell in which both players 1 and 2 are interested, in which case it should be 0. We observe that modal play coincides with the prediction of that equilibrium (the frequency of true messages is above 0.5 in the top cells and in the bottom-left cell of Table 8, and it is below 0.5 in the bottom-right cell), but there are significant deviations. These deviations are evident in the top cells, in which approximately one-third of the uninterested senders (who are expected to report the truth) lie, and in the bottom-right cell, in which approximately one-third of the interested senders (who are expected to lie) report the truth.

Table 8. Player 1’s behavior. Frequency of true messages

	All Rounds	Rounds 21-40	All Rounds	Rounds 21-40
<i>Player 1 not interested</i>	0.6315 (0.4834) [228]	0.6111 (0.4897) [108]	0.6497 (0.4782) [197]	0.6593 (0.4765) [91]
<i>Player 1 interested</i>	0.7419 (0.4387) [186]	0.8085 (0.3955) [94]	0.3439 (0.4762) [189]	0.3551 (0.4808) [107]
	<i>Player 2 not interested</i>		<i>Player 2 interested</i>	

(Std. Dev.), [number of observations]

In Table 9, we present the responses of the receivers with respect to the content of the message. Again, we observe that modal choices in each case correspond to the prescriptions of the informative equilibrium: the choice of *left* when the message says that the receiver is interested (*Color – Yes*, with a frequency slightly below 50%) and the choice of *right* when the message says that he is not interested (*Color – No*, with a frequency slightly above 60%). Note the significant use of action center (between 25% and 40% of the observations).

This departure in the observed behavior from the informative equilibrium can be associated with preferences that do not depend solely on monetary payoffs and/or with participants playing other equilibria. We now explore each of these possibilities in turn.

Table 9. Player 2's behavior. Absolute frequencies of choices

	Left		Center		Right	
	<i>All Rounds</i>	<i>Rounds 21-40</i>	<i>All Rounds</i>	<i>Rounds 21-40</i>	<i>All Rounds</i>	<i>Rounds 21-40</i>
<i>Color– No</i>	65 (13.68)	31 (12.81)	109 (22.95)	62 (25.62)	301 (63.37)	149 (61.57)
<i>Color – Yes</i>	161 (49.54)	76 (48.10)	110 (33.85)	61 (38.61)	54 (16.62)	21 (13.29)

% over total in the row in brackets, theoretical predictions in bold

4.3 Social preferences and risk attitudes

In what follows, we explore to what extent the behavior of senders is related to social preference considerations. To this end, we will use the responses of subjects to the dictator games (*a la* Bartling et al. (2009)) that we implemented at the end of each session. We now describe these games. Each subject had to make four decisions (one of them, randomly chosen, was paid). Each decision consists of a choice between *distribution 1* and *distribution 2*. The choice of a distribution determines a payoff for the player and a payoff for another player. These payoffs are shown in Table 10.

Table 10. Games for the elicitation of social preferences

<i>Game</i> <i>(All payoffs in euros)</i>	<i>Distribution 1</i> <i>self: other</i>	<i>Distribution 2</i> <i>self: other</i>
(I) Pro-sociality	2: 2	2: 1
(II) Costly pro-sociality	2: 2	3: 1
(III) Envy	2: 2	2: 4
(IV) Costly envy	2: 2	3: 5

According to the choices in these games, we can classify the subjects according to their pro-sociality and envy attitudes. Regarding pro-sociality (games I and II), those subjects choosing distribution 1 in game I and distribution 2 in game II are classified as *weakly pro-social* and those choosing distribution 1 in both games are classified as *strongly pro-social*. However, those choosing distribution 2 in both games are classified as *non-pro-social*.

Regarding envy (games III and IV), the subjects choosing distribution 1 in game III and distribution 2 in game IV are classified as *weakly envious*, while those choosing distribution 1 in both games are classified as *strongly envious*. In contrast, those choosing

distribution 2 in both games are classified as *non-envious*.²² In Table 11, we report the distribution of social preferences in our population of experimental subjects. We observe that 42.5% of the subjects are classified as envious and 30% of the subjects are classified as non-pro-social. We will explore the extent to which these attitudes are associated with the deviations from the informative equilibrium identified in Table 8.

Table 11. Distribution of social preferences

	<i>Weakly Envious</i>	<i>Strongly Envious</i>	<i>Non envious</i>	<i>Total</i>
<i>Weakly pro-social</i>	6 (25.00) [100]	5 (20.83) [45.45]	13 (54.16) [56.52]	24 (100) [60.00]
<i>Strongly pro-social</i>	0 (0) [0]	0 (0) [0]	4 (100) [17.39]	4 (100) [10.00]
<i>Non pro-social</i>	0 (0) [0]	6 (50.00) [54.54]	6 (50.00) [26.08]	12 (100) [30.00]
<i>Total</i>	6 (15.00) [100]	11 (27.50) [100]	23 (57.50) [100]	40 (100) [100]

(% over total row), [% over total column]

In Table 12 we perform a logit estimation of the probability that player 1 sends a true message. The explanatory variables are combinations of the social preference variables and dummies that determine whether the sender and the receiver are interested in the object. In particular, we define *Int1* (*Int2*) as a dummy that takes value 1 if player 1 (player 2) is interested in the object, i.e., if his assigned color and the color of the object coincides, and takes value 0 otherwise. We define $NoInt1 = 1 - Int1$ and $NoInt2 = 1 - Int2$. Similarly, we define *Prosoc* (*Env*) as a dummy that takes value 1 if player 1 is pro-social (envious), either weakly or strongly. We define $NoProsoc = 1 - Prosoc$ and $NoEnv = 1 - Env$. We also include in the regression the variables *Round* (from 1 to 40) and *Risk*, which corresponds to the choice of the subject in the risk test.

In panel *A*), we present the results of the model estimation, and in panel *B*) we report the marginal effects of envious and pro-social attitudes on the probability of sending a truthful report (measured at round 20 and the average risk level). In particular,

²² Note that a subject choosing distribution 2 in game I and distribution 1 in game II would be inconsistent in terms of pro-sociality. Similarly, a subject choosing distribution 2 in game III and distribution 1 in game IV would be inconsistent in terms of envy. We do not find any of these inconsistencies in our sample.

on the left-hand side of panel *B*), we present the marginal effect of envy (i.e., of $Env = 1$ vs. $Env = 0$). This is measured separately for pro-social and non-pro-social subjects and for the different combinations of preferences for the object of the sender and receiver (i.e., the four cells of Table 8). On the right-hand side of panel *B*), we present the marginal effect of pro-sociality (i.e., of $Prosoc = 1$ vs. $Prosoc = 0$). This is measured separately for envious and non-envious subjects, again differentiating among the four different combinations of preferences of the sender and receiver for the object.

The marginal effect of the risk variable measured at the average value of the remaining regressors (not reported in panel *B*) of Table 12) is -0.02665 and is significant at the 5% level. As the risk variable represents the number of euros (from 0 to 5) invested in the risky asset in the Charness and Gneezy (2010) test (see Footnote 21), this indicates that, the more risk averse an individual is, the more likely it is that he sends a true message. In particular, on average, each additional euro invested by an individual in the risky asset reduces the probability of sending the true message by 2.66%.²³

The first implication we draw from the marginal effects is that social preference attitudes are associated with deviations in the behavior of the sender from the informative equilibrium in the cases in which such equilibrium prescribes that senders report the truth, i.e., in the first three rows of panel *B*). In these cases, the results on the left-hand side of panel *B*) show that, for those subjects who are non-pro-social, the fact that they are also envious significantly reduces the probability of telling the truth. In contrast, envy does not have a significant effect for pro-social subjects.²⁴ Similarly, for those subjects who are envious, the fact that they are also non-pro-social significantly reduces the probability of telling the truth. However, pro-sociality does not have a significant effect for non-envious subjects. Thus, our results suggest that it is the combination of envious and non-pro-social attitudes that leads subjects to lie in situations in which they would be expected to report the truth.²⁵

²³ We also evaluated the marginal effect of the risk variable in each of the four possible information sets of player 1 (i.e., regarding whether player 1 and player 2 are interested), as well as in each of the four possible combinations of social preferences of player 1 (i.e., regarding whether player 1 is pro-social and envious). We find that in the first case, the value of the marginal effect of risk ranges from -0.02843 to -0.02213, while in the second, it ranges from -0.02895 to -0.02216. In both cases, the differences between these values are not significant.

²⁴ Only marginally at the 10% level if the sender is interested and the receiver is not.

²⁵ See also Morgan, Steiglitz and Reis (2003) for an analysis of the effect of social preferences on bidding behavior in auctions.

