Disclosing Information to Interacting Agents*

Kfir Eliaz† and Françoise Forges‡

December 2012

Abstract

This paper analyzes the problem of an informed principal who decides at each state of nature what verifiable information to disclose (privately) to two uninformed agents who interact in a strategic form game. We focus on a class of games in which payoffs are quadratic and the agents' actions are substitutes. It is shown that when the principal can commit to a disclosure policy, it is optimal to disclose all information to one agent and disclose no information to the other agent. We then identify sufficient conditions under which this optimal policy can be implemented in the absence of commitment.

JEL Classification: C72, D82

Keywords: Disclosure, Verifiable information, Sender-receiver games

1 Introduction

The outcome of strategic interaction often depends on a state of nature, which may not be observed by all involved parties. For example, economic fundamentals are often not fully known to firms or investors at the time they choose their strategies. Similarly, managers within a firm may be uncertain about the state of the firm when planning their actions. Many studies in the literature assume an *exogenously given* information structure and examine the effect of the amount or precision of information available to agents on the strategic interaction between them (examples of such works will be discussed below). This paper is concerned with the *endogenous* determination of information structures. The question we pose is, what information structures will

*We thank Frédéric Koessler for helpful comments and discussions.

†Brown University, Providence RI, kfir_eliaz@brown.edu.

‡Member of Institut Universitaire de France, CEREMADE and LEDa, Université Paris-Dauphine, francoise.forges@gmail.com.

players in a Bayesian game have, if these structures are designed optimally by an informed planner?

We consider situations in which a planner observes the realization of a random variable (the state of nature) that affects the payoffs of a game played by two players who do *not* observe the state. The planner, whose payoffs depends on the outcome of the game, then needs to decide what information, if any, to pass on to the players, taking into account how the players would play in the induced Bayesian game. In other words, at each state, the planner optimally decides on the Bayesian game that the players will face, taking as given the state-dependent payoffs from each profile of actions. We study the weak perfect Bayesian Nash equilibrium of the extensive-form game in which the planner first sends private messages to the two agents who then simultaneously choose actions.

One example of situations we have in mind is a government agency or official who has inside information on some important economic fundamentals (e.g., the interest rate to be set by the central bank or information about fiscal or monetary policy). This information affects the profitability of firms in the economy or of large institutional investors. Since the government entity cares about the decisions of the firms or investors, it may want to affect their decisions by passing some of its private information to each of the parties.

A similar situation also occurs within organizations. A senior manager may be aware of some of important information regarding the firm's profitability (e.g., upcoming merges or aquisitions, entry to new markets, layoffs or hiring decisions), which are unknown to managers working under him. These managers may have some conflicting interests, which involve their own promotion or the profitability of their departments/divisions within the firm. As a result, the senior manager may be concerned about how his subordinates may strategically use information that is revealed to them. The senior manager may therefore want to affect the strategic interaction between his subordinates by carefully choosing what information to reveal to each one.

This paper focuses on environments that have the following features. First, information is transmitted via hard, verifiable evidence: In each state the planner can report the set of possible states, which must include the true state. Thus, we consider situations in which the state of nature will eventually be revealed to the players, and it is prohibitely costly for the planner to be caught telling a lie (i.e., declaring a state to be impossible when in fact that state was realized). The planner, however, can manipulate the accuracy of the information he transmits by increasing the number of states that he declares possible. Second, we follow the literature on the social value of information (e.g., Morris and Shin (2002) and Angeletos and Pavan (2007)) and assume that payoffs are quadratic, which ensures the existence of pure strategy equilibria. We also assume that the actions of the players are substitutes. As illustrated in Angeletos and Pavan (2007), this class of payoffs accommodates a number of different applications that have been discussed in the literature.

Our main results are as follows. First, we show that when the planner can commit in advance to what information he will pass in each state, then his ex-ante expected payoff is maximized when one agent is fully informed and the other agent is fully uninformed. That is, the first-best outcome calls for the planner to "collude" with one of the parties by sharing its private information, while keeping the other party in the dark. When this occurs, the informed agent has no interest to share information with the uninformed agent. Second, we give sufficient conditions under which this first-best outcome can be achieved in weak perfect Bayesian Nash equilibrium when the planner cannot commit to a strategy.

