

Communication with Evidence in the Lab^{*}

Jeanne HAGENBACH[†] Eduardo PEREZ-RICHET[‡]

July 31, 2016

Abstract

We study communication with evidence in the lab. Our experimental design involves a collection of sender-receiver games with various payoffs and permits partial disclosure. We use local and global properties of the sender's incentive graph to uncover behavioral regularities and explain performance across games. Sender types whose interests are aligned with those of the receiver fully disclose, while sender types whose interests are not aligned with those of the receiver remain silent or partially disclose. When partially disclosing senders mostly disclose favorable pieces of evidence and hide unfavorable ones. But the cognitive cost of partial disclosure, as measured by response times, is higher for both senders and receivers. Receivers take evidence into account and tend to be skeptical about vague messages in games whose graph is acyclic. They perform better in acyclic games, whereas senders perform better in cyclic games.

Keywords: Sender-receiver game; hard evidence; information disclosure; masquerade relation; skepticism.

JEL classification: C72; C91; D82.

^{*}We thank Larbi Alaoui, Sourav Bhattacharya, Raicho Bojilov, Olivier Gossner, Guillaume Hollard, Asen Ivanov, Nicolas Jacquemet, Philippe Jehiel, Navin Kartik, Frdric Koessler, Laurent Lamy, Anol Llorente Sagner, Daniel Martin, Sri Srikandan, Olivier Tercieux, Tristan Tomala and Arne Uhendorff. We also thank seminar audiences at Ecole Polytechnique, Paris School of Economics, University Pompeu Fabra, University Autònoma de Barcelona, University of Essex, London School of Economics, HEC Lausanne and participants to the Cowles Foundation Economic Theory Conference and the 30th annual meeting of the European Economic Association. The authors are grateful to Investissements d'Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047), to the French National Research Agency (ANR, project DYNAMITE) and to the iCODE (Institute for Control and Decision, research project of the Idex Paris-Saclay), for financial support.

[†]CNRS – Ecole Polytechnique, e-mail: jeanne.hagenbach@polytechnique.edu

[‡]Ecole Polytechnique, e-mail: eduardo.perez@polytechnique.edu

1 Introduction

Exchanging information is an important part of any economic, political or social activity. Recent technological evolutions have impacted information transmission in two major ways. First, they have made information more easily verifiable, due to its wide availability and to the development of cross checking technologies. Second, they have made huge amounts of data available to economic agents, who can then control access to this data by other agents. Altogether, these evolutions make the study of information transmission through evidence as important as ever. They also call for policies and regulations that should be based on a theoretical and empirical understanding of communication behavior.

In this paper, we study communication with evidence in a collection of sender-receiver games in the lab. Empirical and experimental papers on disclosure in the literature focus on situations with monotonic incentives for the sender, such that the sender always wants to appear as having the highest possible informational type. By contrast, the games in our experiment span a variety of masquerading incentives for the sender. This is important because non monotonic incentives are omnipresent in economic situations. For example, think of the following situation. Company A and company B compete in the cell phone market, where A is dominant. Company B is the dominant player on the tablet market, where A is a potential entrant. A has just held its board meeting and it is known that they have taken a decision about entering the tablet market. This decision is private information but generates evidence that can be disclosed or hidden. If A has decided to enter, it would prefer B to believe that it has not, so B does not engage into preventive R&D on a new tablet. If A has decided not to enter, it would prefer B to believe it has decided to enter so B invests on R&D to develop a better tablet rather than better cell phones, and A is certain to keep its dominant position in the cell phone market. In this case, the incentives of the sender, company A , are cyclic. If instead company A and company B do not compete in the cell phone market, then A does not want to pretend it has decided to enter if it has not, and we are back to a situation with monotonic incentives. More generally, cyclic incentives are natural in all strategic situations with a Colonel Blotto flavor.¹ When having

¹In Colonel Blotto games, two players simultaneously allocate their resources across battlefields. On each

decided to attack on one front, a player then prefers the opponent to think she will attack on the other fronts. Non-monotonic incentives also arise naturally in situations in which the receiver does not know the direction of the sender’s bias,² or in situations where the sender is addressing a heterogeneous audience. In fact, in situations where the sender’s bias is fixed and known, incentives to masquerade are not necessarily monotonic either; the sender may want to appear as a slightly higher type but not as the highest. The payoffs of our experimental games include cyclic incentives as well as acyclic but non-monotonic incentives.

The incentives of the sender are best depicted as a graph, the masquerade relation, whose vertices represent the different informational states (types) of the sender, and whose oriented arcs represent masquerading incentives. Thus, there is an arc from type t_1 to type t_2 if t_1 wants to masquerade as t_2 , that is, if the sender is better off by convincing the receiver that her information is t_2 when it really is t_1 . An important local feature of the graph is whether the type of the sender has an incentive to masquerade as another type (an outgoing arc) or not. We call sender types with an incentive to masquerade *envious*, and others *satisfied*. An important global feature of the graph is whether it is acyclic or cyclic. One of the main contributions of this paper is to uncover how such local and global features of the graph explain behavior and performance of the players across different games in the lab. More broadly, our approach is to use the incentive graph to rigorously formulate and answer general questions about behavior in various games. We believe that this is a novel approach to experiments. We also contribute to the literature with an original experimental design in which the information of the sender is materialized as a set of two colored cards that she can observe, and reveal as she sees fit. The main advantages of this simple design are, first, that most subjects are familiar with card games, and, second, that it provides a natural way to allow for partial disclosure. Whenever the sender reveals one or two cards, the receiver can take the evidence into account by ruling out some of the possible informational types of the sender.

Our main results are as follows. First, we find that receivers are overwhelmingly consistent

field, the player with the higher force wins. The game has been applied to many economic problems such as research and development portfolio selection, political campaign resource allocation and auctions with simultaneous bidding.

²In Hagenbach et al. (2014), Example 2 details an example where a hidden bias creates a masquerading cycle.

with evidence both in beliefs and actions. That is, they report beliefs that do not put any weight on informational types that are ruled out by the evidence, and do not take actions that reflect inconsistent beliefs. This is important as it validates a restriction imposed by perfect Bayesian equilibrium, the usual solution concept for the analysis of disclosure games. In particular, receivers act optimally when information is fully revealed to them. Second, we find that satisfied senders, those with no incentive to masquerade as another type, fully disclose their information, whereas envious senders, those with an incentive to masquerade as another type, try to obfuscate their information by remaining silent or partially disclosing. When partially disclosing, they generally disclose information that is favorable to them. However, partial disclosure seems to be cognitively more costly as measured by response times. Third, we find that, in games where the masquerade is acyclic, the receivers tend to act skeptically. That is, when faced with a vague message, they tend to choose an action that weakly punishes all types of the sender that could have sent this message. However, they do sometimes fail to do so, thus giving the envious senders the opportunity to actually gain from obfuscation. The senders report beliefs that correctly anticipate skepticism from the receivers, but in some cases underestimate the level of skepticism.

Finally, we run regressions to explain the performance of envious senders and receivers across games. A receiver performs well when she acts optimally given the sender's true type, and a sender performs well when the receiver takes the action she prefers. We show that the existence of cycles affects these performances significantly, even controlling for the existence of a fully revealing equilibrium. The existence of cycles can be interpreted as a measure of complexity that favors the sender and penalizes the receiver. In acyclic games indeed, a simple heuristic can lead the receiver to a skeptical interpretation which enforces full revelation: (i) start by attributing the message to a possible sender type, (ii) ask which other possible types could benefit from this interpretation, (iii) if there is no such type, stick to the initial belief, otherwise attribute the message to one of the benefitting types and go back to (ii). This heuristic does not work in cyclic games even if there exists a fully revealing equilibrium. But other features of the masquerading graph can also be interpreted as complexity measures. In particular, we

find that the number of arcs in the graph, that is the number of masquerading incentives, has a negative impact on the performances of both senders and receivers. This suggests that with more masquerading incentives to track, it becomes more difficult for the receiver to figure out how to interpret vague messages, and for the sender to figure out how to obfuscate.

While our goal is not to test a particular theory, our results do to a certain extent validate some theoretical assumptions about communication with evidence, such as the fact that receiver take evidence into account. They also provide insights as to which equilibria are closer to actual behavior. For example, our results show that fully revealing equilibrium is most often a reasonable prediction in acyclic games, even though we document some small departures. But, even among fully revealing equilibria, theory is mostly mute about which particular messages will be used: senders could fully disclose, or use vague messages that will be correctly interpreted by receivers. Our results show that satisfied senders fully disclose, whereas envious ones use vague messages. Our observations can also be used to shape theory. Indeed, while there are well known selection criteria in signalling and cheap talk games, these criteria generally have no bite in hard information games, and the literature has not tried to formulate specific criteria. Our results do suggests such a refinement: by eliciting beliefs, we observe that receivers put little or no weight on satisfied senders. Thus a refinement imposing that receivers put less or no weight on satisfied types, both on and off the equilibrium path would be reasonable. This actually has the flavor of the D1 criterion, but contrary to D1, it is important to apply it both on and off path. It does have bite: for example it eliminates the fully revealing equilibrium in the cyclic games of our sessions 18 to 21.

Related Literature. While a significant amount of experimental papers³ test the predictions of cheap talk models (Crawford and Sobel, 1982), communication with evidence⁴ (Grossman, 1981; Milgrom, 1981) has given rise to a much thinner experimental literature. Early papers are mostly concerned with disclosure by a seller in a market environment. Forsythe et al. (1989)

³See for example Dickhaut et al. (1995); Blume et al. (1998, 2002); Cai and Wang (2006); Wang et al. (2010) for sender-receiver games, Battaglini and Makarov (2014) for a setup with multiple receivers, Vespa and Wilson (2015) for multiple senders, and Lai et al. (2015) for multidimensional cheap talk.