Table 12. Determinants of the probability that player 1 sends a true message

A) Logit estimation

<i>NoInt1 × NoInt2 × Prosoc × NoEnv</i>	0.4020 (0.3575)
<i>NoInt1 × NoInt2 × NoProsoc × Env</i>	-0.8442** (0.4294)
<i>NoInt1 × NoInt2 × NoProsoc × NoEnv</i>	0.2168 (0.4534)
<i>NoInt1 × Int2 × Prosoc × Env</i>	0.0843 (0.3999)
<i>NoInt1 × Int2 × Prosoc × NoEnv</i>	0.4415 (0.3695)
<i>NoInt1 × Int2 × NoProsoc × Env</i>	-1.0487** (0.4714)
<i>NoInt1 × Int2 × NoProsoc × NoEnv</i>	0.6609 (0.5180)
<i>Int1 × NoInt2 × Prosoc × Env</i>	0.3446 (0.4069)
<i>Int1 × NoInt2 × Prosoc × NoEnv</i>	1.0806*** (0.4141)
<i>Int1 × NoInt2 × NoProsoc × Env</i>	-0.5787 (0.4610)
<i>Int1 × NoInt2 × NoProsoc × NoEnv</i>	1.5388** (0.6766)
<i>Int1 × Int2 × Prosoc × Env</i>	-0.9592** (0.3900)
<i>Int1 × Int2 × Prosoc × NoEnv</i>	-0.9540*** (0.3674)
<i>Int1 × Int2 × NoProsoc × Env</i>	-1.6444*** (0.5284)
<i>Int1 × Int2 × NoProsoc × NoEnv</i>	-1.6643*** (0.5149)
<i>Round</i>	0.0014 (0.0067)
<i>Risk</i>	-0.1276** (0.0560)
<i>Constant</i>	0.8704** (0.3578)

[Number of obs = 800], [***, **, * significant at the 1%, 5%, and 10% level, respectively]

B) Marginal Effects of Envy and Pro-sociality (at round 20 and the average risk level) on the probability of sending a true message

	Marginal effect of <i>Env</i> = 1 vs. <i>Env</i> = 0 (by values of <i>Prosoc</i> , <i>Int1</i> and <i>Int2</i>)		Marginal effect of <i>Prosoc</i> = 1 vs. <i>Prosoc</i> = 0 (by values of <i>Env</i> , <i>Int1</i> and <i>Int2</i>)	
	<i>Prosoc</i> = 0	<i>Prosoc</i> = 1	<i>Env</i> = 0	<i>Env</i> = 1
<i>Int1</i> = 0 & <i>Int2</i> = 0	-0.2550** (0.1111)	-0.0883 (0.0795)	0.0392 (0.0905)	0.2058** (0.1017)
<i>Int1</i> = 0 & <i>Int2</i> = 1	-0.3924*** (0.1175)	-0.0768 (0.0815)	-0.0419 (0.0928)	0.2736** (0.1098)
<i>Int1</i> = 1 & <i>Int2</i> = 0	-0.4027*** (0.1103)	-0.1296* (0.0755)	-0.0557 (0.0766)	0.2173** (0.1102)
<i>Int1</i> = 1 & <i>Int2</i> = 1	0.0036 (0.1133)	-0.0012 (0.0852)	0.1498 (0.0963)	0.1449 (0.1040)

Note that the cases considered in the first two rows of Panel B) are conceptually different from that in the third row because in those cases the monetary payoff of the sender is not affected by the action of the receiver. Thus the sender does not have a direct monetary interest in affecting the receiver's choice. In contrast, in the situation considered in the third row, the monetary payoff of the sender depends on the action of the receiver.

When instead we direct our attention to situations in which the informative equilibrium prescribes the sender to lie (i.e., when both the sender and the receiver are interested in the object – as in the fourth line of panel B)), we find that the marginal effects of envy and pro-sociality on the probability of sending a true message are not significant. In such a case, the coefficient representing the marginal effect of envy is essentially 0 and the coefficient representing the marginal effect of pro-sociality is positive but not significant (p-values of 0.16 and 0.12 for not envious and envious subjects, respectively). This suggests that the observation of truth-telling by the sender in situations in which the informative equilibrium would prescribe him to lie cannot be explained in a substantial way by social preferences and that, instead, an argument of a different nature is necessary. The main candidate is subjects' aversion to lying.

At this point, it is useful to discuss in greater detail the relationship with the findings of Gneezy (2005). In the situation considered in that paper, a sender has to tell a receiver, who is uninformed of her payoffs, what the receiver's most profitable action is. The action of the receiver, in turn, also has implications for the sender. There are three treatments. In all three treatments, the sender is better off when the receiver takes action A, and the receiver is better off taking action B. Thus, recommending B is a deception

that, if followed (which the receiver, in fact, tends to do), is beneficial for the sender. The key difference between the treatments is given by the extent to which the sender and receiver benefit from the two actions. In treatment 1, both agents benefit very little from the correct actions. In treatment 2, the sender benefits very little from her preferred action, but the receiver benefits substantially from it. Finally, in treatment 3, both players benefit considerably from their preferred actions. Importantly, the receiver is not informed of the conflict of interest, and hence, we cannot be certain whether he knows about it. In any case, in the observed behavior, she tends to follow the advice of the sender, and thus, the priors regarding the presence of that conflict cannot be very large. The payoffs in Gneezy (2005) are summarized in Table 13.

Table 13. Payoffs in Gneezy (2005)

<i>Treatment</i>	<i>Option</i>	<i>Payoff to</i>	
		<i>Sender</i>	<i>Receiver</i>
1	A	5	6
	B	6	5
2	A	5	15
	B	6	5
3	A	5	15
	B	15	5

As one can easily see, treatment 1 in Gneezy (2005) is somewhat similar to the situation arising in the game considered in our *Simplified* treatment when both sender and receiver are not interested in the object. Treatment 2 is closer to the situation in our game when the sender is not interested in the object but the receiver is. Treatment 3 is then similar to the situation in our game when both sender and receiver are interested in the object. Interestingly, the frequency with which senders recommend option A to the receiver (recall, this is the best for them but not for the receivers) is 17% in treatment 1, 36% in treatment 2, and 52% in treatment 3. From Table 8 we see that the frequency of lies in the situation in our experiment that we argued is similar to treatment 1 (both not interested) is 37%, for the one similar to treatment 2 (sender not interested, receiver interested) is 35%, and for the one similar to treatment 3 (both interested) is 66%. Given how different the experiments are, this similarity is rather remarkable, particularly because Gneezy (2005) demonstrates, in a separate experiment conducted with a different subject pool, that individuals are more willing to accept a lie if it benefits a lower-income person.

4.4 Other equilibria

As mentioned above, the game of the *Simplified* treatment has other equilibria, besides the informative one. Thus, it is important to examine to what extent the observed individual behavior is consistent with subjects using strategies from other equilibria.

In this section, we present a maximum likelihood error-rate analysis of senders' choices following the econometric model used in Costa-Gomes, Crawford and Broseta (2001) adapted to our framework. This is a mixture model in which each sender's type is drawn from a common prior distribution over types and remains constant for the periods in which the player acts as a sender. In our analysis, a sender's type is associated with playing one of the 16 available strategies when acting as sender.²⁶ Our final objective is to classify experimental subjects according to types (sender's strategies) and identify whether the strategy assigned to an individual is a best response to the (aggregate) behavior of the receivers that he experienced during the experiment. In this sense, we can obtain an indicator of the frequency of equilibrium behavior in the simplified game (differentiating among equilibria), as well as of non-equilibrium behavior due to other motives.²⁷

A strategy for the sender is a vector of four components $(c_{ni,ni}, c_{ni,i}, c_{i,ni}, c_{i,i}) \in \{0,1\}^4$, where $c_{ni,ni}$ is the sender's choice in the information set in which both the sender and the receiver are not interested in the object, $c_{ni,i}$ is the choice in the information set in which the sender is not interested and the receiver is interested, $c_{i,ni}$ represents the choice in the information set in which the sender is interested and the receiver is not, and $c_{i,i}$ represents the choice in the information set in which both the sender and the receiver are interested. For each information set, the choice $c = 1$ indicates that the sender sends a true message (i.e., the content of the message is the true color of the object), while the choice $c = 0$ indicates that the sender sends a false message (i.e., the content of the

²⁶ A sender has four information sets and two actions in each one. This results in 16 available strategies, and therefore, there are 16 possible types.