The problem studied in this paper may be viewed as a "reverse implementation" problem. First, with respect to implementation under complete information, the informational roles of planner and agents are reversed: the planner is informed but the agents are not. Second, while implementation theory is concerned with the design of effective institutions to achieve a given goal, our planner is stuck with the institution, and he can only resort to decide how much information he should leak to the agents. While implementation theory has achieved an impressive degree of generality, we have focused on a particular class of environments in the current paper.

The paper is organized as follows. Section 2 presents the model, while Section 3 presents the results. Related literature is discussed in Section 4 and concluding remarks are given in Section 5.

2 Model

A state θ is a discrete random variable distributed over $\Theta \subseteq \mathbb{R}_+$. Let μ be the mean of θ and denote $\theta_L = \min \Theta$ and $\theta_H = \max \Theta$. For any $S \subseteq \Theta$, let $\underline{\theta}(S)$ and $\bar{\theta}(S)$ denote the minimal and maximal elements in S.

There are two agents and a principal. Each agent can either exit or choose an action in $\mathbb{R}_+$. We let (x, y) denote the actions of players 1 and 2. If both agents opt in, then agent i's payoff has the property that for every profile of actions for the other agent, i has a pure best response, which is essentially affine in both θ and the action y of the other agent. Specifically, the best response increases in θ and decreases in y.

Hence, by denoting as x agent i's own action, agent i's payoff from (x, y) in state θ is the quadratic function:

$$u_i((x, y), \theta) = -\alpha_1 x^2 + (\alpha_2 \theta - \alpha_3 y)x$$

where $(\alpha_k)_{k=1,2,3}$ are positive parameters. If at least one agent opts out, then each receives a payoff of $\varepsilon > 0$, where ε is arbitrarily close to zero.

If both agents opt in, then the principal's payoff from (x, y) in state θ is a quadratic function of the sum of actions:

$$v((x, y), \theta) = -\beta_1 (x + y)^2 + \beta_2 \theta (x + y) \tag{1}$$

where β_1 and β_2 are positive parameters. If at least one agent opts out, then the principal gets a payoff of zero. As recalled in the introduction, these payoff functions are well-behaved (they represent single-peaked preferences, lead to pure-strategy equilibrium in each state, etc.) and many applications have this structure.

To simplify the exposition, we focus on the case $\alpha_k = \beta_k = 1$ for all k, namely

$$\begin{aligned} u_i((x, y), \theta) &= x\,[\theta - (x + y)] \\ v((x, y), \theta) &= (x + y)\,[\theta - (x + y)] \end{aligned}$$

With the normalization, we thus face the same payoff functions as in a Cournot duopoly with uncertain (inverse) demand.[1] The results that follow do not crucially depend on the normalization.

The principal and the agents play the following three-stage game. In stage zero a state of nature θ is realized and is observed only by the principal. In stage 1, the principal sends each agent a private message $S \subseteq \Theta$ such that $\theta \in S$. This message is interpreted as the set of states that the principal views as possible. In stage 2, after agents observed their signals, they simultaneously choose their actions. We analyze the weak perfect Bayesian Nash equilibria (PBE) of this game.

As explained in the introduction, we interpret the principal's strategy as capturing situations in which the cost of lying is excessively high. The principal must present hard evidence that must contain a "grain of truth" in the sense that it must include

[1] As pointed out by Einy et al. (2010), as soon as there is uncertainty in a Cournot duopoly, possibly negative prices have an impact on existence of an equilibrium. We take the payoff functions seriously, namely, implicitly allow negative prices, but insist on positive quantities. Our interpretation is that, in case of overproduction (i.e., $x+y > \theta$), the agents face a loss, *as if* the price were negative. We thus see the payoff function as $u_i((x, y), \theta) = x \max\{\theta - (x + y), 0\} - x\kappa(x+y)$ with $\kappa(z) = \max\{z - \theta, 0\}$.

the true state. What we have in mind are situations where the state of nature will be revealed eventually, and the principal does not want to be caught saying in state θ that this state is impossible.

We impose the following restrictions on agents' out-of-equilibrium beliefs. Let S_i denote an out-of-equilibrium message received by agent i. Let $S_i^+ \subseteq S_i$ be the set of states in S_i to which i assigns positive probability. Let $T_j(S_i)$ be the set of states that agent i considers possible for agent j when agent i himself gets S_i. We require that $S_i^+ \subseteq T_j(S_i)$ to reflect that fact that even out-of-equilibrium agent i knows that the principal must be truthful.