⁴A related model is Dye (1985), in which there is uncertainty about whether the sender is informed.

and King and Wallin (1991a) consider an informed seller, as in the models of Grossman (1981) and Milgrom (1981), and show that the unraveling principle works well as the players learn to fully disclose.⁵ King and Wallin (1991a) also observe that departures from full disclosure arise as the number of disclosure options to the sender increases, and there is no knowledge of these options by the buyer. In this case, receivers are insufficiently skeptical. Jin et al. (2015) run an experiment about quality disclosure in which they reach the different conclusion that the unraveling principle fails as lower types do not disclose, and receivers are too optimistic about sellers who choose not to disclose. However, this may be due to the fact that, even though players interact repeatedly, they do not get any feedback about their interaction. Benndorf et al. (2015) study an environment with multiple senders, interpreted as workers, with the option to disclose their productivity at a cost. There is no strategic receiver, but payoffs are designed to see how far disclosure of higher types pushes disclosure of lower types. Their subjects under-reveal, especially when of lower productivity.

All these experiments consider environments with monotonic masquerading incentives for senders: regardless of her true type, the sender wants to appear as having the highest possible type. By contrast, we are interested in understanding the behavior of senders and receivers in situations in which masquerading incentives can be more complex. This paper is rooted in a body of theoretical work that extends Grossman (1981) and Milgrom (1981) to richer setups. Seidmann and Winter (1997) look at communication with evidence with preferences as in Crawford and Sobel (1982) that lead to non-monotonic incentives, and introduce the notion of worst-case types, which gives an operational definition of skepticism in non-monotonic environments.⁶ Okuno-Fujiwara et al. (1990) generalize Milgrom (1981) to environments in which players exchange verifiable information before playing an incomplete information game, provide conditions for fully revealing equilibria to exist and examples in which full revelation is not an equilibrium. Van Zandt and Vives (2007) perform a similar exercise in supermodular games. Giovannoni and Seidmann (2007) show that there is a unique pooling equilibrium in

⁵See also Forsythe et al. (1989); King and Wallin (1991b) and Dickhaut et al. (2003) who add uncertainty about the seller's information as in Dye (1985).

⁶Mathis (2008) looks at conditions on the evidence structure that make it possible to extent the result of Seidmann and Winter (1997) to partial provability.

one-dimensional action games if the bias of the sender is strongly centripetal. Hagenbach et al. (2014) provide results that encompass all previous results in this literature, and introduce a new approach based on properties of the masquerade relation, i.e. the incentive graphs of informed players. Koessler and Perez-Richet (2015) and Perez-Richet and Tercieux (2015) use these ideas in the contexts of mechanism design and matching.

Our experimental approach also differs from previous papers in design and in scope. The particularities of our card-game design are that types are not ordered numbers, and that we allow for a natural partial disclosure technology. King and Wallin (1991a) also allow for partial disclosure through disclosure of given subsets of possible values.⁷ The difference in scope arises from the possibilities of our design, and our distinct theoretical anchor. When looking at partial disclosure, for example, we are not only interested in checking whether sender disclose according to the prediction of full revelation in acyclic games, as in King and Wallin (1991a), but also in whether senders try to obfuscate in ways that can benefit them if the receiver responds to evidence consistently. In addition, the variety of games that we use, enables us to focus on how local and global properties of the incentive graph affect behavior and performances across games.

The empirical literature on disclosure has also focused on environments with monotonic masquerading incentives. Dranove and Jin (2010) gives a nice overview of the literature on quality disclosure. Brown et al. (2012) look at cold openings in the movie industry, and observe that consumers fail to correctly infer the average quality of movies that are withheld from critics.

2 Theoretical Background

In the experiment, we consider simple sender-receiver games that can be described in the following framework. The sender has a finite set T from which her type is drawn according to a distribution $p(\cdot)$ with full support. The set of messages, or pieces of evidence, $M(t)$ to which

⁷In the papers that work with the Dye (1985) model, King and Wallin (1991b) and Dickhaut et al. (2003) also allow for partial disclosure.

she has access is contingent on her type. Hence a message m in $M = \cup_{t \in T} M(t)$ certifies that the type of the sender lies in the subset $M^{-1}(m) = \{t \in T : m \in M(t)\} \subseteq T$. A subset of types $S \subseteq T$ is *certifiable* if there exists a message m such that $M^{-1}(m) = S$. We assume *Own Type Certifiability*, that is, for every t , the singleton $\{t\}$ is certifiable.

The sender first chooses a message m to send to the receiver, and the receiver then chooses an action a from a finite set A after observing m . The payoffs of the sender and the receiver are given by $u_S(a, t)$ and $u_R(a, t)$. In this section and in the experiment, we focus on games such that, for every t , the utility function of the receiver is uniquely maximized by an action $a^*(t)$. All the games of the experiment satisfy the additional assumption that the receiver's payoff is positive if the receiver plays her optimal action and 0 otherwise. This is summarized in the following assumption. In the remainder of the paper, we refer to $a^*(t)$ as the *receiver optimal action*, or *optimal action* when the context is clear.

Assumption 1. *For every t , the receiver optimal action $a^*(t)$ is the unique action such that $u_R(a^*(t), t) > 0$, and, for every $a \neq a^*(t)$, $u_R(a, t) = 0$.*

We refer to any message that certifies T as *silence* since it provides no information. If the sender certifies a singleton, we say that her message is *fully disclosing*. In intermediate situations, we talk about *partial disclosure*.

Masquerade Relation. In games of communication with evidence, it is useful to introduce the *masquerade relation* as in Hagenbach et al. (2014). We say that type t wants to masquerade as type $s \neq t$, and write $t \rightarrow s$, if

$$u_S(a^*(s), t) > u_S(a^*(t), t).$$

This relation is best pictured as an incentive graph whose vertices are the elements of T , and where there is an oriented arc between t and s if $t \rightarrow s$. The evidence structure can be represented on the same graph, by drawing all certifiable subsets. This kind of representation is illustrated in [Figure 1](#), in which S_1 , S_2 and T are the certifiable subsets.⁸

⁸Note that this evidence structure does not satisfy the Own Type Certifiability assumption we make in the




Figure 1: A masquerade and evidence structure graph.

We say that a type t is *satisfied* if there is no other type s that t wants to masquerade as. We refer to non satisfied types as *envious*. We say that a sequence of types t^1, \dots, t^k forms a *masquerading cycle* if, for every $1 \leq \ell \leq k$, $t^\ell \rightarrow t^{\ell+1}$, where, by convention, $t^{k+1} = t^1$. A masquerade relation is cyclic if it admits a masquerading cycle, and acyclic otherwise. By extension, we talk about cyclic and acyclic games.

Worst-case types, skepticism and acyclicity. The notion of *worst-case type*, first introduced by Seidmann and Winter (1997), plays an important role in enforcing revelation. A worst-case type of a subset $S \subseteq T$ is a type $s \in S$ that no other type $s' \in S$ wants to masquerade as. Graphically, it is a vertex in S with no incoming arc from other types in S . The set of worst-case types of S is denoted by

$$\text{wct}(S) = \{s \in S : t \rightarrow s \Rightarrow t \notin S\}.$$

Worst-case types are key to full revelation as they provide an operational definition of skepticism. Indeed, attributing vague messages to worst-case types allows the receiver not to reward obfuscation attempts while remaining consistent with the evidence. However, worst-case types need not exist. In Figure 1, for example, the certifiable sets T and S_1 admit a worst-case type, but S_2 does not. This is because the types in S_2 form a masquerading cycle. In fact, the existence of worst-case types, and hence the possibility of skepticism, is crucially related to the

rest of the paper.

absence of masquerading cycles, as shown by the following lemma.

Lemma 1. *Every subset of types $S \subseteq T$ admits a worst-case type if and only if the masquerade relation is acyclic.*

Strategies and equilibrium concepts. In this section we consider only pure strategy equilibria. Our results can be extended to mixed strategies for all games in the experiment, but to show that, we need to introduce another assumption that is satisfied by all games in the experiment (see [Appendix A](#)). A pure messaging strategy of the sender is a mapping $\mu : T \rightarrow M$ such that, for every t , $\mu(t) \in M(t)$. A strategy of the receiver is a mapping $\alpha : M \rightarrow A$. Denote by $\beta_m \in \Delta(t)$ the belief of the receiver following message m , and let $\beta = (\beta_m)_{m \in M}$. Our equilibrium concept is perfect Bayesian equilibrium (PBE). A triple (μ, α, β) forms a PBE if, for every $m \in M$ and every $t \in T$, it satisfies

$$\alpha(m) \in \arg \max_a \sum_{t \in T} u_R(a, t) \beta_m(t) \quad (\text{receiver optimality})$$

$$\mu(t) \in \arg \max_{m \in M(t)} u_S(\alpha(m), t) \quad (\text{sender optimality})$$

$$m \in \cup_{t \in T} \mu(t) \Rightarrow \beta_m(t) = \frac{p(t)}{\sum_{s: \mu(s)=m} p(s)} \mathbb{1}_{\mu(t)=m} \quad (\text{Bayesian rationality})$$

$$\beta_m(t) > 0 \Rightarrow t \in M^{-1}(m) \quad (\text{evidence consistency})$$

For the remainder of the paper, it is useful to extend the idea of consistency to receiver strategies. We say that a strategy α of the receiver is a *consistent strategy* if for every message $m \in M$, we have $\alpha(m) = a^*(t)$ for some type $t \in M^{-1}(m)$. By [Assumption 1](#), this is equivalent to $\alpha(m)$ being optimal for some consistent belief β_m . Otherwise, we say that α is inconsistent. Similarly, we say that action a is *consistent with message m* if $a = a^*(t)$ for some type $t \in M^{-1}(m)$. We refer to a receiver taking an inconsistent action (consistent, resp.) as a receiver *ignoring the evidence* (*taking the evidence into account*, resp.). The consistency of receiver strategies will prove central for the theoretical observations of the next section.

Fully revealing equilibria. We say that a PBE is a fully revealing equilibrium (FRE) if the strategy of the sender is separating, that is $\mu(t) \neq \mu(t')$ if $t \neq t'$. Skepticism plays an important role for the existence of FRE. Indeed, results in Hagenbach et al. (2014) imply that the existence of a worst-case type is necessary and sufficient for the existence of a FRE.