²⁷ An alternative modeling strategy to identify subjects' types was pursued by Cai and Wang (2006) and Wang, Spezio and Camerer (2010). They fit a level- k model, where level-0 for player 1 is truth-telling and for player 2 is to naively believe player 1. The behavior of players at all other levels is then strictly pinned down in their game by this choice of level 0 behavior. We did not pursue this route for two reasons. Using the same specification for level 0 does not uniquely pin down behavior in our game for other levels of $k > 1$ because of the insensitivity of player 1's payoff to player 2's action when player 1 is not interested. Furthermore, the observed behavior of player 1 when both players are interested is inconsistent with that kind of model.

message is the opposite color of that of the object). Hence, we consider each strategy $(c_{ni,ni}, c_{ni,i}, c_{i,ni}, c_{i,i}) \in \{0,1\}^4$ to be a sender's type.

For the estimation of the mixture model, let $i = 1, \dots, N$ index the different players and $k = 1, \dots, K$ index our types. We assume that a type- k player normally makes a type k decision, but in each period, he makes an error with probability $\varepsilon_k \in [0, 1]$, constituting type k 's *error rate*, in which case he chooses to send a true or a false message with equal probability $\frac{1}{2}$. For a type- k player, the probability of a type k decision in any information set is then $1 - \frac{1}{2}\varepsilon_k$. Hence, the probability of a non-type- k decision is $\frac{\varepsilon_k}{2}$. We assume that errors are independently and identically distributed across periods and players and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_K)$.

The likelihood function can be constructed as follows. Let T_i denote the total number of periods in which player i acted as sender. Next, let x_{ik} denote the number of player i 's decisions that equal type k 's in periods in which he acts as a sender and $x_i = (x_{i1}, \dots, x_{iK})$, $x = (x_1, \dots, x_i, \dots, x_N)$. Let p_k denote the common probability that a player is of type k , $\sum_{k=1}^K p_k = 1$ and $p = (p_1, \dots, p_K)$. As each period has one type- k decision and one non-type- k decision, the probability of observing a particular sample with x_{ik} type- k decisions when player i is type k can be written as follows:

$$L_k^i(\varepsilon_k | x_{ik}) = \left[1 - \frac{1}{2}\varepsilon_k\right]^{x_{ik}} \left[\frac{1}{2}\varepsilon_k\right]^{T_i - x_{ik}}$$

Weighting the right-hand side by p_k , summing over k , taking logarithms, and summing over i yields the log-likelihood function for the entire sample:

$$\ln L(p, \varepsilon | x) = \sum_{i=1}^N \ln \sum_{k=1}^K p_k L_k^i(\varepsilon_k | x_{ik})$$

This function is maximized by the EM algorithm.²⁸ We find that the most parsimonious model is the one with the following four types:²⁹ T1 playing strategy (1,1,1,0) associated with the informative equilibrium, T2 playing strategy (0,1,1,0) associated with the *extra* equilibrium, T3 playing strategy (1,0,0,0), and T4 playing strategy (1,1,1,1). Thus, types 1 and 2 are equilibrium types, whereas types 3 and 4 are

²⁸ As proposed in the seminal paper by Dempster, Laird and Rubin (1977).

²⁹ To select the model with four types, we apply the following iterated procedure. We begin by estimating this model using all ($n = 16$) available types and compute the BIC and AIC. Then, we estimate all models with $n-1$ types, choose the best one and compute AIC and BIC. If the last one performs better, we continue estimating all models with $n-2$ types; otherwise, we stop.

not. Type 3 only tells the truth when no one likes the object, and type 4 always tells the truth.

Table 14 shows the estimated parameters of this model. The two types associated with equilibrium strategies account for more than half of the observed behavior. Furthermore, type T3 is quite noisy, as the estimated error rate is 0.80 (implying that 40% of the observed choices are not representative of type-T3 behavior).

Table 14. Error-rate model

Type	Coef.	Probabilities of types		
		Std Err.	P> z	# subjects
T1: (1,1,1,0)	0.32130	0.08775	0.00025	14
T2: (0,1,1,0)	0.24680	0.10528	0.01907	8
T3: (1,0,0,0)	0.33279	0.11293	0.00321	14
T4: (1,1,1,1)	0.09912	0.06008	0.09897	4

Type	Coef.	Error rates	
		Std Err.	P> z
T1: (1,1,1,0)	0.24654	0.06887	0.00000
T2: (0,1,1,0)	0.60924	0.12789	0.00225
T3: (1,0,0,0)	0.80886	0.12319	0.12076
T4: (1,1,1,1)	0.17248	0.18238	0.00001

Loglikelihood -458.6315

(For error rates, the null hypothesis is $\varepsilon_k = 1$)

Next, as we anticipated at the beginning of this section, using the above estimated parameters, we compute the probability with which each subject is assigned to a specific type, conditional on the observed pattern of choices. The procedure is as follows: for each subject, we calculate the probabilities of observing his pattern of choices conditional on type. Then, using Bayes' rule, we can compute the probability that an individual i is of type k , given the observed choices. Finally, we assign the subject to the type (T1, T2, T3 or T4) with the highest probability. We find that 22 subjects out of 40 are assigned to one of the two equilibrium types, i.e., T1 or T2 (the number of subjects assigned to each type is reported in the last column of Table 14).

We then check whether the assigned type is a best response to the (aggregate) behavior of the receivers they encounter during the experiment. Considering the 22 subjects classified in one of the equilibrium types (T1 or T2), for 20 of them (i.e., 90.91%), the assigned strategy is a best response to the aggregate behavior of the receivers they encounter. Conversely, for the 18 subjects classified in the two non-equilibrium

types (T3 or T4), the strategy is not a best response to the behavior observed from the receivers they faced (taking into account only the senders' monetary payoffs). As type T3 lies, with the only exception being when neither the sender nor the receiver is interested, such behavior might conceivably be influenced by envious and/or non-pro-social preferences. Moreover, T4 always tells the truth, including when the sender and the receiver have conflicting interests, and this behavior may then be related to aversion to lying.

Even within the more than 50% of the subjects whose behavior is consistent with equilibrium, social preferences might still play a role. Approximately one-third of those subjects choose a strategy that corresponds to the extra equilibrium, where the sender's payoff is the same as in the informative equilibrium, while the receiver has a lower payoff – hence, we may argue that this equilibrium selection may also be due to social preferences.

5. Seller of information without conflict of interest

In this section, we are interested in analyzing the role played by a conflict of interest between buyers and sellers of reports in the collapse of the market for information in the experiments we have described thus far. To this end, we analyze another treatment in which a fourth player (player 0) is added. This player cannot participate in the auction (he is uninterested in the object) and is the only one allowed to sell reports. In stage 1, all players decide whether to pay a cost of 20 to acquire information, with player 0 being the first to decide and the other three deciding in sequence. In stage 2, only player 0 can sell reports; the rest of the game proceeds as in the *Base* treatment.

The *informative* equilibrium for the game we consider in this treatment, referred to as *Uninterested*³⁰ is as follows: player 0 acquires information and sells his report to two buyers, at a price of 12.5 ECUs. The report sent is always truthful, and thus, two of the players participating in the auction learn the true color of the object, and the remaining bidder remains uninformed.

Table 15 presents the results of the auction stage (third stage). Analogously to Table 2, we divide behavior according to the information of the subject: informed players, uninformed players and buyers of reports.

³⁰ See *CG* for a proof.

In comparing Table 15 to Table 2, we should stress that the behavior of buyers of reports is much more sensitive to the content of the reports, which we can view as higher trust in their informational content.

Table 15. Average bids by type of player and block of 10 rounds in treatment *Uninterested*

		<i>Rounds 1-10</i>	<i>Rounds 11-20</i>
<i>Informed players</i>	<i>Color - Yes</i> [Prediction: 200]	138.14 (67.67)	183.73 (45.74)
	<i>Color - No</i> [Prediction: 100]	72.16 (53.38)	97.55 (50.66)
<i>Uninformed players</i> [Prediction: 150]		120.15 (70.69)	141.85 (66.17)
<i>Buyers of report</i>	Content: <i>Color - Yes</i> [Prediction: 200]	132.46 (67.86)	171.54 (69.37)
	Content: <i>Color - No</i> [Prediction: 100]	81.58 (56.62)	89.17 (58.69)
	Content: <i>No info</i> [Prediction: 150]	73.80 (65.73)	77.34 (81.92)

(In brackets, Std. Dev.)

In Table 16, analogously to Table 4 for the *Base* treatment, we present the frequencies of the choices with which participants in the experiment (i) acquire information in the first stage (*Inform*), (ii) buy a report in the second stage (*Buy rep*) and (iii) remain uninformed (*Uninf*), distinguishing by mover and by block of 10 rounds.

We find that, as in our *Base* treatment, the market for reports is never very large and ultimately collapses. However, in a sense, the collapse is even more significant in this case because the prediction now is that two-thirds of the potential buyers should have bought the report, whereas it was only one-third in the previous treatment. Given the prices at which reports are sold, the number of buyers of reports is far lower (approximately one-third) of what would be necessary for the market to be viable.