Let $s(\theta) = (s_1(\theta), s_2(\theta))$ denote a strategy for the principal, where we denote by $s_i(\theta) = \Theta$ the action "no signal to agent i when the state is θ". To simplify notation we let $x(\theta)$ and $y(\theta)$ denote the best responses of players 1 and 2, respectively, to $s(\theta) = (s_1(\theta), s_2(\theta))$. Agent 1 chooses $x(\theta)$, $s_1(\theta)-$measurable, to maximize

$$E[u_1(x(\theta), y(\theta), \theta) \mid s_1(\theta)] = x(\theta) \left[E[\theta - y(\theta) \mid s_1(\theta)] - x(\theta)\right]$$

namely a parabola with roots 0 and $E[\theta - y(\theta) \mid s_1(\theta)]$. If $E[\theta - y(\theta) \mid s_1(\theta)] \leq 0$, the max over $\mathbb{R}_+$ is 0. We interpret the action zero as opting out.

We will typically assume that the distribution of θ cannot put much weight on relatively high values, in the sense that

$$\mu < 3\theta_L \tag{A1}$$

$$\text{and } \mu < \frac{3}{7}\theta_H \tag{A2}$$

(A1) is equivalent to $\theta - \frac{1}{3}\mu > 0$ for every θ, which means that an agent who knows the state θ and conjectures that the other agent will choose $\frac{1}{3}\mu$ always opts in. (A1) is always satisfied if $\theta_H < 3\theta_L$ and is assumed throughout this note. (A2) is equivalent to $\mu - \frac{1}{2}\left[\theta_H - \frac{1}{3}\mu\right] < 0$, which means that an agent who does not know anything about the state and makes the pessimistic conjecture that the other agent will choose $\frac{1}{2}\left[\theta_H - \frac{1}{3}\mu\right]$ opts out. (A2) is not assumed unless explicitly mentioned.

3 Analysis

We begin by considering the case in which the principal can commit to a strategy, i.e., he decides what pair of private messages to send at each state of nature *before* the state is realized. A commitment strategy is optimal if it maximizes the ex-ante expected payoff of the principal, taking into account the second-stage game between the agents.

Proposition 1 *Assume the principal commits to a strategy before observing θ. Then an optimal strategy is to reveal θ to one agent and reveal no information to the other one.*

Proof. The objective function of the principal is

$$E_\theta\{[x(s_1(\theta)) + y(s_2(\theta))][\theta - x(s_1(\theta)) + y(s_2(\theta))\}$$

where

$$x(s_1(\theta)) = \max\left\{0, \frac{1}{2}E[\theta - y(s_2(\theta)) \mid s_1(\theta)]\right\}$$

and similarly for agent 2. Define

$$\delta(\theta) \equiv x(s_1(\theta)) + y(s_2(\theta)) - \frac{\theta}{2}$$

Then the objective function may be rewritten as

$$E_\theta\{[\frac{\theta}{2} + \delta(\theta)][\frac{\theta}{2} - \delta(\theta)]\} = E_\theta[(\frac{\theta}{2})^2 - \delta^2(\theta)]$$

Step 1. We show that $E_\theta[\delta(\theta)] \geq \frac{\mu}{6}$. Given the above expression for $x(s_1(\theta))$ and the similar expression for $y(s_2(\theta))$, we have that $E(x) \geq \frac{1}{2}(E(\theta) - E(y))$ *and* $E(y) \geq \frac{1}{2}(E(\theta) - E(x))$. Hence $E(x+y) \geq \frac{2}{3}E(\theta) = \frac{2}{3}\mu$ and $E[\delta(\theta)] \geq \frac{\mu}{6}$.

Step 2. By convexity, $[E_\theta(\delta(\theta))]^2 \leq E_\theta\delta^2(\theta)$. Note that this implies that $E_\theta\delta^2(\theta) \geq (\frac{\mu}{6})^2$. Thus, for the objective of the principal, we obtain that

$$E_\theta[(\frac{\theta}{2})^2 - \delta^2(\theta)] \leq E((\frac{\theta}{2})^2) - (\frac{\mu}{6})^2 = \frac{1}{4}E(\theta^2) - \frac{1}{36}\mu^2$$

Step 3. We show that the lower bound on $E_\theta\delta^2(\theta)$ is achieved by having one agent fully informed and the other agent uninformed. Suppose one agent is fully informed while the other one is fully uninformed. The informed agent then chooses $\frac{1}{2}\theta - \frac{1}{6}E(\theta) \geq 0$ for every θ iff $\frac{1}{3}E(\theta) \leq \theta_L$ (namely, (A1)), while the uninformed one chooses $\frac{1}{3}E(\theta) \geq 0$. It follows that the principal's ex-ante expected payoff is:

$E\left[(x+y)(\theta - (x+y)\right] = E\left[(\frac{1}{2}\theta + \frac{1}{6}E(\theta))(\frac{1}{2}\theta - \frac{1}{6}E(\theta))\right] = E\left[\frac{1}{4}\theta^2 - \frac{1}{36}E^2(\theta)\right] = \frac{1}{4}E(\theta^2) - \frac{1}{36}E^2(\theta) = \frac{1}{4}E(\theta^2) - \frac{1}{36}\mu^2$. ■

Assume next that the principal cannot commit to a strategy prior to the realization of the state.

Proposition 2 *There exists a PBE with full disclosure.*

Proof. Consider the following profile of strategies. In each state θ, the principal sends each agent the singleton message $\{\theta\}$. When an agent receives a singleton message $\{\theta\}$, he chooses $\theta/3$. If an agent receives any other message S, he believes the state is $\bar{\theta}(S)$ and that the other agent receives the singleton message $\{\bar{\theta}(S)\}$. Therefore, he chooses $\bar{\theta}(S)/3$. If the principal always discloses the state, then the agents' actions constitute a Nash equilibrium. In the proposed equilibrium the sum of actions is $\frac{2}{3}\theta$ in each state. If in any state θ the principal would send agent i a non-singleton message S, then i would choose $\frac{1}{3}\bar{\theta}(S) \geq \frac{1}{3}\theta$. Since the principal's payoff decreases with the sum of actions, when this sum exceed $\frac{1}{2}\theta$, the principal has no profitable deviation.[2] ■

Proposition 3 *If $\theta_L \leq \frac{2}{3}\mu$, then there does not exist a PBE with no disclosure.*

Proof. Assume, by contradiction, that there exists a PBE in which in every state θ, the principal sends each agent the null message Θ and each agent then chooses $\frac{1}{3}\mu$. Suppose that in the lowest state θ_L, the principal deviates only to agent 1 by sending him the singleton message $\{\theta_L\}$. Agent 1 must believe the state is θ_L and would therefore choose $\max\{0, \frac{1}{2}\theta_L - k\}$, where k is agent 1's deterministic, degenerate belief about agent 2's action, conditional on receiving the message $\{\theta_L\}$. If $k \geq \frac{1}{2}\theta_L$, then the sum of actions would be $\frac{1}{3}\mu$. This is a profitable deviation for the principal if $\frac{1}{2}\theta_L - \frac{1}{3}\mu < \frac{2}{3}\mu - \frac{1}{2}\theta_L$. Since $\mu > \theta_L$ this inequality holds. If $k < \frac{1}{2}\theta_L$, then the sum of actions (as a result of the deviation) would be $\frac{1}{3}\mu + \frac{1}{2}\theta_L - k$. The deviation is profitable if $\frac{1}{2}\theta_L - k < \frac{1}{3}\mu$. Since $\theta_L \leq \frac{2}{3}\mu$, and since by definition, $k \geq 0$, the above inequality must hold. It follows that the principal's deviation is necessarily profitable. ■

Proposition 4 *Under (A2), there exists a PBE in which one agent is fully informed and the other agent is completely uninformed.*

Proof. Consider the following strategies. In every state θ, the principal sends the message $\{\theta\}$ to agent 1 and the null message Θ to agent 2. When agent 1 receives a message $S \subseteq \Theta$ he believes the true state is $\bar{\theta}(S)$ and that the principal did not deviate against agent 2. He, therefore chooses $x(\{\theta\}) = \frac{1}{2}\bar{\theta}(S) - \frac{1}{6}\mu$ (which is positive by (A1)). When agent 2 receives no message, he chooses $\mu/3$. If he receives $S \subset \Theta$

[2] Observe that the proof does not make use of (A1).

with $\underline{\theta}(S) \geq \mu$, then he believes the state is $\bar{\theta}(S)$ and that the principal did not deviate against agent 1. He therefore chooses

$$y(S) = \frac{1}{2}\bar{\theta}(S) - \frac{1}{2}[\frac{1}{2}\bar{\theta}(S) - \frac{1}{6}\mu] = \frac{1}{4}\bar{\theta}(S) + \frac{1}{12}\mu$$