Proposition 1. *Under Assumption 1 and Own Type Certifiability, a sender-receiver game admits a FRE if and only if every certifiable subset of types admits a worst-case type.*

By Lemma 1, we have the following corollary, relating the existence of FRE to acyclicity of the masquerade relation.

Corollary 1. *Under Assumption 1, a sender-receiver game admits a FRE for every evidence structure that satisfies Own Type Certifiability if and only if it is acyclic.*

This implies that all acyclic games in our experiment have a FRE. But cyclic games may or may not have a FRE depending on the evidence structure.⁹

3 Experimental Design

Our design consists of a collection of sender-receiver games. In each of these games, the sender observes two cards, each of which is either yellow, Y , or blue, B . She chooses which cards to reveal, if any. The receiver observes the cards revealed by the sender, and then takes a decision a , b or c . Revealing evidence in this way is novel and has several advantages. First, and most importantly, it is a neat and simple way to allow partial disclosure. Second, it avoids the caveat of giving numerical values to types, which might frame behavior. Finally, most people are familiar with card games which makes the rules easy to understand.

Upon their arrival to the lab, subjects were given written instructions and had to fill a comprehension test before taking any decision.¹⁰ Subjects then played 20 or 30 rounds of a

⁹Consider the masquerade relation $t_1 \rightarrow t_2 \rightarrow t_3 \rightarrow t_1$ with certifiable subsets $\{t_1\}$, $\{t_2\}$, $\{t_3\}$, $\{t_1, t_2\}$, $\{t_2, t_3\}$, and $\{t_3, t_1\}$. Full revelation is an equilibrium as every certifiable subset has a worst-case type. However, adding the possibility of remaining silent, which certifies $\{t_1, t_2, t_3\}$, prevents existence of a FRE.

¹⁰The english translation of the instructions and comprehension test are available in the supplementary appendix.

fixed sender-receiver game preceded by two rounds of practice. All decisions were taken on computers and screenshots are provided in [Appendix E](#). At the beginning of an experimental session, each subject was randomly assigned a role, labeled as *player 1* (sender) or *player 2* (receiver), for the whole session. At the end of each round, both players in every pair were given full feedback (cards of the sender, actions and payoffs) about their interaction. Senders and receivers were then randomly rematched for the next round.¹¹ The final payment of a subject was the sum of a fixed show up fee, a variable payment based on her total score in the sender-receiver games she played and, for some sessions, a payment for a belief-elicitation task. For each session, [Table 3](#) gives the number of subjects, the number of rounds played and the value of a point in euros. The last column of the Table states whether beliefs were elicited or not.

Information and messages. The type set corresponding to our design is given by the feasible combinations of cards $T = \{YY, YB, BB\}$. The three combinations were drawn with equal probabilities by the computer and independently in each round. The sender was allowed to reveal both cards, none or any one of her choice. The corresponding evidence structure, which satisfies Own Type Certifiability, is described in [Figure 2](#). Note that among all subsets of types, $\{BB, YY\}$ is the only one which is not certifiable. A subject *fully discloses*, abbreviated as *FD*, if she reveals both cards, *partially discloses*, abbreviated as *PD*, if she reveals one card, and remains *silent*, abbreviated as \emptyset , if she reveals no card. We also refer to silence or one-card messages as *vague messages*.




Figure 2: *Evidence structure of the experiment.*

¹¹We did not use a perfect stranger matching process, so rematch was in principle allowed. On average, over all sessions, a given pair of subjects was formed 3.9 times. Each combination of two subjects and two cards was formed 2.3 times on average.

Payoffs and masquerade relations. In each session, subjects could at all time see two matrices showing sender and receiver payoffs in the game they were playing. The payoff matrix of the receiver was the same across all sessions, whereas the payoff matrix of the sender was the unique part of the design that varied across sessions. As an illustration, [Table 1](#) shows the payoff matrix of the sender in sessions 4, 5 and 6. [Table 2](#) shows the payoff matrix of the receiver for all sessions. Note that for every couple of type and action, players payoff always equals either 0 or 3.

	<i>a</i>	<i>b</i>	<i>c</i>
<i>YY</i>	0	3	0
<i>YB</i>	0	0	3
<i>BB</i>	0	0	3

Table 1: *Sender Payoff Matrix*

	<i>a</i>	<i>b</i>	<i>c</i>
<i>YY</i>	3	0	0
<i>YB</i>	0	3	0
<i>BB</i>	0	0	3

Table 2: *Receiver Payoff Matrix*

To each of the games corresponds a particular masquerade relation, and [Figure 3](#) shows all the masquerading graphs that we used. For example, the masquerading graph that corresponds to [Table 1](#) is the Line. [Table 3](#) shows the masquerading graph and sender payoffs associated to each of the 21 experimental sessions. Over all these sessions, we collected data for 3690 one-shot sender-receiver games.¹²

Theoretical observations. [Table 4](#) shows the distribution of acyclic and cyclic games. As explained in the theory section, there is a difference between acyclic and cyclic games regarding the existence of fully revealing equilibria. By [Corollary 1](#), all acyclic games of the first ten sessions have a FRE. According to [Proposition 1](#), there also exists a FRE in cyclic games of sessions 18 to 21 as every certifiable set of types has a worst-case type (the set $\{YY, BB\}$ has no worst-case type, but it is not certifiable). On the contrary, cyclic games of sessions 11 to 17 have no FRE. All games have at least one equilibrium which involves some pooling. For each session, [Table 3](#) reports whether the game is cyclic or acyclic and whether it has an FRE.

¹²[Appendix B](#) discusses the selection process of our masquerading graphs and reports 4 additional control sessions.

Session	Masquerade	Sender Payoff	Game Type	FRE	Rounds	Subjects	Exch. Rate	Beliefs
1		$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 3 & 0 \end{pmatrix}$	Acyclic	Yes	20	18	0.20	No
2		$\begin{pmatrix} 0 & 0 & 3 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.25	No
3		$\begin{pmatrix} 3 & 0 & 0 \\ 3 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.20	No
4		$\begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.20	No
5		$\begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.25	Yes
6		$\begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.25	Yes
7		$\begin{pmatrix} 0 & 3 & 3 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	18	0.25	No
8		$\begin{pmatrix} 0 & 3 & 3 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	20	0.25	Yes
9		$\begin{pmatrix} 0 & 3 & 3 \\ 0 & 0 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	14	0.25	Yes
10		$\begin{pmatrix} 3 & 0 & 0 \\ 3 & 0 & 0 \\ 3 & 3 & 0 \end{pmatrix}$	Acyclic	Yes	20	16	0.25	No
11		$\begin{pmatrix} 0 & 3 & 0 \\ 3 & 0 & 0 \\ 0 & 0 & 3 \end{pmatrix}$	Cyclic	No	20	20	0.66	No
12		$\begin{pmatrix} 0 & 3 & 0 \\ 3 & 0 & 0 \\ 0 & 0 & 3 \end{pmatrix}$	Cyclic	No	20	18	0.25	Yes
13		$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 3 & 0 \end{pmatrix}$	Cyclic	No	30	10	0.20	No
14		$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 3 & 0 \end{pmatrix}$	Cyclic	No	20	14	0.25	Yes
15		$\begin{pmatrix} 0 & 3 & 0 \\ 3 & 0 & 3 \\ 0 & 3 & 0 \end{pmatrix}$	Cyclic	No	20	20	0.25	No
16		$\begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	No	30	10	0.20	No
17		$\begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	No	20	16	0.25	No
18		$\begin{pmatrix} 0 & 0 & 3 \\ 0 & 3 & 0 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	Yes	30	10	0.20	No
19		$\begin{pmatrix} 0 & 0 & 3 \\ 0 & 3 & 0 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	Yes	20	20	0.25	No
20		$\begin{pmatrix} 0 & 0 & 3 \\ 0 & 3 & 0 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	Yes	20	16	0.25	Yes
21		$\begin{pmatrix} 0 & 0 & 3 \\ 0 & 3 & 0 \\ 3 & 0 & 0 \end{pmatrix}$	Cyclic	Yes	20	14	0.25	Yes

Table 3: *Summary of sessions.*




Figure 3: *Masquerading graphs of the experiment.*

n	Envious Sender	Satisfied Sender	Total
Acyclic Game	1039	821	1860
Cyclic Game	1372	458	1830
Total	2411	1279	3690

Table 4: *Distribution of the data across game and sender categories.*

Table 4 also gives the number of games in which the sender was of an envious or satisfied type. We now provide a set of observations that will make clear how these types affect, in theory, the disclosure behavior of senders. More generally, these observations will guide our empirical analysis of the data. We start with a statement about receiver's behavior. It relies on the fact that, in our games, there exists a unique optimal action of the receiver for each type of the sender, and all other actions yield a payoff of zero ([Assumption 1](#)).

Observation 1. *For the receiver, any inconsistent strategy is weakly dominated.*

Indeed, a receiver who ignores evidence in m will obtain a zero payoff, while there exists a strategy that accounts for m and would yield a positive payoff. Hence, consistency is more than a mere reasonable equilibrium requirement in our games. [Observation 1](#) implies that

strategically sophisticated senders should expect receivers to take evidence into account, and disclose accordingly. The next observation concerns the behavior of senders of satisfied type. In our games, every satisfied type t gets a payoff 3 if the receiver takes action $a^*(t)$ and 0 otherwise.

Observation 2. *Provided that the receiver only uses consistent strategies, it is a weakly dominant strategy for satisfied types to fully disclose.*

Receiver consistency also puts a restriction on the kind of disclosure strategies that envious senders should use. To clarify this, we define a notion of *advantageous obfuscation*. Intuitively, an envious sender obfuscates advantageously if she uses a message that can lead the receiver to choose an action that she strictly prefers to the one induced by full disclosure. Formally, given a messaging strategy $\mu : T \rightarrow M$ of the sender, we say that type t *obfuscates advantageously* if there exists $s \in M^{-1}(\mu(t))$ such that $t \rightarrow s$. In the games of our experiment, every envious sender has a simple way to obfuscate advantageously: staying silent.