To understand why potential buyers of reports buy so few of them, it is important to analyze the informative content of the reports sent by sellers in treatment *Uninterested*. In this case, as the seller of reports (mover 0) cannot participate in the auction, his monetary payoff cannot be affected by the buyers' behavior in the auction. Hence, there is no way that lies or uninformative reports can affect their monetary payoffs. However, when we analyze the data, we find that of the 69 reports issued by the sender (mover 0) 58 (84.06%) were truthful, and the remaining 11 (15.94%) were either 0 reports or false

reports. This lack of truthful behavior clearly plays a role in the collapse of the market for information.

Table 16. Behavior in information markets – Absolute number of observations (treatment *Uninterested*)

	<i>Rounds 1-10</i>			<i>Rounds 11-20</i>		
	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>
<i>Mover 0</i>	46 (38.33)	74 (61.67)	---	71 (59.17)	49 (40.83)	---
<i>Mover 1</i>	61 (50.83) [36.31]	38 (31.67) [27.94]	21 (17.50) [37.50]	88 (73.33) [38.60]	23 (19.17) [23.71]	9 (7.50) [25.71]
<i>Mover 2</i>	50 (41.67) [29.76]	54 (45.00) [39.71]	16 (13.33) [28.57]	59 (49.17) [25.88]	49 (40.83) [50.52]	12 (10.00) [34.29]
<i>Mover 3</i>	57 (47.50) [33.93]	44 (36.67) [32.35]	19 (15.83) [33.93]	81 (67.50) [35.53]	25 (20.83) [25.77]	14 (11.67) [40.00]
<i>Total1-3</i>	168 (46.67) [100]	136 (37.78) [100]	56 (15.56) [100]	228 (63.33) [100]	97 (26.94) [100]	35 (9.72) [100]

(% over total row) [% over total column excluding Mover 0]

6. Conclusion

In this paper, we study experimentally the viability of markets for information, where information is transmitted via cheap-talk reports. This type of game has equilibria with and without information transmission, and hence an empirical assessment of the viability of information transmission seems necessary. Furthermore, previous results in the experimental literature on cheap-talk games suggests that agents in the lab may tell the truth even when theory predicts that reports should be uninformative.

In the laboratory, we find that in the last iterations of the game played in the experiment, very few reports are sold. We observe that some agents indeed tell the truth when their monetary payoffs could be increased by sending deceptive reports. However, a novel finding in our experiment is that some agents lie when doing so does not increase their monetary payoff. We show that the agents who are deceptive in the game we consider are either non-pro-social or envious by studying their choices in a different game designed to elicit these traits. This deceptive behavior is a main reason for the collapse of the market for information.

To investigate the role of the conflict of interest between senders and receivers in our findings, we run an additional treatment in which the monetary payoff of the seller of reports is independent of the action of the receiver. We find that even in this treatment, the seller of reports sometimes misinforms the buyers of reports, again leading to the collapse of the market for information.

Our results are also robust to other extensions. In particular, we considered two other treatments (denoted *Option* and *Unint-Opt*) in which we again give subjects the option to directly acquire the information once they have observed the prices of the reports. This option might make subjects believe that it is “more secure” not to directly acquire the information and wait for the market for reports, to determine whether it is worthwhile to buy a report (if the prices of the reports are very high, subjects can still acquire the information directly). We found (for details, see the online Appendix) that this option is almost never used, and the results of treatments *Option* and *Unint-Option* are not significantly different from their respective counterparts in which the option is not available.

We believe that our paper provides important insights for real-life settings. For example, one important mission of organizations is information transmission (see, e.g., Garicano (2000) and Alonso, Dessein and Matouscheck (2008)). Our results suggest that social preferences within organizations could be detrimental to this important function of firms even in cases in which the information does not appear to lead to material conflicts of interest. Clearly, more evidence is needed, hopefully from a field setting, to determine whether this is indeed the case, but we hope that our work forms the first step in an important agenda.

References

- Alonso, R., Dessein, W., & Matouschek, N. (2008). When does coordination require centralization? *The American Economic Review*, 98(1), 145-179.
- Bartling, B., Fehr, E., Maréchal, M. A., & Schunk, D. (2009). Egalitarianism and competitiveness. *The American Economic Review*, 99(2), 93-98.
- Berry, S. T. (1992). Estimation of a Model of Entry in the Airline Industry. *Econometrica*, 889-917.
- Blume, A., DeJong, D. V., Kim, Y. G., & Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *American Economic Review*, 1323-1340.

- Blume, A., DeJong, D. V., Kim, Y. G., & Sprinkle, G. B. (2001). Evolution of communication with partial common interest. *Games and Economic Behavior*, 37(1), 79-120.
- Bolton, P., Freixas, X., & Shapiro, J. (2007). Conflicts of interest, information provision, and competition in the financial services industry. *Journal of Financial Economics*, 85(2), 297-330.
- Brandts, J., & Charness, G. (2003). Truth or consequences: An experiment. *Management Science*, 49(1), 116-130.
- Cabrales, A., & Gottardi, P. (2014). Markets for information: Of inefficient firewalls and efficient monopolies. *Games and Economic Behavior*, 83(1), 24-44.
- Cai, H., & Wang, J. T. Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1), 7-36.
- Charness, G., & Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6), 1579-1601.
- Charness, G., & Gneezy, U. (2010). Portfolio choice and risk attitudes: An experiment. *Economic Inquiry*, 48(1), 133-146.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5), 1193-1235.
- Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6), 1431-1451.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), 1-38.
- Dickhaut, J. W., McCabe, K. A., & Mukherji, A. (1995). An experimental study of strategic information transmission. *Economic Theory*, 6(3), 389-403.
- Erev, I., & Rapoport, A. (1998). Coordination, "magic," and reinforcement learning in a market entry game. *Games and Economic Behavior*, 23(2), 146-175.
- Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1), 53-82.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5), 874-904.
- Gneezy, U. (2005). Deception: The role of consequences. *The American Economic Review*, 95(1), 384-394.
- Hurkens, S., & Kartik, N. (2009). Would I lie to you? On social preferences and lying aversion. *Experimental Economics*, 12(2), 180-192.

- Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, 76(4), 1359-1395.
- Kartik, N., Ottaviani, M., & Squintani, F. (2007). Credulity, lies, and costly talk. *Journal of Economic Theory*, 134(1), 93-116.
- Kawagoe, T., & Takizawa, H. (1999). Instability of babbling equilibrium in cheap talk games. Saitama University.
- Kawagoe, T., & Takizawa, H. (2005). Why lying pays: Truth bias in the communication with conflicting interests. *Available at SSRN 691641*.
- Morgan, J., Steiglitz, K., & Reis, G. (2003). The spite motive and equilibrium behavior in auctions. *Contributions in Economic Analysis & Policy*, 2(1), Art. 5.
- Sánchez-Pagés, S., & Vorsatz, M. (2007). An experimental study of truth-telling in a sender–receiver game. *Games and Economic Behavior*, 61(1), 86-112.
- Sobel, J. (2013). Giving and receiving advice, in *Advances in Economics and Econometrics*, D. Acemoglu, M. Arellano, and E. Dekel (eds.).
- Sutter, M. (2009). Deception through telling the truth?! Experimental evidence from individuals and teams. *The Economic Journal*, 119(534), 47-60.
- Wang, J. T. Y., Spezio, M., & Camerer, C. F. (2010). Pinocchio's Pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3), 984-1007.

ONLINE APPENDIX

A) EQUILIBRIUM ANALYSIS OF THE SIMPLIFIED SENDER-RECEIVER GAME

In Table 17 below, we present the expected payoffs to players 1 and 2 associated with all of the pure strategy profiles and identify the Bayes-Nash equilibria (in grey), i.e., the candidates for perfect Bayesian equilibria.³¹ In the rows of Table 17, we list all of the strategy profiles of player 1 and in the columns those of player 2. A strategy of player 1 (the sender) is a vector of four components $(c_{ni,ni}, c_{ni,i}, c_{i,ni}, c_{i,i}) \in \{0,1\}^4$, as defined in Section 4.1. A strategy of player 2 (the receiver) is a vector of two components $(c_{m_{ni}}, c_{m_i}) \in \{L, C, R\}^2$. The component $c_{m_{ni}}$ represents player 2's choice in the information set in which player 1's message says that player 2 is not interested in the object (i.e., the color reported in the message does not coincide with player 2's assigned color); c_{m_i} represents player 2's choice in the information set in which player 1's message says that player 2 is interested in the object. For each information set, the choice $c = L$ indicates that player 2 chooses action *left*, the choice $c = C$ indicates that player 2 chooses action *center*, and the choice $c = R$ indicates that player 2 chooses action *right*.