If he receives $S \subseteq \Theta$ with $\underline{\theta}(S) \leq \mu$ and $\bar{\theta}(S) \geq \mu$, he believes that the true state is distributed according to a probability distribution $f(S)$ over S such that $E_{f(S)}(\theta \mid S) = \mu$. He also believes that the principal did not deviate against agent 1. He therefore chooses

$$y(S) = \frac{1}{2}\mu - \frac{1}{2}(\frac{1}{2}\mu - \frac{1}{6}\mu) = \frac{1}{3}\mu$$

If he receives $S \subset \Theta$ with $\bar{\theta}(S) \leq \mu$ he believes the state is $\underline{\theta}(S)$ and that the principal sent agent 1 the message $\{\underline{\theta}(S), \theta_H\}$. He then chooses the maximum between zero and $\frac{1}{2}\underline{\theta}(S) - \frac{1}{2}(\frac{1}{2}\theta_H - \frac{1}{6}\mu)$. Since $\theta_H/\mu > 7/3$ (by (A2)) this last expression is negative, and so agent 2 would choose zero.

It is easy to verify that when the principal carries out his equilibrium strategy, the agents' strategies constitute a BNE in the ensuing subgame. It remains to show that the principal has no incentive to deviate.

Suppose the state is $\theta > \mu$. An agent who receives an out-of-equilibrium message will never reduce his action and for some messages, he will strictly increase it. The distance between the equilibrium sum of actions and the ideal sum of action (which is $\frac{1}{2}\theta$) is $\frac{1}{6}\mu$. If the principal deviates this distance may either remain unchanged or even increase. Consequently, the principal cannot profit from such a deviation.

Suppose the state is $\theta \leq \mu$. If the principal deviated only against agent 1, then agent 1 would increase his action and this would lower the principal's payoff (because the distance between the sum of actions and the ideal sum would increase). Suppose the principal deviated against agent 2 (and possibly against agent 1). If the principal sends agent 2 a message containing a state $\theta' > \mu$, then 2 would not change his action (and agent 1 would only increase his). If the principal sends agent 2 a message S with $\bar{\theta}(S) \leq \mu$ then 2 would choose zero, lowering the principal's payoff. ■

The agents' ex-post participation constraint plays an important role in the proof of Proposition 4. Relaxing this constraint complicates the analysis. To illustrate this, suppose the state of nature can be either θ_H with probability p or θ_L with probability $1-p$. Assume that agents cannot opt out and that the principal's payoff is given by (1). We show that to apply the same reasoning as in Proposition 4, we need to make further assumptions on the distribution of θ. A possible formulation is proposed below.[3]

Observation 1. *There exists a PBE that attains the maximal ex-ante expected payoff to the principal if*

$$\max\{\frac{7-p}{3-p}, \frac{5-2p}{3-2p}\} < \frac{\theta_H}{\theta_L} < \frac{2+p}{p} (< 3)$$

Proof. Agent 1 uses the following strategy. If he gets a singleton message $\{\theta\}$ he chooses $\theta/2 - \mu/6$, which is non-negative under (A1), namely $\frac{\theta_H}{\theta_L} < \frac{2+p}{p}$. If he gets the uninformative message $\{\theta_L, \theta_H\}$ he believes the state is θ_H for sure and that the principal did not deviate to agent 2. Agent 1 then chooses the same action he chooses when he receives the message $\{\theta_H\}$.

Agent 2 uses the following strategy. If he receives the message $\{\theta_L, \theta_H\}$ he chooses $\mu/3$. If he gets the message $\{\theta_H\}$ he believes the state is θ_H and the principal did not deviate to agent 1. Consequently, he chooses

$$\frac{\theta_H}{2} - \frac{1}{2}[\frac{\theta_H}{2} - \frac{\mu}{6}] = \frac{1}{4}\theta_H + \frac{\mu}{12} \geq 0$$

If he gets the message $\{\theta_L\}$ he believes the state is θ_L and that the principal sent agent 1 the message $\{\theta_L, \theta_H\}$. Consequently, he chooses the action

$$\max\{0, \frac{\theta_L}{2} - \frac{1}{2}[\frac{\theta_H}{2} - \frac{\mu}{6}]\}$$

We first verify that the principal has no incentive to unilaterally deviate against each agent. Consider agent 1 first. In state θ_H the agent will choose the same action regardless of the message he receives. In state θ_L a deviation only to agent 1 will only raise the aggregate action, reducing the principal's payoff.