Observation 3. *Provided that the receiver only uses consistent strategies, any strategy μ such that an envious type t does not obfuscate advantageously is weakly dominated by a strategy μ' such that the type t obfuscates advantageously.*

This is true because, under μ' , there is at least one consistent strategy of the receiver such that she will take a decision that type t envies (that is, a decision $a^*(s)$ for an s such that $t \rightarrow s$) whereas the best outcome for type t under μ is if the receiver chooses $a^*(t)$.¹³ An interesting corollary is that envious types should never fully disclose in our games.

Observation 4. *Provided that the receiver only uses consistent strategies, any strategy μ such that an envious type t fully discloses is weakly dominated by a strategy μ' such that the type t obfuscates advantageously.*

¹³Observation 3 does not mean that disadvantageous obfuscation cannot exist in equilibrium. Consider, for example, the game of session 4 in the experiment. Suppose the sender fully discloses if her type is BB , to which the receiver responds by c , and shows a yellow card if her type is either YY or YB , to which the receiver responds by b . Suppose in addition that the receiver responds to silence by choosing a , and to a blue card by choosing b . This is an equilibrium even though YB obfuscates disadvantageously.

Finally, going back to the link between FRE and skepticism, we can note that skepticism is a necessary feature of FRE. Since we do not observe beliefs, we say that a receiver *acts skeptically* if her action following a vague message m corresponds to the optimal action $a^*(t)$ associated with a worst-case type of that message, that is, with $t \in \text{wct}(M^{-1}(m))$. This must hold in any fully revealing equilibrium, both for on path and off path messages. Such a choice of action is supported by any consistent belief that puts more probability on t than on any other type in $M^{-1}(m)$.

Observation 5. *A PBE is fully revealing if and only if the receiver acts skeptically following any (on-or-off-path) message of the sender.*¹⁴

Therefore a measure of the extent to which receivers act skeptically can be interpreted as a measure of how well FRE explains the data, on the dimension of receiver behavior, and in games that have a FRE.

Belief elicitation. We elicited the beliefs of all subjects in 8 of the 21 sessions (4 with acyclic games, 4 with cyclic games, 136 subjects in total). In each of these sessions, we elicited beliefs once, after the final round of play. Each sender was asked to indicate the frequency with which she thought that each action was played after each possible message. Each receiver was asked to indicate the frequency with which she thought that each card combination had been the true one given each possible message. The tables used for beliefs elicitation are given in [Appendix E](#) and were filled on the computers. We then compared the reported frequencies to the empirical frequencies in the session, and remunerated the players in points according to the scoring rule

$$s(f, f^e) = 16 \left(1 - \frac{1}{18} \sum_{k=1}^{18} \delta_k \left(\frac{f_k - f_k^e}{100} \right)^2 \right),$$

where f_k and f_k^e were, respectively, the believed and empirical frequencies in square k of the belief elicitation tables. If a message had never occurred in a session, the indicator δ_k was set to 0 for the squares corresponding to this message. A subject whose beliefs matched the empirical

¹⁴This is a consequence of the proof of [Proposition 1](#). Intuitively, if any message is followed by a non skeptical action, there is a type of the sender that will deviate to sending this message.

frequency exactly could earn a maximum of 16 points (equivalent to 4 euros). The scoring rule was given in the written instructions together with an explanation that it was in each subject's interest to declare her true beliefs about the frequencies.

Implementation. The experiments were conducted at the Experimental Laboratory in Ecole Polytechnique, France, with a total of 354 subjects who were mostly science students from Ecole Polytechnique and ENSTA, two French engineering schools, and some administrators at Ecole Polytechnique. The average payoff per subject was 15 euros plus a 5 euros show-up fee. For the sessions in which beliefs were elicited, subjects could earn an additional payoff of 4 euros at most.

4 Gains from Communication

To get a general sense of the data, we start by looking at the average scores of senders and receivers, as described on [Figure 4](#). In every game, the receiver gets a payoff of 3 if she takes her unique optimal action, and 0 otherwise, and each of her three possible actions is equally likely to be optimal at the outset. Therefore, her average score in the absence of communication would be 1. Hence, we say that receivers gain from communication whenever their average payoff is above 1. The average payoff of the sender if the receiver chooses her action according to the uniform distribution is between 1 and $4/3$, depending on the game. This range is depicted as the red hatched region in [Figure 4](#). We say that senders gain from communication if their average payoff is above $4/3$, and lose from communication if their average payoff is below 1.

According to these definitions, both senders and receivers gain from communication overall. Satisfied senders gain, but envious senders lose from communication. The losses of envious senders come mainly from acyclic games, in which they obtain a payoff of 0.35 on average against 1.15 in cyclic games. Receivers gain from communication regardless of the precision of the message they receive, and of cyclicity. But their scores are lower after vaguer messages, and in cyclic games.

The fact that receivers obtain an average payoff close to 3 when they face full disclosure




Figure 4: *Sender and receiver average payoffs.*

confirms that the subjects understood the games. The fact that satisfied senders obtain an average payoff close to 3 both in cyclic and acyclic games shows that subjects managed to convey information effectively when their interests were aligned with those of the senders.

As the theory suggests, acyclic games favor receivers and harm envious senders. The poorer performance of receivers when facing vague messages, even in acyclic games, however, is not compatible with FRE. Indeed, while FRE does not preclude the use of vague message, it implies that the receivers should guess the type of the sender anyway.

5 Receiver Behavior

Do receivers take evidence into account? PBE requires consistent receiver strategies. In addition, inconsistent strategies are weakly dominated in our games ([Observation 1](#)). To measure consistency in the data, we look at the rate of consistent receiver decisions. A decision taken by the receiver after seeing message m is consistent if this decision is optimal for some type who could have sent m . Overall, excluding the case of silence, after which any choice of the receiver is necessarily consistent, we find that more than 98% of receivers decisions are consistent whether facing FD or PD, and whether the game is cyclic or acyclic.

How accurate are receivers' decisions? We say that a receiver's decision is accurate if it corresponds to the optimal action associated with the true sender type. Overall, decisions are accurate 71% of the time, confirming that receivers benefit from communication with evidence. The accuracy rate jumps to 80% in acyclic games, which is close to the prediction of FRE. In cyclic games, receivers achieve an accuracy rate of only 62%. Following full disclosure, receivers have an accuracy rate of 99% over all games. Vaguer messages lead to lower accuracy rates, irrespective of cyclicity. In acyclic games, the accuracy rates following partial disclosure is 76%, and it is 58% after silence. A receiver who would choose randomly with uniform weights after considering the evidence would obtain accuracy rates of 50% after partial disclosure, and 33% after silence. These rates are below the performance of receivers in acyclic games, but comparable to their performance in cyclic games, with 54% after partial disclosure, and 44% after silence. If we isolate cyclic games with a FRE (sessions 18 to 21), we obtain accuracy rates of 60% following partial disclosure, and 51% following silence. Interestingly, while receivers do perform better in these games than in other cyclic games, they remain lower than their performance in acyclic games despite the existence of a FRE.

Are receivers skeptical? Following [Observation 5](#), we can interpret the extent to which receivers act skeptically as a measure of how well FRE explains the data, on the dimension of receiver behavior. In games with no FRE, skepticism is not possible after some messages, and generally not a necessary feature of equilibrium. [Figure 5](#) shows the frequency of skeptical actions. Following full disclosure, receivers overwhelmingly choose the correct action, and therefore act skeptically, so we exclude full disclosure from the reported results. We also eliminate observations for which the receiver does not have the option of acting skeptically because the message she faces has no worst-case type, as might happen in cyclic games. This leaves 1641 data points.

When receivers face silence or partial disclosure, they act skeptically 60% of the time. In acyclic games, receivers act skeptically 79% of the time. Hence FRE fits the data relatively well on the dimension of skepticism. However, our results indicate that, in acyclic games,




Figure 5: *Frequency of skeptical choices excluding full disclosure.*

receivers tend to be less skeptical following silence (73%) than partial disclosure (90%), which is a departure from FRE.

In cyclic games with no FRE, acting skeptically is impossible following the messages which have no worst-case type. When it is possible, it is often not a necessary feature of equilibrium. Even in cyclic games with a FRE, skepticism may be unnatural. For example, consider silence in the game of sessions 18 to 21. There is a worst-case type for this message, which is YB , and playing the corresponding action is a necessary feature of FRE in this game. Yet, the belief that the silent message indeed comes from YB is unnatural: it is weakly dominant for YB to fully disclose once inconsistent strategies of the receiver are eliminated, whereas BB and YY may gain from the response of the receiver to the silent message.¹⁵

In cyclic games, we find that the share of skeptical responses is only 28%. But the share of skeptical responses is not very informative as skepticism plays no particular role in cyclic games in general. It is more interesting to focus on sessions 18 to 21, which have cycles and a FRE. In these games, as explained above, skepticism following silence is unnatural even though it is

¹⁵The reasoning is reminiscent of the intuitive criterion or D1, but this criterion does not eliminate all FRE. First think of the following FRE: all types fully disclose, and the receiver attributes any vague message to YB or puts a higher probability on YB than any other type. This FRE does not satisfy D1, since, if we consider the silent message, for example, YY and BB stand to benefit from any response of the receiver whereas YB stands to lose from any action other than b . There is another FRE, however, which is not eliminated by D1. Suppose indeed that BB and YY both mix between full disclosure and partial disclosure, while YB mixes between silence and full disclosure. In this case every message is on the equilibrium path, and the equilibrium is separating. But, while this FRE cannot be eliminated by D1, it does not feel like a good prediction, and indeed makes use of weakly dominated strategies.




Figure 6: Average response time of receivers.

a necessary response to enforce FRE. In sessions 18 to 21, receivers act skeptically only 8% of the time after silence, so FRE fits the data quite poorly for this game.

Response times. Receivers’ response times provide two interesting observations. First, the response time in cyclic games is higher than in acyclic games whenever the message is vague, which confirms the intuition that cyclic games require more thinking on the receiver side. Second, response times are non-monotonic in the amount of disclosure: they are lowest for full disclosure, and highest for partial disclosure. On the face of it, this observation is slightly surprising as receivers face more possible responses after a silent message, than after partial disclosure. It suggests that deciphering the intentions of the sender following partial disclosure requires more thinking. This increased difficulty, however, does not translate into more mistakes, as the skepticism and accuracy data show.