Each cell of Table 17 contains a vector with the (*ex ante*) expected payoffs to player 1 and player 2, associated with the respective strategies of players 1 and 2 indicated by row and column. The expected payoffs are computed using the payoffs contained in Table 7, taking into account that each of the four possible combinations of players 1 and 2 being interested/not interested in the object has an *ex ante* probability of 1/4.

In this table, we also mark in bold the best responses of players 1 and 2 and fill in grey those cells in which the strategies of players 1 and 2 are mutual best responses, i.e., the (pure strategy) Bayes-Nash equilibria.

We find that there are 10 Bayes-Nash equilibria that can be grouped into 3 classes (*informative* equilibria, *babbling* equilibria and *extra* equilibria). The equilibria within each class are informationally equivalent and only differ: (i) in the use of colors, as each equilibrium has a reverse one, and (ii) in the case of the babbling and extra equilibria, also in the choice of player 2 in one of his information sets, which can be either *left* or *center*.

³¹ We then check which Bayes-Nash equilibria are perfect Bayesian equilibria.

Table 17. Bayes-Nash equilibria of treatment *Simplified*

	(L,L)	(L,C)	(L,R)	(C,L)	(C,C)	(C,R)	(R,L)	(R,C)	(R,R)
(1, 1, 1, 1)	(20.0, 67.5)	(32.5, 55.0)	(45.0, 35.0)	(32.5, 92.5)	(45.0, 80.0)	(57.5, 60.0)	(45.0, 110)	(57.5, 97.5)	(70.0, 77.5)
(1, 1, 1, 0)	(20.0, 67.5)	(20.0, 57.5)	(20.0, 40.0)	(45.0, 90.0)	(45.0, 80.0)	(45.0, 62.5)	(70.0, 105)	(70.0, 95.0)	(70.0, 77.5)
(1, 1, 0, 1)	(20.0, 67.5)	(45.0, 70.0)	(70.0, 57.5)	(20.0, 77.5)	(45.0, 80.0)	(70.0, 67.5)	(20.0, 87.5)	(45.0, 90.0)	(70.0, 77.5)
(1, 0, 1, 1)	(20.0, 67.5)	(32.5, 65.0)	(45.0, 62.5)	(32.5, 82.5)	(45.0, 80.0)	(57.5, 77.5)	(45.0, 82.5)	(57.5, 80.0)	(70.0, 77.5)
(0, 1, 1, 1)	(20.0, 67.5)	(32.5, 65.0)	(45.0, 55.0)	(32.5, 82.5)	(45.0, 80.0)	(57.5, 70.0)	(45.0, 90.0)	(57.5, 87.5)	(70.0, 77.5)
(1, 1, 0, 0)	(20.0, 67.5)	(32.5, 72.5)	(45.0, 62.5)	(32.5, 75.0)	(45.0, 80.0)	(57.5, 70.0)	(45.0, 82.5)	(57.5, 87.5)	(70.0, 77.5)
(1, 0, 1, 0)	(20.0, 67.5)	(20.0, 67.5)	(20.0, 67.5)	(45.0, 80.0)	(45.0, 80.0)	(45.0, 80.0)	(70.0, 77.5)	(70.0, 77.5)	(70.0, 77.5)
(0, 1, 1, 0)	(20.0, 67.5)	(20.0, 67.5)	(20.0, 60.0)	(45.0, 80.0)	(45.0, 80.0)	(45.0, 72.5)	(70.0, 85.0)	(70.0, 85.0)	(70.0, 77.5)
(1, 0, 0, 1)	(20.0, 67.5)	(45.0, 80.0)	(70.0, 85.0)	(20.0, 67.5)	(45.0, 80.0)	(70.0, 85.0)	(20.0, 60.0)	(45.0, 72.5)	(70.0, 77.5)
(0, 1, 0, 1)	(20.0, 67.5)	(45.0, 80.0)	(70.0, 77.5)	(20.0, 67.5)	(45.0, 80.0)	(70.0, 77.5)	(20.0, 67.5)	(45.0, 80.0)	(70.0, 77.5)
(0, 0, 1, 1)	(20.0, 67.5)	(32.5, 75.0)	(45.0, 82.5)	(32.5, 72.5)	(45.0, 80.0)	(57.5, 87.5)	(45.0, 62.5)	(57.5, 70.0)	(70.0, 77.5)
(1, 0, 0, 0)	(20.0, 67.5)	(32.5, 82.5)	(45.0, 90.0)	(32.5, 65.0)	(45.0, 80.0)	(57.5, 87.5)	(45.0, 55.0)	(57.5, 70.0)	(70.0, 77.5)
(0, 1, 0, 0)	(20.0, 67.5)	(32.5, 82.5)	(45.0, 82.5)	(32.5, 65.0)	(45.0, 80.0)	(57.5, 80.0)	(45.0, 62.5)	(57.5, 77.5)	(70.0, 77.5)
(0, 0, 1, 0)	(20.0, 67.5)	(20.0, 77.5)	(20.0, 87.5)	(45.0, 70.0)	(45.0, 80.0)	(45.0, 90.0)	(70.0, 57.5)	(70.0, 67.5)	(70.0, 77.5)
(0, 0, 0, 1)	(20.0, 67.5)	(45.0, 90.0)	(70.0, 105)	(20.0, 57.5)	(45.0, 80.0)	(70.0, 95.0)	(20.0, 40.0)	(45.0, 62.5)	(70.0, 77.5)
(0, 0, 0, 0)	(20.0, 67.5)	(32.5, 92.5)	(45.0, 110)	(32.5, 55.0)	(45.0, 80.0)	(57.5, 97.5)	(45.0, 35.0)	(57.5, 60.0)	(70.0, 77.5)

The equilibria for this game are as follows:

1.- *Informative equilibrium:*

- ((1,1,1,0), (R, L))
- ((0,0,0,1), (L, R))

The expected payoffs for players 1 and 2 are 70 and 105, respectively.

2.- *Babbling equilibrium:*

- ((1,0,1,0), (C, L or C))
- ((0,1,0,1), (L or C, C))

The expected payoffs for players 1 and 2 are 45 and 80, respectively.

3.- *Extra equilibrium:*

- ((0,1,1,0), (R, L or C))
- ((1,0,0,1), (L or C, R))

The expected payoffs for players 1 and 2 are 70 and 85, respectively

We now check that all the Bayes-Nash equilibria we found in Table 22 also constitute (weak) perfect Bayesian equilibria, hereafter PBE.

1.- *Informative equilibrium*

As both profiles are informationally equivalent (they only differ in the use of colors), let us consider the profile $((1,1,1,0), (R, L))$.

1.i) *Player 2's beliefs*

If the message of player 1 says that player 2 is not interested in the object, then, given the strategy of player 1, the beliefs of player 2 derived from Bayes' rule assign equal probability (1/3) to the following three events: (i) player 2 is not interested in the object and player 1 is interested, (ii) neither player 2 nor player 1 is interested in the object, and (iii) both players 2 and 1 are interested in the object.

If the message of player 1 says that player 2 is interested in the object, then, given the strategy of player 1, the beliefs of player 2 assign probability 1 to the following event: player 2 is interested in the object, and player 1 is not interested.

1.ii) *Sequential rationality of player 2*

If the message of player 1 says that player 2 is not interested in the object, then, given the beliefs above, the expected payoffs of player 2 associated with the choices L , C and R are $110/3$, $200/3$ and $260/3$, respectively (see Table 7). Thus, the choice $c_{m,ni} = R$ prescribed by the strategy of player 2 is sequentially rational.

If the message of player 1 says that player 2 is interested in the object, then, given the beliefs above, the expected payoffs of player 2 associated with the choices L , C and R are 160, 120 and 50, respectively. Thus the choice $c_{m,i} = L$ prescribed by the strategy of player 2 is sequentially rational.

1.iii) *Sequential rationality of player 1*

If player 1 is not interested in the object, his payoff to player 1 is 20, regardless the choices of players 1 and 2. Thus, the choices $c_{ni,ni} = 1$ and $c_{ni,i} = 1$ prescribed by the strategy of player 1 are sequentially rational.

If player 1 is interested in the object and player 2 is not interested, then, given the strategy of player 2, (R, L) , the payoffs to player 1 associated with sending a true and false message are, respectively, 120 and 20. Thus the choice $c_{i,ni} = 1$ prescribed by the strategy of player 1 is sequentially rational.

If both player 1 and player 2 are interested in the object, then, given the strategy of player 2, (R, L) , the payoffs to player 1 associated with sending a true and a false

message are, respectively, 20 and 120 (see Table 7). Thus the choice $c_{i,i} = 0$ prescribed by the strategy of player 1 is sequentially rational.

Hence, $((1,1,1,0), (R, L))$ is a PBE, and therefore, $((0,0,0,1), (L, R))$ is also a PBE.