Consider agent 2 next. Suppose the state is θ_H. If the principal deviates only to agent 2 and sends him the message $\{\theta_H\}$ agent 2 will raise his action if

$$\frac{\theta_H}{2} - \frac{1}{2}[\frac{\theta_H}{2} - \frac{\mu}{6}] \geq \frac{\mu}{3}$$

This inequality holds iff

$$\frac{\theta_H}{4} \geq \frac{\mu}{3} - \frac{\mu}{12} = \frac{\mu}{4}$$

and this inequality is satisfied. Suppose the state is θ_L and the principal deviates only to agent 2. If agent 2 were to choose zero, the aggregate action to the principal would

[3] The condition $\frac{7-p}{3-p} < \frac{\theta_H}{\theta_L}$ is equivalent to $\theta_L < \frac{1}{2}(\theta_H - \frac{1}{3}\mu)$ and is thus weaker than (A2), while $\frac{\theta_H}{\theta_L} < \frac{2+p}{p}$ is just (A1).

be $\frac{1}{2}\theta_L - \frac{1}{6}\mu$ and hence the principal would be indifferent between deviating or not. For agent 2 to choose zero when receiving the message $\{\theta_L\}$ it must be that

$$\frac{\theta_L}{2} - \frac{1}{2}[\frac{\theta_H}{2} - \frac{\mu}{6}] < 0$$

Since $\mu = p\theta_H + (1-p)\theta_L$ this inequality reduces to $(7-p)/(3-p) < \theta_H/\theta_L$, which holds by assumption (note that $(7-p)/(3-p) < (2+p)/p$ if and only if $p < 1$).

Finally, it remains to check that the principal has no incentive to deviate against both agents. Since in state θ_H agent 1 chooses the same action for any message, the principal cannot gain from deviating in this state. If the principal deviates against both agents in state θ_L, agent 2 would choose zero and agent 1 would choose $\theta_H/2 - \mu/6$. If this action is at least as high as $\theta_L/2 + \mu/6$, then the deviation is not profitable for the principal. This inequality holds if $\theta_H/\theta_L \geq (5-2p)/(3-2p)$, which holds by assumption (note that $(5-2p)/(3-2p) < (2+p)/p$ if and only if $p < 1$). ■

Throughout this note, we have assumed that the principal has to provide hard evidence of the information that he transmits to the agents. We now briefly investigate to what extent the principal can implement his ex ante first best payoff (as described by proposition 1) by just using *cheap talk*. In this case, the principal has further possible deviations, as he can report a false state, but at the same time, weak perfect Bayesian Nash equilibria do not impose more restrictions than plain Nash equilibria. As above, we assume that the state of nature takes two values, θ_H with probability p and θ_L with probability $1-p$ and that agents cannot opt out. Cheap talk is precisely modeled by modifying stage 1 of the three-stage game as follows: the principal sends each agent a private message in $\{\Theta, \{\theta_H\}, \{\theta_L\}\}$.

Observation 2. *There exists a PBE of the cheap talk game that attains the maximal ex-ante expected payoff to the principal iff*

$$\frac{5-2p}{3-2p} \leq \frac{\theta_H}{\theta_L} \leq \frac{2+p}{p} (\leq 3)$$

Proof. Consider the following strategies: agent 1 chooses $\frac{\theta}{2} - \frac{\mu}{6}$ (resp., $\frac{\theta_H}{2} - \frac{\mu}{6}$) if he receives the message θ (resp., θ_H) from the principal; agent 2 chooses $\frac{\mu}{3}$ whatever the message of the principal; the principal tells the true state to agent 1 and sends the non-informative message Θ to agent 2.

These strategies achieve the principal's maximal ex-ante expected payoff and unilateral deviations are not profitable to the agents. The only unilateral deviation that can be profitable to the principal is to send agent 1 the message θ_L when the true state is θ_H. If the principal does so, agent 1 chooses $\frac{\theta_L}{2} - \frac{\mu}{6}$, so that the aggregate action is $\frac{\theta_L}{2} + \frac{\mu}{6}$. If the principal does not deviate, the aggregate action is $\frac{\theta_H}{2} + \frac{\mu}{6}$. The deviation is not profitable to the principal iff

$$\frac{\theta_L}{2} + \frac{\mu}{6} \leq \frac{\theta_H}{2} - \frac{\mu}{6}$$

which is equivalent to

$$\theta_L \leq \theta_H - \frac{2}{3}\mu$$

or the first inequality in the statement of observation 2. The second inequality is just (A1) which guarantees that agent 1 always chooses positive actions. ■