Beliefs and best responses. We elicited subjects beliefs in 8 sessions (see Table 3). Following full disclosure, receivers overwhelmingly reported beliefs that put probability 1 on the disclosed type: the average mass receivers put on the correct type was above 0.99 for each of the FD messages. Their average beliefs following vague messages are given session by session in Appendix D, Table 8, as well as the true empirical frequencies. The first remarkable feature of these elicited beliefs is that they are fully consistent with evidence for PD messages, as they

never put any weight on the type of the sender that is ruled out by evidence. The second remarkable feature is that, even though belief accuracy is hard to measure in any meaningful way, average reported beliefs never seem to be too far off the empirical frequencies.¹⁶ Third, elicited beliefs allow us to take a finer look at the question of skepticism in acyclic games. In acyclic games, average beliefs almost always put the highest mass on a worst-case type,¹⁷ but they are far from putting all the mass on worst-case types. Finally, for all messages that contain a satisfied type, the average belief mass that receivers put on this type is small: after silence, the average weight put on the satisfied type is less than 1%; after partial disclosure, there is always less weight on the satisfied type than on the envious one. In sessions 20 and 21, where the game is cyclic but has a FRE, the worst case type of silence is a satisfied type. The average mass receivers put on this type following silence is close to 0. This confirms the point we made about unnatural skepticism in these games.

Next, we used elicited beliefs to see whether receivers best responded to their reported beliefs. To that purpose, we created the indicator BR-b which, for every observation, takes value 1 if the receiver best responded to her reported beliefs in the period and the session corresponding to the observation, and 0 otherwise. We also created the corresponding indicator, BR-e, for best responses to the true empirical frequencies of the session. These two indicators allow us to decompose behavior into observations for which the receiver best responded to both her reported beliefs and the true frequencies (BR-b AND BR-e), situations in which she failed to best respond to her own beliefs (Not BR-b), and situations in which she did best respond to but was lead astray by her beliefs (BR-b AND Not BR-e). The corresponding rates when facing vague messages¹⁸ are given for all sessions in Figure 7. The rate at which receivers failed to best respond to their belief is between 20% and 30% for most sessions, except sessions 20 and 21 for which it is lower. By contrast, the rate at which receivers best responded to inaccurate beliefs varies more widely across sessions corresponding to different games. This rate is higher

¹⁶The only notable exception concerns the B message in session 12. This message was used only 8 times over the 180 observations for this session.

¹⁷The only exception is following \emptyset in session 9.

¹⁸When facing FD, receivers best responded to beliefs and true frequencies at a rate close to 100% in all sessions.




Figure 7: Best-response rates for receivers facing with vague messages. *BR-b and BR-e* measures the average rate at which receivers best-responded to the elicited beliefs and to the empirical frequencies. *BR-b and not BR-e* measures the average rate at which receivers best-responded to the elicited beliefs but not to the empirical frequencies. *Not BR-b* measures the rate of failure of best-response to the elicited beliefs.

in sessions with cyclic games.

6 Sender Behavior

What do senders disclose? The disclosure strategies of senders are illustrated in Figure 8.

Unsurprisingly, satisfied types overwhelmingly use full disclosure whether the game is cyclic or acyclic, which validates Observation 2. Envious senders, on the other hand, understand that they need to obfuscate in order to stand a chance of obtaining their first-best payoff, as their empirical behavior validates Observation 4. Indeed, their frequency of full disclosure is only 5%. In acyclic games, envious types do not use full disclosure more frequently than in cyclic games (respectively 5.1% and 4.5%). This is not inconsistent with FRE in acyclic games, and shows that envious senders try to take advantage of receivers' possible mistakes: if receivers sometimes fail to be skeptical following vague messages (as the data indicates), it is better to try an obfuscation strategy. Finally, envious senders use silence more often than partial disclosure (respectively 68% and 32%, conditional on using vague messages). The higher frequency of silence is observed both in cyclic and acyclic games.




Figure 8: *Sender disclosure.*

Do envious senders obfuscate to their advantage? Observation 3 shows that rational senders can be expected to obfuscate advantageously. This observation implies, in particular, that envious type should not fully disclose (see Observation 4). We have already shown that envious senders use vague messages 95% of the time. But it does not imply that they obfuscate advantageously. In order to look at this, we isolate the situations in which an envious type has the choice between obfuscating advantageously or not when using partial disclosure. These situations occur for type YB of the sender in sessions 4 to 14, 16 and 17 and they are depicted on the left panel of Figure 9. In any of these situations, obfuscating with a red (dashed) message is disadvantageous for type YB as the set of types certified includes no type that YB wants to masquerade as. On the contrary, obfuscating with a green (non dashed) message is advantageous as it certifies a set of types which includes a type that YB wants to masquerade as. Note that silence is advantageous in all situations. The right panel of Figure 9 shows the frequency with which each message is used. Overall, senders obfuscate advantageously 92% of the time. The rate of disadvantageous obfuscation of 12.2% in acyclic games is relatively important, and suggests that partially disclosing may be a complicated task for envious senders.




Figure 9: *Obfuscation by senders of type YB. Red (dashed) messages correspond to disadvantageous obfuscation. Green (non dashed) messages correspond to advantageous obfuscation. The graph on the right panel shows the frequency with which silence, advantageous PD and disadvantageous PD are used. The numbers indicate the size of the data.*

Response times. Unsurprisingly, response times of envious senders are considerably higher than those of satisfied senders. In addition, the response times of envious senders exhibit the same non-monotonicity in the amount of disclosure as those of receivers: they are highest for partial disclosure, and lowest, and similar, for silence and full disclosure. Envious senders rarely fully disclose. The observation that partial disclosure requires more thinking time than silence is interesting.¹⁹ It suggests a higher cognitive burden, which is compatible with the significant rate of disadvantageous partial disclosure that we observe. It might explain that envious senders use silence more often than partial disclosure while partial disclosure leaves the receiver with fewer consistent actions than silence.

Beliefs and best responses. We elicited subjects beliefs in 8 sessions (see Table 3). Senders overwhelmingly reported beliefs that put probability 1 on receivers choosing her optimal action if they fully disclosed: the average mass they put on the correct action was above 0.97 for each of the FD messages. Their average beliefs about the actions the receiver would take after vague messages are given session by session in Appendix D, Table 9, as well as the true empirical

¹⁹A possible explanation could be that, for the YB types, there are two ways of partially disclosing, but it is invalidated by the fact that other types of envious senders exhibit the same profile of response times. An alternative explanation is that silence is always an advantageous way to obfuscate while partial disclosure may not be.




Figure 10: Average response time of senders.

frequencies. The senders put very little weight on the receiver choosing inconsistent actions following PD messages.²⁰ Elicited beliefs are mostly quite accurate with some remarkable mistakes. Senders seem to systematically underestimate the degree of skepticism of the receivers following the Y message in acyclic games. Even though they underestimate the degree of skepticism for the Y message, senders seem to correctly anticipate a skeptical behavior of the receiver: their average beliefs almost always put the highest mass on the receiver choosing a skeptical action following a vague message.²¹ In particular they do correctly anticipate the behavior of receivers following the B message. The other mistakes of senders occur in cyclic games and seem less important.²²

As for receivers, we used elicited beliefs to see whether senders best responded to their reported beliefs and empirical frequencies. To that purpose, we created the indicator BR-b which, for every observation, takes value 1 if the sender best responded to her reported beliefs in the period and the session corresponding to the observation, and 0 otherwise. We also created

²⁰The maximum average reported rate on an inconsistent action is 11.9% in session 20. If we exclude session 20, it is never above 1.7%.

²¹Beliefs in session 9 following the Y message is an exception

²²In session 12, senders failed to identify the fact that receivers were choosing b much more often than a following \emptyset . In sessions 20 and 21, elicited beliefs are rather inaccurate on all PD messages, but these messages were produced only 80 times over 360 observations for these two sessions.




Figure 11: *Best-response rates for envious senders. BR-b and BR-e measures the average rate at which envious senders best-responded to the elicited beliefs and to the empirical frequencies. BR-b and not BR-e measures the average rate at which envious senders best-responded to the elicited beliefs but not to the empirical frequencies. Not BR-b measures the rate of failure of best-response to the elicited beliefs.*

the corresponding indicator, BR-e, for best responses to the true empirical frequencies of the session. These two indicators allow us to decompose behavior into observations for which the sender best responded to both her reported beliefs and the true frequencies (BR-b AND BR-e), situations in which she failed to best respond to her own beliefs (Not BR-b), and situations in which she did best respond to but was lead astray by her beliefs (BR-b AND Not BR-e). The corresponding rates for envious senders²³ are given for all sessions in Figure 11. As for receivers, the rate at which envious senders failed to best respond to their belief is roughly between 20% and 30% for most sessions, except sessions 20 and 21 for which it is lower. By contrast, the rate at which senders best responded to inaccurate beliefs varies more widely across sessions corresponding to different games. It is higher in sessions with acyclic games.²⁴

²³Satisfied senders best responded to beliefs and true frequencies at a rate close to 100% in all sessions.

²⁴Session 12 is an exception, but this can be explained by the unbalance in the true frequencies of a and b actions following silence which can be noted in Appendix D, Table 9. Senders failed to pick up on this unbalance which they could have benefitted from, and instead best-responded to their reported beliefs.




Figure 12: *Game Variables: Sessions with acyclic games are on the top line. The boxes show for which sessions the corresponding indicator takes value 1.*

7 Explaining Performance

The former sections suggest that the existence of masquerading cycles affects the performance and behavior of senders and receivers. In this section, we try to get a more precise understanding of this effect through regressions that control for the effects of other game characteristics. Some of these characteristics, such as the number of envious types, or the number of masquerading incentives (the number of arcs in the graph), can be understood as measures of game complexity related to the incentive graph. The existence of cycles itself can be interpreted as a measure of complexity, since games with no cycles lend themselves to sequential reasoning, while games with cycles do not, as we pointed out in the introduction. Our regressions can therefore be interpreted as an attempt to compare the effects of different notions of complexity on the performances of the players. The expected direction of these effects is not immediately obvious. While some notions of complexity, such as the existence of cycles, are expected to make the task of the sender easier and that of the receiver harder, other types of complexity may increase both the difficulty for the sender of obfuscating smartly, and the difficulty for the receiver of making the right inferences.