2.- Babbling equilibrium

Because we have two sets of profiles that are informationally equivalent (they only differ in the use of colors), let us focus on the profiles $((1,0,1,0), (C, C))$ and $((1,0,1,0), (C, L))$

2.i) Player 2's beliefs

In this case, the message of player 1 is not correlated with the state of the world: it always says that player 2 is not interested. Thus, if the message of player 1 says that player 2 is not interested in the object, then the beliefs of player 2 using Bayes' rule assign equal probability (1/4) to each of the four possible events regarding whether players 1 and 2 are interested in the object.

If the message of player 1 says that player 2 is interested in the object (which does not happen on the equilibrium path), then beliefs cannot be determined by Bayes' rule and are specified below.

2.ii) Sequential rationality of player 2

If the message of player 1 says that player 2 is not interested in the object, then, given the beliefs above, the expected payoffs to player 2 associated with the choices L , C and R are $270/4$, $320/4$ and $310/4$, respectively. Thus the choice $c_{m,ni} = C$ prescribed by the strategy of player 2 is sequentially rational.

If the message of player 1 says that player 2 is interested in the object, then we can find beliefs such that both the choices (i) $c_{m,i} = C$ and (ii) $c_{m,i} = L$ are sequentially rational (the beliefs are free in this case). For instance, if the beliefs in this information set are the same as in the former one (i.e., all four possible events have the same probability), then the choice $c_{m,i} = C$ is sequentially rational. Alternatively, if the beliefs in this information set assign probability 1 to the event in which player 2 is interested in the object and player 1 is not, then the choice $c_{m,i} = L$ is sequentially rational.

2.iii) *Sequential rationality of player 1*

If player 1 is not interested in the object, then the payoff to player 1 is 20, regardless of the choices of players 1 and 2. Thus the choices $c_{ni,ni} = 1$ and $c_{ni,i} = 0$ prescribed by the strategy of player 1 are sequentially rational.

If player 1 is interested in the object and player 2 is not interested, then we have the following:

- Given the strategy of player 2 in the first equilibrium in this class, i. e., (C, C) , the payoff of player 1 is 70 regardless of whether he sends a true or a false message. Thus, the choice $c_{i,ni} = 1$ prescribed by the strategy of player 1 is sequentially rational.
- Given the strategy of player 2 in the second equilibrium in this class, i. e., (C, L) , the payoffs of player 1 associated with sending a true and a false message are 70 and 20, respectively. Thus, the choice $c_{i,ni} = 1$ prescribed by the strategy of player 1 is sequentially rational.

If both players 1 and 2 are interested in the object, then we have the following:

- Given the strategy of player 2 in the first equilibrium in this class, i. e., (C, C) , the payoff of player 1 is 70 regardless of whether he sends a true or a false message. Thus, the choice $c_{i,ni} = 0$ prescribed by the strategy of player 1 is sequentially rational.
- Given the strategy of player 2 in the second equilibrium in this class, i. e., (C, L) , the payoffs of player 1 associated with sending a true and a false message are, 20 and 70, respectively. Thus, the choice $c_{i,ni} = 0$ prescribed by the strategy of player 1 is sequentially rational.

Hence, the profiles $((1,0,1,0), (C, L \text{ or } C))$ are PBE, and therefore, $((0,1,0,1), (L \text{ or } C, C))$ also are PBE.

3.- *Extra equilibrium*

As we have two sets of profiles that are informationally equivalent (they only differ in the use of colors), let us focus on the profiles $((0,1,1,0), (R, L))$ and $((0,1,1,0), (R, C))$

3.i) *Player 2's beliefs*

If the message of player 1 says that player 2 is not interested in the object, then, given the strategy of player 1, the beliefs derived from Bayes' rule of player 2 assign

equal probability (1/2) to the following two events: (i) player 2 is not interested and player 1 is interested in the object, and (ii) both players 2 and 1 are interested in the object.

If the message of player 1 says that player 2 is interested in the object, then, given the strategy of player 1, the beliefs of player 2 assign equal probability (1/2) to the following two events: (i) neither player 2 nor player 1 is interested in the object, and (ii) player 2 is interested in the object and player 1 is not interested.

3.ii) *Sequential rationality of player 2*

If the message of player 1 says that player 2 is not interested in the object, then, given the beliefs above, the expected payoffs of player 2 associated with the choices L , C and R are $90/2$, $140/2$ and $160/2$, respectively. Thus the choice $c_{m,ni} = R$ prescribed by the strategy of player 2 is sequentially rational.

If the message of player 1 says that player 2 is interested in the object, then, given the beliefs above, the expected payoffs of player 2 associated with the choices L , C and R are $180/2$, $180/2$ and $150/2$, respectively. Thus, the choices $c_{m,i} = L$ and $c_{m,i} = C$ are sequentially rational.

3.iii) *Sequential rationality of player 1*

If player 1 is not interested in the object, then his payoff is 20, regardless of the choices of players 1 and 2. Thus, the choices $c_{ni,ni} = 0$ and $c_{ni,i} = 1$ prescribed by the strategy of player 1 are sequentially rational.

If player 1 is interested in the object and player 2 is not interested, then we have the following:

- Given the strategy of player 2 in the first equilibrium in this class, i.e. (R, L), the payoffs of player 1 associated with sending a true and a false message are 120 and 20, respectively. Thus, the choice $c_{i,ni} = 1$ prescribed by the strategy of player 1 is sequentially rational.
- Given the strategy of player 2 in the second equilibrium in this class, i.e., (R, C), the payoffs of player 1 associated with sending a true and a false message are 120 and 70, respectively. Thus, the choice $c_{i,ni} = 1$ prescribed by the strategy of player 1 is sequentially rational.

If both player 1 and player 2 are interested in the object, then we have the following:

- Given the strategy of player 2 in the first equilibrium in this class, i. e., (R, L), the payoffs of player 1 associated with sending a true and a false message are 20 and

120, respectively. Thus, the choice $c_{i,i} = 0$ prescribed by the strategy of player 1 is sequentially rational.

- Given the strategy of player 2 in the second equilibrium in this class, i.e., (R, C) , the payoffs of player 1 associated with sending a true and a false message are 70 and 120, respectively. Thus, the choice $c_{i,i} = 0$ prescribed by the strategy of player 1 is sequentially rational.

Hence, $((0,1,1,0), (R, L \text{ or } C))$ is a PBE, and therefore, $((1,0,0,1), (L \text{ or } C, R))$ is also a PBE.


B) EXPERIMENTAL INSTRUCTIONS

B.1 Experimental instructions of treatment *Base*³²


The aim of this experiment is to study how individuals make decisions in certain contexts. The instructions are simple. If you follow them carefully you will earn a non-negligible amount of money in cash (euros) at the end of the experiment. During the experiment, your earnings will be in ECUs (experimental currency units). Individual payments will remain private, as nobody will know the other participants' payments. Any communication among you is strictly forbidden and will result in immediate exclusion from the experiment.

1. The experiment consists of 20 rounds. You will be randomly assigned to a group of 4 participants. This group is determined randomly at the beginning of the experiment and remains the same for all rounds. Moreover, you will be randomly assigned a player number within your group: you will be either player 1, player 2 or player 3. Your player number will remain the same throughout the experiment.
2. At the beginning of each round
 - a. You will be endowed with 250 ECUs that you can use to make the decisions within the round, as explained below.
 - b. You will be assigned a color (that will be immediately revealed to you) whose value for you is explained below.
3. At each round, you and the other players in your group will have the possibility to buy one object, by bidding in an auction (the auction rules will be detailed below). There will be one auctioned object, which can be either **orange** or **green**. The earnings of a player in case of getting the auctioned object depend on the color of the object:
 - If the object is equal to the player's assigned color, then the player will earn 200 ECUs.
 - If the object is different from the player's assigned color, then the player will earn 100 ECUs.
4. At the beginning of each round, the object to be auctioned is randomly drawn by the computer from an (virtual) urn containing two objects: one **orange object** and one **green object**. Each object is picked with equal probability (50%).

³² We omit the experimental instructions of treatments *Uninterested*, *Option* and *Unint-Opt*, which are variations of the instructions of treatment *Base* (as explained in Sections 5 and 6). These instructions are available from the authors upon request.




The assigned colors of players 1, 2 and 3 for the round are determined in a similar way. There is one (virtual) urn for each of these three players, containing two pieces of paper: one **orange** and one **green**. The computer randomly (and independently) draws one piece of paper from each urn. Each piece of paper is picked with equal probability (50%). The piece of paper selected for each player determines that player's assigned color for the round.




- At each round, each player will take his/her decisions knowing his/her preferred color but not others' preferred colors.