The statement in observation 2 does not extend when there are $n > 2$ states. By proceeding as above, a necessary incentive compatibility condition for the principal is

$$\theta_L \leq \theta_H - \frac{2}{3}\mu \quad \text{for every } \theta,\, \theta' \text{: } \theta' < \theta$$

which implies that

$$\theta_L \leq \theta_H - \frac{2(n-1)}{3}\mu$$

This shows that incentive compatibility for the principal becomes harder (or even impossible) as n gets larger. For instance, if $n \geq 4$, it implies $\frac{\theta_H}{\theta_L} \geq 3$. But, as we pointed out above, a strong form of (A1) is the reverse inequality.

4 Related literature

This paper studies a game in which an informed sender sends private, verifiable messages to uninformed receiver who interact stragically. A related game was studied by Eliaz and Serrano (2012). These authors examine a situation in which the receivers play a generalized multi-action prisoner's dilemma. They show that under some conditions, the first-best outcome for the sender calls for no information transmission. However, the authors also present sufficient conditions under which the first-best involves asymmetric treatment of the two players, i.e., in each state the receiver get different information from the sender and hence, are asymmetrically informed (but in contrast to the current paper, both receivers have some information).

Our paper is closely related to the cheap-talk literature that studies games between an informed "sender" and uninformed "receivers".[4] The feature common to our model and to cheap-talk is the absence of commitment on the part of the sender. However, there are two salient distinctions. First, we make an important assumption: the planner cannot lie. That is, sooner or later he will be pressed to provide "hard facts" or evidence about the true state of nature, and punishments to lying would be prohibitive. Thus, his typical message will be of the form: "here's the set of possible states," and such a message will always contain the true state of nature as one of the possibilities. Second, we consider two *strategic* receivers of the message, while most of the cheap-talk literature is concerned with only one receiver. Two well-known exceptions are Farrell and Gibbons (1989) and Stein (1989). However, the first paper assumes that the payoff of each receiver is independent of the actions of the other receiver, while the second paper models the set of receivers as a single representative agent.[5] Hence, neither paper examines the effect of externalities across the receivers' actions on the sender's incentives. Recent works that study cheap talk in auctions include Jewitt and Li (2012), who consider public announcements, and Azacis and Vida (2012), who consider private messages.

Another related literature analyzes voluntary disclosure of verifiable information (see Milgrom (1981), Milgrom and Roberts (1986) and Okuno-Fujiwara, Postlewaite and Suzumura (1990) for early papers and Hagenbach et al. (2012) for a recent one). In these models there is a set of players who possess verifiable information on the state of nature and need to decide how much of this information to disclose (e.g., the parties may announce a set of states that include the true state). A large part of this literature characterizes conditions for "unraveling", whereby all private information is revealed in equilibrium.[6]

The current paper is concerned with the effect of varying forms of information on the strategic interactions of the recipients. As such, it is also related to a recent strand of the literature, which explores the social value of information (most notably, Morris and Shin (2002) and Angeletos and Pavan (2007)). The aim of these papers is to examine the equilibrium and welfare effects of changes in the precision and form of information (private versus public) on certain classes of economies or games. A central feature of these models is that the information structure is exogenously given, whereas our main

[4] For surveys on cheap-talk games see, e.g., Krishna and Morgan (2007) and Koessler and Forges (2008).

[5] Two recent investigations of sender-multiple-receivers games are Goltsman and Pavlov (2011) and Koessler (2008). Both papers look at the case in which the payoff of a receiver is independent of the actions of other receivers.

[6] Full revelation of information may not occur if some parties are not informed or cannot verify their private information (see Shavell (1989) and Farrell (1986)).

interest is in understanding what information structures would arise endogenously in equilibrium.

Our paper is also related to a recent literature on the design of procedures for information transmission. Like us, this literature is concerned with situations in which the principal cannot directly affect the players' payoff, but can do so only indirectly via information revelation. some examples of recent works include Rayo and Segal (2010), Gentzkow and Kamenica (2011), and Hörner and Skrzypacz (2011).

Finally, another way to view our study is in comparison to mechanism design or implementation theory. In that theory, the planner is uninformed about the state of nature and tries to elicit information about it from the agents by a clever design of the institution. Usually, in that framework, the planner simply sets up the mechanism and is not a player in it, although two notable exceptions are Baliga, Corchón and Sjöström (1997) and Baliga and Sjöström (1999). For a point of comparison, our planner is also a player in the information leakage game, but he is the informed party while the agents are not.