Theoretical and empirical payoffs. Before showing the regressions, it is useful to start with a game-by-game overview of how players perform. In [Figure 13](#) and [Figure 14](#) of [Appendix C](#), we plotted, for each session, the range of theoretical equilibrium payoffs, and the empirical payoffs, for each sender type and the receiver. In each case, we computed the average payoff, and confidence intervals for the whole session, its first half and its second half. The first remark

is that no session seems to strongly rule out PBE as a possible explanation of the data on the payoff dimension,²⁵ except perhaps for sessions 7, 17 and 18. In sessions 7 and 18, satisfied types get an average payoff lower than 3. In these two sessions, receivers always interpreted full disclosure correctly, so the discrepancy must come from a failure of some satisfied senders to fully disclose. The result can indeed be explained by the behavior of a few bad players who use vague messages when satisfied. In session 17, the unique equilibrium is such that all types pool and the receiver’s expected outcome is 1. In fact, senders use silence less than predicted by the theory, explaining the better performance of receivers.

Envious Senders. Having confirmed that satisfied senders perform optimally nearly all the time, we naturally focus on envious senders. To measure the performance of envious senders, we construct a binary success variable that takes value 1 if the receiver takes an action that the sender prefers to the one that the receiver would take knowing her true type, and 0 otherwise. We then estimate a linear probability model to evaluate the determinants of performance. All our variables are binary indicators. The results of the regression are given in [Table 5](#).

The first five variables in [Table 5](#) are indicators for game characteristics. The first two variables indicate whether the game is cyclic or acyclic, and whether it has a FRE. This allows us to separate the effect of cycles from the existence of a FRE. However, these two indicators only differ for sessions 18 to 21. The next variable, *Several Envious*, indicates whether there is more than one envious type in the masquerading graph. The next two variables are indicators of the number of masquerading incentives. The first variable, *Many Arcs Cyc*, takes value 1 if the corresponding graph has three or more arcs, and if the game is cyclic, while *Many Arcs Acyc* takes value 1 if the corresponding graph has three or more arcs, and if the game is acyclic.²⁶

[Figure 12](#) shows how our game variables partition the set of sessions.

The next two variables are indicators of type characteristics, sometimes coupled with char-

²⁵The test is obviously far from perfect. For example, the vector of average payoffs may lie out of the convex hull of equilibrium payoffs, even though each average payoff lies in the range of possible equilibrium payoffs for this dimension.

²⁶Note that the distinction between acyclic graphs with one or with two arcs is redundant as it is already done by crossing *Several Envious* and *Many Arcs Acyc*. There are no cyclic graph with a single arc. Note also that a variable encoding whether all types are envious in cyclic graphs is redundant: our cyclic graphs with three or four arcs always exhibit three envious types.

Table 5: *Envious Sender Performance*

Estimator Dep. var.	(1)	(2)	(3)
	Linear Probability Model Sender Success		
Cyclic	0.229*** (0.029)		0.113*** (0.041)
No FRE		0.261*** (0.026)	0.201*** (0.039)
Several Envious	0.140*** (0.028)	0.203*** (0.028)	0.147*** (0.027)
Many Arcs Cyc.	-0.014 (0.035)	-0.108 (0.032)	-0.106*** (0.032)
Many Arcs Acy.	-0.059 (0.024)	-0.112 (0.026)	-0.060** (0.024)
Masq. as Two	0.113*** (0.027)	0.108*** (0.027)	0.110*** (0.027)
Type YB	-0.079*** (0.022)	-0.141*** (0.020)	-0.118*** (0.021)
Const.	0.035 (0.023)	0.056*** (0.021)	0.048** (0.022)
R^2	0.104	0.116	0.121

2411 observations. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Standard errors clustered at the individual level appear in parenthesis.

acteristics of the masquerading graph. *Masq. as Two* indicates whether the type of the sender is such that she wants to masquerade as any of the other two types rather than only one. The variable *Type YB* indicates whether the type of the sender is *YB* or one of the two other types. We include this variable because of the evidence structure which implies that type *YB* has more choices and therefore faces a more complex choice situation, potentially affecting performance.

As mentioned above, most of these variables can be interpreted as indicators of increased complexity along different dimensions: cycles make sequential logic ineffective, FRE is a simple equilibrium to coordinate on, with multiple envious senders there are multiple types whose strategic behavior needs to be understood, more arcs in the graph mean that incentives are

harder to track, and, finally, being of type YB means more messages to choose from for the sender.

In the last specification of [Table 5](#), all effects are significant. First we see that, controlling for other game characteristics, envious senders benefit from cycles. The positive effect of cycles is strongly enhanced by the absence of FRE as shown by the significant effect of *No FRE*. This is confirmed by comparing all sessions in which a 2-cycle is played. [Figure 14](#) shows that in 2-cycle games with no FRE (sessions 11 to 14), envious senders get a higher average payoff than in 2-cycle games with a FRE (sessions 18 to 21). The multiplicity of envious senders also has a positive effect on their performance. This indicator singles out sessions 1 to 3, in which there is a single envious sender, as particularly simple games that make the task of receivers very easy and the task of envious senders very difficult, as the average payoffs in [Figure 13](#) further illustrate.

Interestingly, the effect of the number of masquerade relations goes in the opposite direction. As the graph becomes more complex in the sense that there are more masquerading incentives to contemplate, envious senders perform less well. This is true both in cyclic and acyclic games. For acyclic games, the variable singles out monotonic games (sessions 7 to 10) as more complex ones. For cyclic games, it singles out the 3-cycle and double cycle games (sessions 15 to 17) as more complex. As we will see with the next regression, the comparative loss of performance of envious senders in these games is not to the benefit of the receivers. Thus these variables seem to capture a notion of complexity that hinders the performance of both players. Finally, the significant variable *Masq. as Two* shows that envious senders benefit from more masquerading options. The significant negative effect of *Type YB* indicates that envious senders perform less well when they have more disclosure options to choose from.

Receivers. For receivers, performance is measured by a binary indicator that takes value 1 if the receiver chooses her optimal action given the true type of the sender. As for senders, we estimate a linear probability model. Since the evidence received limits possible choices if accounted for, we introduce two additional variables to explain the performance of the receiver:

Silence takes value 1 if the sender shows no card and 0 otherwise, and *FD* takes value 1 if the sender shows all cards and 0 otherwise.²⁷

Table 6: *Receiver Performance.*

Estimator Dep. var.	(1)	(2)	(3)
	Linear Probability Model		
	Receiver Accuracy		
Cyclic	-0.067*** (0.023)		-0.038* (0.029)
No FRE		-0.076*** (0.024)	-0.057** (0.029)
Several Envious	-0.096*** (0.023)	-0.116*** (0.021)	-0.096*** (0.023)
Many Arcs Cyc.	-0.110*** (0.031)	-0.083** (0.032)	-0.083*** (0.032)
Many Arcs Acy.	-0.020 (0.026)	-0.001 (0.023)	-0.021 (0.026)
Silence	-0.122*** (0.020)	-0.135*** (0.020)	-0.132*** (0.020)
FD	0.308*** (0.020)	0.297*** (0.020)	0.299*** (0.020)
Const.	0.779*** (0.023)	0.789*** (0.023)	0.787*** (0.023)
R^2	0.252	0.253	0.253

3690 observations. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Standard errors clustered at the individual level appear in parenthesis.

The results are shown on [Table 6](#). All effects are significant, except for *Many Arcs Acyc*. The variables *Silence* and *FD* have strong and significant effects on the performance of the receiver. As expected, the effect of *FD* is positive,²⁸ and the effect of *Silence* is negative, confirming

²⁷Since the realization of these variables results from a strategic choice by the sender, this creates a potential endogeneity bias. Indeed, aspects of the game that are unaccounted for by our game variables, and thus appear in the error term of our regression, may also explain the choice of the sender. While our set of game variables seems sufficiently rich to limit the extent of the bias, we also ran regressions with session fixed effects and the *Silence* and *FD* variables as a robustness check. The estimated coefficients for both variables are very close: -0.140 (0.021) for *Silence*, and 0.296 (0.020) for *FD*.

²⁸Note that the effect is so strong as to predict probabilities greater than 1 when all other variables are set to

our results of [Section 5](#). The existence of cycles and the absence of FRE have negative effects on receiver performance. However, these effects are of small magnitude, especially the effect of cycles once we control for the absence of FRE. This is confirmed by [Figure 14](#), where we can see that receivers achieve similar payoffs in acyclic games, which all have a FRE, and in cyclic games with FRE (session 18 to 21). The multiplicity of envious senders has a strong and significant negative effect on receivers performance whereas it helps senders.

Everything else equal, the negative effect of *Many Arcs Cyc.* mostly captures the fact that the 3-cycle and double cycle games of sessions 15 to 17 are more difficult than 2-cycle games for receivers. This is compatible with equilibrium theory which gives an expected payoff of $1/3$ for the receivers in these games, and of $2/3$ in 2-cycle games. [Figure 14](#) actually shows that receivers' empirical payoffs are close to the theoretical predictions for all cyclic games. As already noted, the effect of *Many Arcs Cyc* is negative for both senders and receivers. *Many Arcs Acyc.*, on the other hand, is not significant for the receiver.

8 Conclusion

We propose a novel and exploratory experiment to investigate hard information transmission in various settings. Our design involves a sender who can reveal any subset of cards she has seen to the receiver, allowing for partial disclosure in a simple and portable way. The experiment is not organized around one or two treatments but consists of a collection of sender-receiver games which span a rich set of payoffs. To analyze the data, we then rely on a common and reduced representation of the sender incentives by a graph on her type set, which was developed in [Hagenbach et al. \(2014\)](#). Our objective is to uncover some local and global properties of the incentive graph that affect players behavior and performance in all games.