FOR EXAMPLE,

if in a round the selected colors for players 1, 2 and 3 are: 

- Then, Player 2 will know: . In such a case, what player 2 will know about the colors of players 1 and 3 is that one of the next four combinations has been drawn, each of them with equal probability (25%):



- An analogous reasoning holds for players 1 and 3 (they only know their own assigned color).

- Initially, no player knows which object (orange or green) has been selected by the computer for the round. However, prior to the auction, in sequence, you and the other players in the group will have the possibility to become informed of the color of the object to be auctioned by paying 20 ECUs. These decisions take place according to the following sequence: first player 1, then player 2 (knowing player 1's choice), and finally player 3 (knowing players 1 and 2's choices).
- Then, if at least one player has decided to acquire the information and at least one player has decided not to acquire it, there is a market for reports. In such a case, prior

to the auction, the players that have acquired the information can sell a report about the color of the object to the uninformed players. In all other cases, the color of the object is revealed to the players who decided to acquire the information, and all the players directly participate in the auction.

These are the rules of the market for reports:

- i. First, the players who have acquired the information, not knowing the color of the object yet, set a price for their report, and all players observe this price.
 - ii. The price of the report cannot exceed 20 ECUs.
 - iii. Then, according to the sequence (player 1 – player 2 – player 3), the uninformed players decide whether to buy one of the reports. When a player makes his/her choice, he/she will know the decisions of those players who acted before him/her in the sequence.
 - iv. The color of the object is revealed to all the players who decided to be informed (see point 6). The players who have sold a report decide the content of the report. The content can be: "*The object is orange*", "*The object is green*" or "*The object is orange or green*". Thus, the report can contain the true color, contain the false one, or be uninformative.
 - v. The buyers of the report receive it, and all the players participate in the auction.
8. Auction rules: Simultaneously, each player makes a bid for the object. The player that makes the highest bid gets the object. However, this player will not pay his bid, but the second-highest bid. The other players neither get the object nor pay anything.

For example: If player 1 bids 8 ECUs, player 2 bids 55 ECUs, and player 3 bids 18 ECUs, then player 2 (the highest bidder) receives the object and pays 18 ECUs for it (the second-highest bid). Players 0, 1 and 3 neither receive the object nor pay anything.

In case of ties in the highest bids, the computer randomly picks (with equal probability) the player who receives the object among those players who have made the highest bid. In such a case, the player who receives the object pays his/her own bid and the remaining players neither receive the object nor pay anything.

For example: If player 1 bids 55 ECUs, player 2 bids 55 ECUs, and player 3 bids 18 ECUs, then either player 1 or player 2 gets the object, with equal probability. The player who gets the object pays 55 ECUs.


9. Bidding rules:
 - a. Players who acquired the information: prior to knowing the color of the object (and thus prior to deciding the content of their report) must choose their bid
 - i. In case the object is green and
 - ii. In case the object is orange
 - b. Players who acquired the report: prior to knowing the content of the report must choose their bid
 - i. In case the message is “The object is orange”,
 - ii. In case the message is “The object is green” and
 - iii. in case the message is “The object is orange or green”
 - c. Players who did not acquire the report and did not acquire the information must choose their bid.

The bids are implemented automatically according to the true color of the object and the message received.
10. Summary of round payoffs. The round payoff of a player has three parts:
 - a. The endowment (250 ECUs) minus the payments (if any) incurred by the player either to be informed or to buy a report.
 - b. In the event of having sold reports, the player gets the agreed price from each buyer.
 - c. In the event of getting the auctioned object, the player gets either 200 ECUs (if the object is of his/her assigned color) or 100 ECUs (if it is not) minus his/her payment in the auction.
11. After the auction, and before proceeding to the next round, each player will receive the following ex post information:
 - a. The bids made by each player in the auction.
 - b. The player who obtained the auctioned object and the price paid for it.
 - c. The color of the auctioned object.
 - d. His/her round payoff (disaggregated).
12. Payments. At the end of the experiment, you will be paid your payoffs from 4 of the 20 rounds. These rounds will be randomly selected by the computer. The payoffs that you obtained in the selected 4 rounds will be converted into euros at the rate 100 ECUs = 1 euro and will be paid to you in private.

B.2 Experimental instructions of treatment *Simplified*


The aim of this experiment is to study how individuals make decisions in certain contexts. The instructions are simple. If you follow them carefully you will earn a non-negligible amount of money in cash (euros) at the end of the experiment. During the experiment, your earnings will be in ECUs (experimental currency units). Individual payments will remain private, as nobody will know the other participants' payments. Any communication among you is strictly forbidden and will result in immediate exclusion from the experiment.

1. The experiment consists of 40 rounds. In each round, you will be randomly assigned to a group of 2 participants (including yourself). This group is determined randomly at the beginning of the round. Therefore, the group you are assigned to changes at each round. In this room, there are 4 participants (including yourself) who are potential members of your group. That is, at every round, your group is selected among these 4 participants, each of them being equally likely to be in your group. You will not know the identities of any of these participants. In each round, you will only interact with the other participant in your group, and your payoff will only depend on your choice and the choice of the other participant in your group.
2. In each round, one of the two participants in your group will have the role of player 1 and the other one will have the role of player 2. The roles will be randomly assigned, and both participants in a group are equally likely to have each role assigned. At the beginning of the round, each participant will be informed of his/her assigned role.
3. At the beginning of the round, the computer randomly draws one object from an (virtual) urn containing two objects: one white object and one black object. Each object is picked with equal probability (50%).



The color of the object is revealed to player 1 but not to player 2 in your group.

4. At each round, each player is assigned a color. At the beginning of the round, the color assigned to each player is determined in the following way. There is one (virtual) urn for each player, containing two pieces of paper: one white and one black. The computer randomly (and independently) draws one piece of paper from each urn. In each urn, each piece of paper is picked with equal probability (50%). The piece of paper selected for each player determines that player's assigned color.



In each group, player 1 is informed both of his/her assigned color for the round and of the color assigned to player 2. Player 2 is only informed of his/her assigned color but not of the color assigned to player 1.

5. At each round, in each group, player 1 will be the first to make his/her decisions, knowing the color of the object drawn from the computer, his/her assigned color, and the assigned color of player 2. Player 1 has to decide what message to send to player 2 regarding the color of the object (which is unknown by player 2). The message can be either "*The object is white*" or "*The object is black*". Thus, the message can contain the true color or the false one.
6. Then, player 2, being informed of his/her assigned color (but neither of the color of the object nor of the color assigned to player 1), observes the content of the message sent by player 1 and decides which action to take: *Left*, *Center* or *Right*.
7. Round payoffs. At each round, the payoff to each player depends on whether the color of the object did or did not match his/her assigned color and on the action chosen by player 2:
 - i. At each round, the payoff to player 1 is determined as follows.
 - If the color of the object is equal to the color assigned to player 1, then his/her payoff depends on the choice of player 2 in the following way:
 - 20 ECUs if player 2 has chosen *Left*
 - 70 ECUs if player 2 has chosen *Center*
 - 120 ECUs if player 2 has chosen *Right*
 - If the color of the object is different from the color assigned to player 1, then his/her payoff is 20 ECUs, regardless of the action chosen by player 2.
 - ii. At each round, the payoff of player 2 is determined as follows.
 - If the color of the object is equal to the color assigned to player 2, then action *Left* provides him/her with a higher payoff than action *Center* or *Right*, and action *Center* provides him/her with a higher payoff than action *Right*.
 - If the color of the object is different from the color assigned to player 2, then action *Right* provides him/her with a higher payoff than action *Center* or *Left*, and action *Center* provides him/her with a higher payoff than action *Left*.

- The payoff of player 2 also depends on the correspondence between his/her assigned color and the color assigned to player 1: the payoff for player 2 when his/her assigned color is the same than the assigned color of player 1 is *lower* than in the case in which his/her assigned color is different from the color assigned to player 1.

The four tables below provide the payoffs of player 1 and player 2 in all possible situations:

- The top-left table corresponds to the cases in which both players have the same assigned color, which is different from the color of the object;
- The top-right table corresponds to the cases in which the color of the object is equal to the color assigned to player 2 but different from the color assigned to player 1;
- The bottom-left table corresponds to the cases in which the color of the object is equal to the color assigned to player 1 but different from the color assigned to player 2;
- The bottom-right table corresponds to the cases in which both players have the same assigned color, which is equal to the color of the object.