We start with a basic observation that, in all games, receivers respect the evidence contained in the messages sent, and find that senders expect such consistent reactions. Next, we show that sender types can be split into two categories which strongly affect disclosure strategies: types

0. This is due to our linear probability model. However, it is possible to get rid of this problem while keeping a linear probability model by defining a FD indicator specific to each session. With this specification, all predicted probabilities are below 1, and other effects are unchanged.

with no outgoing masquerading arrow are said satisfied and fully disclose; types with at least one outgoing masquerading arrow are said envious and use vague messages. We also find that partial disclosure takes more time to provide by senders and more time to interpret by receivers than any other message. Finally, theoretical works have established that the existence of a fully revealing equilibrium (FRE) relies on skeptical reactions of the receiver to every (on-path and off-path) message. In games with an acyclic graph where a FRE always exists, we show that receivers indeed act skeptically most of the time, although these games have other equilibria in which skepticism plays no role. We also highlight that skepticism is particularly hard to exert when it consists in attributing a vague message to a satisfied type.

Appendix

A Theoretical Appendix

Proofs and Additional Results. We start by proving the main results for pure strategies, and then we consider the extension to mixed strategies. [Lemma 1](#) is proved in Hagenbach et al. (2014). We say that (μ, α, β) is a PBE with extremal beliefs if, in addition, it satisfies the property that for every off path message $m \notin \cup_{t \in T} \mu(t)$, β_m puts probability 1 on a single type. [Proposition 1](#) and [Corollary 1](#) are proved in Hagenbach et al. (2014) for equilibria with extremal beliefs. The following lemma shows that the restriction to extremal beliefs is without loss of generality for the class of games that satisfy [Assumption 1](#).

Lemma 2. *Under [Assumption 1](#), if (μ, α, β) is a PBE, then there exists a system of beliefs β' such that $\beta_m = \beta'_m$ for every on-path message $m \in \cup_t \mu(t)$, and such that (μ, α, β') is a PBE with extremal beliefs.*

Proof. Consider a PBE (μ, α, β) . Let m be an off path message. Note that, by assumption, any action $a \notin \cup_{t \in M^{-1}(m)} \{a^*(t)\}$ yields a null payoff, while for any belief concentrated on $M^{-1}(m)$, there exists an action in $\cup_{t \in M^{-1}(m)} \{a^*(t)\}$ that yields a strictly positive payoff. Therefore there exists $t \in M^{-1}(m)$ such that $\alpha(m) = a^*(t)$. Then we can define β'_m as the belief that puts mass

1 on t . Let β' be the belief system obtained by doing this operation for every off path message m , and by letting $\beta'_m = \beta_m$ for the remaining messages. It is easy to see that (μ, α, β') is a PBE with extremal beliefs. \square

To extend our results to mixed strategies, we introduce an additional assumption that is satisfied by all the games in the experiment. This assumption says that, when there exists a certifiable subset S with no worst-case type, no type in S can be strictly worse off by masquerading as some other type in S .

Assumption 2. *For every certifiable subset S with no worst-case type, we have that, for every $t, t' \in S$, $u_S(a^*(t'), t) \geq u_S(a^*(t), t)$.*

Note that the only loss of generality is for cycles of length longer than 2. Under this assumption we have the following result.

Proposition 2. *Under Assumption 1 and Assumption 2, there exists a mixed strategy FRE of the sender-receiver game if and only if there exists a pure strategy FRE of the sender-receiver game.*

Proof. Only one of these implications is non obvious. Assume that there exists a mixed strategy FRE. In such an equilibrium, each type mixes over a set of messages. These sets do not overlap, since the equilibrium is separating. And following any of these messages, which is on-path for type t , the receiver takes action $a^*(t)$. We can pick one message m^t in the set of messages sent by type t , and modify the strategy of the sender to sending m^t with probability 1 when she is of type t . The other messages become off path, but we assume that the receiver reacts to this messages with the belief that they come from t with probability 1, and thus picks action $a^*(t)$. These beliefs are consistent, and because we started from an equilibrium, they deter other types from using the same message.

Consider the messages that are off path in the initial equilibrium. Upon seeing such a message m in the initial equilibrium, the receiver forms a belief β_m with support in $M^{-1}(m)$. Then, given Assumption 1, her mixed response can only be to mix in some way between actions $a^*(s)$ corresponding to types $s \in M^{-1}(m)$ such that $\beta^m(s) > 0$.

Suppose first that $M^{-1}(m)$ admits a worst-case type. Then we can change the belief β^m to a belief that puts probability 1 on a worst-case type, and that would deter any type of the sender able to send m from deviating to m .

Suppose next that $M^{-1}(m)$ has no worst-case type. Then any consistent belief following m leads the receiver to play some action $a^*(s)$ with $s \in M^{-1}(m)$ with strictly positive probability. Since $M^{-1}(m)$ has no worst-case type, there exists a type $t \neq s$ in $M^{-1}(m)$, that strictly prefers $a^*(s)$ to $a^*(t)$, and by [Assumption 2](#), this type t must also prefer to $a^*(t)$ any mixture over consistent actions that puts strictly positive probability on $a^*(s)$. But this in turn implies that β^m could not have supported a FRE in mixed strategies. \square

This result allows us to extend the conclusions of [Proposition 1](#) and [Corollary 1](#) to mixed strategy equilibria.

B Experimental Sessions

We ran a total of 25 experimental sessions in three waves taking place between April 2014 and May 2016. We used the 21 sessions presented in [Table 3](#) in the main analysis (354 subjects). 4 additional sessions (68 subjects) are presented in [Table B](#) and include ones in the sender’s payoff matrices. These sessions have been excluded from the main data to keep the value of masquerading constant for envious types but including them does not affect our results. Empirical payoffs of all 25 sessions are reported in [Appendix C](#). Beliefs were elicited only in the third wave of experiments. Subjects played 30 rounds of play in the first wave and 20 in the next two waves to reduce the time spent by subjects in the lab.

A systematic exploration of all possible graphs on three types would have required many more sessions. In total, there are $2^6 = 64$ possible graphs with our type set, and only 36 if we take symmetries into account (that is make YY and BB equivalent, and keep YB distinct as required by the evidence structure). As illustrated on [Figure 3](#), our experimental sessions use 11 distinct masquerading graphs, 9 if we take symmetries into account. When selecting the graphs for our experiments, we aimed at having a similar number of acyclic and cyclic graphs.

We also wanted graphs which exhibited various number of arrows and not too many. Once we had picked the graphs, they were randomly assigned to the experimental sessions.





Session	Masquerade	Sender Payoff	Game Type	FRE	Rounds	Subjects	Exch. Rate	Beliefs
22*		$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 1 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	30	16	0.20	No
23*		$\begin{pmatrix} 1 & 3 & 0 \\ 0 & 1 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	20	18	0.25	No
24*		$\begin{pmatrix} 0 & 1 & 3 \\ 0 & 1 & 3 \\ 0 & 0 & 3 \end{pmatrix}$	Acyclic	Yes	30	14	0.20	No
25*		$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 1 & 3 \\ 0 & 3 & 1 \end{pmatrix}$	Cyclic	No	20	20	0.29	No


Table 7: *Additional sessions not included in the main data.*

C Theoretical and Empirical Payoffs


We have computed all equilibrium payoffs for each of our games. The range for these payoffs are indicated by the grey areas in [Figure 13](#) for sessions with acyclic games and in [Figure 14](#) for sessions with cyclic games. The multiplicity of equilibrium payoffs comes from receiver indifferences over actions due to the uniform probability distribution. In the game of session 1, for example, the envious type BB can generate a payoff of 3 in an equilibrium in which: YY fully discloses; YB and BB both show a blue card certifying the set $\{YB, BB\}$; the receiver interprets all off-path messages skeptically, plays a after seeing two yellow cards, and b after seeing one blue card. This is an equilibrium because the receiver is indifferent between b and c after seeing a blue card, as Bayes rationality prescribes a belief that puts probability $1/2$ on YB and $1/2$ on BB .

Session 5



	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	53.4 (5.6)	46.6 (5.6)	0 (.)	\emptyset	46.7	53.3
Y	70.5 (8.1)	29.5 (8.1)	0 (.)	Y	73.9	26.1
B	0 (.)	89.5 (9.9)	10.5 (9.9)	B	0	100

Session 6


	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	61 (6.4)	39 (6.4)	0 (.)	\emptyset	39.5	60.5
Y	82 (5.0)	18 (5.0)	0 (.)	Y	76.7	23.3
B	0 (.)	98.9 (1.1)	1.1 (1.1)	B	0	93.3

Session 8



	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	60 (5.6)	39.1 (5.5)	0.9 (0.8)	\emptyset	59.4	37.6
Y	57.5 (9.2)	42.5 (9.2)	0 (.)	Y	68	32
B	0 (.)	96.5 (1.5)	3.5 (1.5)	B	0	88.9

Session 9



	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	48.3 (3.8)	51.7 (3.8)	0 (.)	\emptyset	60.7	39.3
Y	64.2 (8.2)	35.8 (8.2)	0 (.)	Y	53.9	46.1
B	0 (.)	90.8 (4.5)	9.2 (4.5)	B	0	100

Session 12



	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	45.8 (4.0)	54 (4.1)	0.2 (0.2)	\emptyset	45.2	54.8
Y	46.2 (2.8)	53.7 (2.8)	0 (.)	Y	53.3	46.7
B	0 (.)	65.7 (17.8)	34.3 (17.8)	B	0	12.5

Session 14


	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	0 (.)	45.3 (3.8)	54.7 (3.8)	\emptyset	0	48.5
Y	28.6 (18.4)	71.4 (18.4)	0 (.)	Y	.	.
B	0 (.)	51.1 (1.7)	48.8 (1.7)	B	0	50.8


Session 20


	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	49.9 (0.1)	0.2 (0.2)	49.9 (0.1)	\emptyset	48.1	0
Y	60.9 (17.0)	39.1 (17.0)	0 (.)	Y	66.7	33.3
B	0 (.)	39.1 (17.0)	60.9 (17.0)	B	0	0


Session 21


	Beliefs			Frequencies		
	YY	YB	BB	YY	YB	BB
\emptyset	48.6 (0.9)	2.8 (1.8)	48.6 (0.9)	\emptyset	52.6	0
Y	93.6 (2.8)	6.4 (2.8)	0 (.)	Y	100	0
B	0 (.)	6.4 (2.8)	93.6 (2.8)	B	0	0


Table 8: Receivers average beliefs and empirical frequencies.

Session 5



	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	52.7 <i>(4.1)</i>	45 <i>(3.8)</i>	2.3 <i>(2)</i>	\emptyset	56.7	43.3	0
Y	57.5 <i>(10.6)</i>	41.5 <i>(10.1)</i>	1 <i>(1)</i>	Y	95.7	4.3	0
B	1.1 <i>(1.1)</i>	94.1 <i>(2.4)</i>	4.8 <i>(1.7)</i>	B	0	94.1	5.9

Session 6


	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	57.8 <i>(6.6)</i>	42.2 <i>(6.6)</i>	0 <i>(0)</i>	\emptyset	65.8	32.9	1.3
Y	55 <i>(9.5)</i>	45 <i>(9.5)</i>	0 <i>(0)</i>	Y	80	20	0
B	0.2 <i>(0.2)</i>	97.8 <i>(1.3)</i>	2 <i>(1.3)</i>	B	0	93.3	6.7

Session 8


	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	67.9 <i>(8)</i>	32 <i>(8)</i>	0.1 <i>(0.1)</i>	\emptyset	70.3	28.7	1
Y	69.3 <i>(9.8)</i>	30.7 <i>(9.8)</i>	0 <i>(0)</i>	Y	88	12	0
B	0.5 <i>(0.5)</i>	96.8 <i>(1.3)</i>	2.7 <i>(1.3)</i>	B	0	83.3	16.7

Session 9



	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	52.4 <i>(8.3)</i>	47.6 <i>(8.3)</i>	0 <i>(0)</i>	\emptyset	52.5	47.5	0
Y	37.4 <i>(12.5)</i>	62.6 <i>(12.5)</i>	0 <i>(0)</i>	Y	69.2	30.8	0
B	1.7 <i>(1.7)</i>	98.3 <i>(1.7)</i>	0 <i>(0)</i>	B	0	91.7	8.3

Session 12



	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	49 <i>(4.8)</i>	47.3 <i>(4.7)</i>	3.7 <i>(2.6)</i>	\emptyset	26.2	73.8	0
Y	45 <i>(3.1)</i>	55 <i>(3.1)</i>	0 <i>(0)</i>	Y	44.2	55.8	0
B	0 <i>(0)</i>	88 <i>(9.1)</i>	12 <i>(9.1)</i>	B	0	75	25

Session 14


	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	0 <i>(0)</i>	54.2 <i>(4.2)</i>	45.8 <i>(4.2)</i>	\emptyset	3	54.5	42.5
Y	42.8 <i>(17)</i>	57.2 <i>(17)</i>	0 <i>(0)</i>	Y	.	.	.
B	0 <i>(0)</i>	53.6 <i>(3.6)</i>	46.4 <i>(3.6)</i>	B	0	46.3	53.7

Session 20


	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	51.9 <i>(3.5)</i>	5 <i>(3.2)</i>	43.1 <i>(5.6)</i>	\emptyset	39.2	3.8	57.0
Y	60.6 <i>(10.4)</i>	35 <i>(9.6)</i>	4.4 <i>(3.2)</i>	Y	33.3	66.7	0
B	11.9 <i>(8.9)</i>	36.9 <i>(9.1)</i>	51.2 <i>(11.9)</i>	B	0	25	75

Session 21


	Beliefs			Frequencies			
	a	b	c	a	b	c	
\emptyset	49.2 <i>(0.7)</i>	1.6 <i>(1.4)</i>	49.2 <i>(0.7)</i>	\emptyset	47.4	9.0	43.6
Y	71 <i>(13.6)</i>	28.7 <i>(13.3)</i>	0.3 <i>(0.3)</i>	Y	66.7	33.3	0
B	0.3 <i>(0.3)</i>	28.7 <i>(13.3)</i>	71 <i>(13.6)</i>	B	40	20	40

Table 9: Senders average beliefs and empirical frequencies.

D Beliefs

E Screenshots

References

- BATTAGLINI, M. AND U. MAKAROV (2014): “Cheap Talk with Multiple Audiences: an Experimental Analysis,” *Games and Economic Behavior*, 83, 147–164.
- BENNDORF, V., D. KÜBLER, AND H.-T. NORMANN (2015): “Privacy Concerns, Voluntary Disclosure of Information, and Unraveling: an Experiment,” *European Economic Review*, 75, 43–59.
- BLUME, A., D. V. DEJONG, G. R. NEUMANN, AND N. E. SAVIN (2002): “Learning and Communication in Sender-Receiver Games: An Econometric Investigation,” *Journal of Applied Econometrics*, 17, 225–247.
- BLUME, A., Y.-G. KIM, AND G. B. SPRINKLE (1998): “Experimental Evidence on the Evolution of Meaning of Messages in Sender-Receiver Games,” *American Economic Review*, 88, 1323–40.
- BROWN, A. L., C. F. CAMERER, AND D. LOVALLO (2012): “To Review or not to Review? Strategic Thinking at the Box Office,” *American Economic Journal: Microeconomics*, 4, 1–26.
- CAI, H. AND J. T.-Y. WANG (2006): “Overcommunication in strategic information transmission games,” *Games and Economic Behavior*, 56, 7–36.
- CRAWFORD, V. AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1451.
- DICKHAUT, J. W., M. LEDYARD, A. MUKHERJI, AND H. SAPRA (2003): “Information Management and Valuation: An Experimental Investigation,” *Games and Economic Behavior*, 44, 25–53.
- DICKHAUT, J. W., K. A. MCCABE, AND A. MUKHERJI (1995): “An Experimental Study of Strategic Information Transmission,” *Economic Theory*, 6, 389–403.
- DRANOVE, D. AND G. Z. JIN (2010): “Quality Disclosure and Certification: Theory and Practice,” *Journal of Economic Literature*, 48, 935–963.
- DYE, R. A. (1985): “Disclosure of Nonproprietary Information,” *Journal of Accounting Research*, 23, 123–145.
- FORSYTHE, R., R. M. ISAAC, AND T. R. PALFREY (1989): “Theories and Tests of “Blind Bidding” in Sealed-Bid Auctions,” *Rand Journal of Economics*, 20, 214–238.

- GIOVANNONI, F. AND D. J. SEIDMANN (2007): “Secrecy, Two-Sided Bias and the Value of Evidence,” *Games and Economic Behavior*, 59, 296–315.
- GROSSMAN, S. J. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law and Economics*, 24, 461–483.
- HAGENBACH, J., F. KOESSLER, AND E. PEREZ-RICHET (2014): “Certifiable Pre-Play Communication: Full Disclosure,” *Econometrica*, 82, 1093–1131.
- JIN, G. Z. (2005): “Competition and Disclosure Incentives: An Empirical Study of HMOs,” *Rand Journal of Economics*, 26, 93–113.
- JIN, G. Z., M. LUCA, AND D. MARTIN (2015): “Is No News (Perceived as) Bad News? An Experimental Investigation of Information Disclosure,” Working paper.
- KING, R. R. AND D. E. WALLIN (1991a): “Market induced Information Disclosure: An Experimental Markets Investigation,” *Contemporary Accounting Research*, 8, 170–197.
- (1991b): “Voluntary Disclosure when Seller’s Level of Information is Unknown,” *Journal of Accounting Research*, 29, 96–108.
- KOESSLER, F. AND E. PEREZ-RICHET (2015): “Evidence Based Mechanisms,” Working paper.
- LAI, E. K., W. LIM, AND J. T.-Y. WANG (2015): “An Experimental Analysis of Multidimensional Cheap Talk,” *Games and Economic Behavior*, 91, 114–144.
- MATHIOS, A. D. (2000): “The Impact of Mandatory Disclosure on Product Choices: An Analysis of the Salad Dressing Market,” *Journal of Law and Economics*, 43, 651–77.
- MATHIS, J. (2008): “Full Revelation of Information in Sender-Receiver Games of Persuasion,” *Journal of Economic Theory*, 143, 571–584.
- MILGROM, P. (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12, 380–391.
- OKUNO-FUJIWARA, A., M. POSTLEWAITE, AND K. SUZUMURA (1990): “Strategic Information Revelation,” *Review of Economic Studies*, 57, 25–47.
- PEREZ-RICHET, E. AND O. TERCIEUX (2015): “Evidence Based Matching,” Working paper.
- SEIDMANN, D. J. AND E. WINTER (1997): “Strategic Information Transmission with Verifiable Messages,” *Econometrica*, 65, 163–169.
- VAN ZANDT, T. AND X. VIVES (2007): “Monotone Equilibria in Bayesian Games of Strategic Complementarities,” *Journal of Economic Theory*, 134, 339–360.
- VESPA, E. AND A. J. WILSON (2015): “Communication with Multiple Senders: An Experiment,” Forthcoming in *Quantitative Economics*.
- WANG, J. T.-Y., M. SPEZIO, AND C. F. CAMERER (2010): “Pinocchio’s Pupil: Using Eye-tracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games,” *American Economic Review*, 100, 984–1007.




Figure 13: *Equilibrium predictions and actual payoffs for sessions with an acyclic game. The grey zones represent the range of equilibrium payoffs, and the data points represent average payoffs and 95% confidence intervals for the 1st and 2nd half of the session.*




Figure 14: *Equilibrium predictions and actual payoffs for sessions with a cyclic game. The grey zones represent the range of equilibrium payoffs, and the data points represent average payoffs and 95% confidence intervals for the 1st and 2nd half of the session.*




Figure 15: Senders' screen.



Figure 16: Table for beliefs elicitation of senders.




Figure 17: Receivers' screen.




Figure 18: Table for beliefs elicitation of receivers.