Object: White Player 2: Black Player 1: Black		or	Object: Black Player 2: White Player 1: White	
Action of player 2	Left	Center	Right	
Payoff to player 1	20	20	20	
Payoff to player 2	20	60	100	

Object: White Player 2: White Player 1: Black		or	Object: Black Player 2: Black Player 1: White	
Action of player 2	Left	Center	Right	
Payoff to player 1	20	20	20	
Payoff to player 2	160	120	50	

Object: White Player 2: Black Player 1: White		or	Object: Black Player 2: White Player 1: Black	
Action of player 2	Left	Center	Right	
Payoff to player 1	20	70	120	
Payoff to player 2	30	90	120	

Object: White Player 2: White Player 1: White		or	Object: Black Player 2: Black Player 1: Black	
Action of player 2	Left	Center	Right	
Payoff to player 1	20	70	120	
Payoff to player 2	60	50	40	

8. At the end of each round, prior to proceeding to the next round, all the players are informed about current and past rounds: assigned role, color of the object, color assigned to each player in the group, the message of player 1, the action chosen by player 2 and the payoff of each player.

9. Payments. At the end of the experiment, you will be paid the earnings that you obtained in 8 rounds (out of 40). These rounds will be randomly selected by the computer: 4 rounds will be selected from the rounds in which you were assigned the role of player 1 and the other 4 rounds will be selected from the rounds in which you were assigned the role of player 2. The earnings that you have obtained in the selected rounds will be converted into cash at the exchange rate of 40 ECUs = 1 euro and will be paid to you in private.

C) TREATMENTS WITH EXTRA OPTIONS TO ACQUIRE INFORMATION DIRECTLY (*Option* and *Unint-Opt*)

Treatment *Option* considers the following variant with respect to the *Base* treatment: in stage 2 of the game, after choosing which report to buy, if any, at the posted prices, buyers now have the option to directly acquire the information by paying 20 ECUs. This extra choice has no effect on the theoretical predictions regarding the equilibrium outcome (which remains the same as that of the *Base* treatment -see Section 3.1). However, it allows us to examine whether subjects acquire information directly less often in the first stage when they no longer need to worry about the possibility of facing excessively high prices to obtain information indirectly via the purchase of reports.

Finally, treatment *Unint-Opt* combines the features of treatment *Uninterested* and treatment *Option*. Because, in equilibrium, the option is never used, the theoretical predictions coincide with that of treatment *Uninterested* described above.

Here, we briefly report the results of treatments *Option* and *Unint-Opt* and compare them with the results of their respective counterparts without the option (i.e., with treatments *Base* and *Uninterested*).

Table 18 reports the (average) behavior by type of player in the auction in treatment *Option*.

Table 18. Average bids by type of player and block of 10 rounds - treatment *Option*

		<i>Rounds 1-10</i>	<i>Rounds 11-20</i>
<i>Informed Pl.</i>	<i>Color - Yes</i> [Prediction: 200]	107.61 (59.00)	142.54 (64.66)
	<i>Color - No</i> [Prediction: 100]	70.30 (45.29)	96.73 (46.29)
<i>Uninformed players</i> [Prediction: 150]		88.14 (53.43)	133.93 (67.66)
<i>Report Buyers</i>	Content: <i>Color - Yes</i> [Prediction: 200]	96.64 (54.24)	149.28 (44.95)
	Content: <i>Color - No</i> [Prediction: 100]	81.46 (43.25)	104.78 (24.44)
	Content: 0 [Prediction: 150]	41.57 (51.32)	63.44 (60.04)

(In brackets Std. Dev.)

Using Mann-Whitney tests, we test the differences between the players' bids in the *Base* treatment (see Table 2) and in treatment *Option* over all periods. We find that none of the differences are significant.

Table 19 reports the (average) behavior by type of player in the markets for information in treatment *Option*. We present the frequencies with which subjects choose to: (i) acquire information directly (column *Inform*), disaggregating the cases in which information was acquired in the first stage (first summand) or using the option (second summand);³³ (ii) buy a report in the second stage (column *Buy rep*); and (iii) remain uninformed (column *Uninf*), distinguishing by mover and by block of 10 rounds.

Table 19. Behavior in information markets in Treatment *Option* – Absolute number of observations

	<i>Rounds 1-10</i>			<i>Rounds 11-20</i>		
	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>
<i>Mover 1</i>	27 (45.00) [37.50]	21+1 (36.67) [27.50]	11 (18.33) [39.29]	43 (71.67) [45.26]	14+0 (23.33) [20.90]	3 (5.00) [16.67]
<i>Mover 2</i>	18 (30.00) [25.00]	30+2 (53.33) [40.00]	10 (16.67) [35.71]	22 (36.67) [23.16]	28+0 (46.67) [41.79]	10 (16.67) [55.56]
<i>Mover 3</i>	27 (45.00) [37.50]	25+1 (43.33) [32.50]	7 (11.67) [25.00]	30 (50.00) [31.58]	25+0 (41.67) [37.31]	5 (8.33) [27.78]
<i>Total</i>	72 (40.00) [100]	76+4 (44.44) [100]	28 (15.56) [100]	95 (52.78) [100]	67+0 (37.22) [100]	18 (10.00) [100]

(% over total row), [% over total column]

Using Mann-Whitney tests, we now test the differences between treatment *Base* (see Table 4) and treatment *Option* in the relative frequencies with which information is acquired (directly or indirectly) over all periods and aggregated across all movers. We find that no differences are significant. We also observe that the option is almost never used (only 4 times in the first block of 10 rounds).

³³ For example, 21+1 in the cell of Table 18 corresponding to players who are mover 1 in the first block of 10 rounds means that, within this block, mover 1 players directly acquired the information 22 times: 21 in the first stage and 1 using the option (i.e., after observing the reports' prices).

Table 20 reports the (average) behavior by type of player in the auction in treatment *Unint-Option*.

Table 20. Average bids by type of player and block of 10 rounds - treatment *Unint-Option*

	<i>Rounds 1-10</i>	<i>Rounds 11-20</i>
<i>Informed players</i>		
<i>Color - Yes</i> [Prediction: 200]	157.83 (65.78)	188.97 (34.28)
<i>Color - No</i> [Prediction: 100]	81.56 (37.63)	92.40 (23.68)
<i>Uninformed players</i>		
[Prediction: 150]	105.16 (71.70)	149.55 (49.22)
<i>Buyers of report</i>		
Content: <i>Color - Yes</i> [Prediction: 200]	154.87 (65.34)	177.64 (60.28)
Content: <i>Color - No</i> [Prediction: 100]	69.12 (30.35)	95.55 (37.45)
Content: 0 [Prediction: 150]	80.19 (53.86)	119.91 (54.85)

(In brackets Std. Dev.)

Using Mann-Whitney tests, we test the differences between the players' bids in the *Uninterested* treatment (see Table 14) and the *Unint-Option* treatment over all periods. We find that no differences are significant.

Table 21 reports the (average) behavior by type of player in the markets for information in treatment *Option*. We present the frequencies with which subjects choose to: (i) acquire information directly (column *Inform*), disaggregating the cases in which information was acquired in the first stage (first summand) or using the option (second summand); (ii) buy a report in the second stage (column *Buy rep*); and (iii) remain uninformed (column *Uninf*), distinguishing by mover and by block of 10 rounds.

Table 21. Behavior in information markets in Treatment *Unint-Option* – Absolute number of observations

	Rounds 1-10			Rounds 11-20		
	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>	<i>Uninf.</i>	<i>Inform.</i>	<i>Buy rep.</i>
<i>Mover 0</i>	32 (53.33)	28 (46.67)	---	46 (76.67)	14 (23.33)	---
<i>Mover 1</i>	32 (53.33) [29.36]	21+1 (38.33) [41.82]	5 (8.33) [31.25]	45 (75.00) [32.37]	12+0 (20.00) [40.00]	3 (5.00) [27.27]
<i>Mover 2</i>	46 (76.67) [42.20]	13+0 (21.67) [23.64]	1 (1.67) [6.25]	53 (88.33) [38.13]	5+0 (8.33) [16.67]	2 (3.33) [18.18]
<i>Mover 3</i>	31 (51.67) [28.44]	19+0 (31.67) [34.55]	10 (16.67) [62.50]	41 (68.33) [29.50]	13+0 (21.67) [43.33]	6 (10.00) [54.55]
<i>Total1-3</i>	109 (60.56) [100]	54+1 (30.56) [100]	16 (8.89) [100]	139 (77.22) [100]	30+0 (16.67) [100]	11 (6.11) [100]

(% over total row) [% over total column excluding Mover 0]

Using Mann-Whitney tests, we test the differences between treatment *Uninterested* (cf. Table 16) and treatment *Unint-Option* in the relative frequencies with which information is acquired over all periods by mover 0 and by all other three movers (aggregated across the three of them). We find that only one difference is marginally significant (at the 10% level), while the others are not significant.³⁴ We also observe that the option is almost never used (just once in the first block of 10 rounds).

³⁴ The only one that is weakly significant at the 10% level is the difference in the relative frequency with which movers 1 to 3 acquire information.