5 Concluding remarks

In this paper, we started to investigate the problem of strategic disclosure of verifiable information by an informed party to interactive uninformed agents. We focused on a specific class of games, which has several economic applications. In these games, we characterized the optimal disclosure policy under commitment ("the first best") and showed that under some conditions, this policy can be implemented in the absence of commitment. An interesting feature of the optimal policy is that it calls for *asymmetric* treatment of two *ex-ante symmetric* agents: one agent should be fully informed, while the other agent should be kept in the dark.

Much further work is needed to answer our basic question, namely what type of information structures in Bayesian games can be rationalized endogenously as the outcome of strategic information disclosure by an informed planner. A number of topics for future research readily emerge from our analysis. For instance, even within the class of games with quadratic payoff functions, one could consider the case where the actions of the players are complements rather than substitutes. Some examples suggest that under commitment, the planner might then fully inform both agents.

References

[1] Angeletos, G.-M. and A. Pavan (2007): "Efficient Use of Information and Social Value of Information," *Econometrica* **75**, 1103-1142.

[2] Azacis, H. and P. Vida (2012): "Collusive Communication Schemes in a First-Price Auction," Working Paper.

[3] Baliga, S. and T. Sjöström (1999): "Interactive Implementation," *Games and Economic Behavior* **27**, 38-63.

[4] Baliga, S., L. Corchón and T. Sjöström (1997): "The Theory of Implementation when the Planner Is a Player," *Journal of Economic Theory* **77**, 15-33.

[5] Einy, E., O. Haimanko, D. Moreno and B. Shitovitz (2010): "On the Existence of Bayesian Cournot Equilibrium," *Games and Economic Behavior* **68**, 77-94.

[6] Eliaz, K. and R. Serrano (2012): "Sending Information to Interactive Receivers Playing a Generalized Prisoners' Dilemma," Working Paper, Brown University.

[7] Farrell, J. (1986): "Voluntary Disclosure: Robustness of the Unraveling Result, and Comments on its Importance," Chapter 4 in *Antitrust and Regulation,* R. Grieson (Ed.), Lexington, Lexington (Mass.), 91-103.

[8] Farrell, J. and R. Gibbons (1989): "Cheap Talk with Two Audiences," *American Economic Review* **79**, 1214-1223.

[9] Goltsman, M. and G. Pavlov (2011): "How to Talk to Multiple Audiences," *Games and Economic Behavior* **72**, 100-122.

[10] Jewitt, I. and D. Li (2012): "Cheap-Talk Information Disclosure in Auctions," Working Paper.

[11] Hagenbach, J., F. Koessler and E. Perez-Richet (2012): "Certifiable Pre-Play Communication: Full Disclosure," Working Paper.

[12] Hörner, J. and A. Skrzypacz (2011): "Selling Information," Working Paper.

[13] Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review* **101**, 2590-2615.

[14] Koessler, F. (2008): "Lobbying with two Audiences: Public vs Private Certification," *Mathematical Social Sciences* **55**, 305–314.

[15] Koessler, F. and F. Forges (2008): “Multistage Communication with and without Verifiable Types,” *International Game Theory Review* **10**, 145-164.

[16] Krishna V. and J. Morgan (2007): “Cheap Talk,” in *The New Palgrave Dictionary of Economics*, 2nd Edition.

[17] Milgrom, P. (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics* **12**, 380-391.

[18] Milgrom, P. and J. Roberts (1986): “Relying on the Information of Interested Parties,” *Rand Journal of Economics* **17**, 18-32.

[19] Morris, S. and H.S. Shin (2002): “The Social Value of Information,” *American Economic Review* **92**, 1521-1534.

[20] Morris, S. and H.S. Shin (2003): “Global Games: Theory and Applications,” in *Advances in Economics and Econometrics* (Proceedings of the Eighth World Congress of the Econometric Society), edited by M. Dewatripont, L. Hansen and S. Turnovsky; Cambridge University Press.

[21] Okuno-Fujiwara, M., A. Postlewaite and K. Suzumura (1990): “Strategic Information Revelation,” *Review of Economic Studies* **57**, 25-47.

[22] Rayo, L. and I. Segal (2010): “Optimal Information Disclosure,” *Journal of Political Economy* **118**, 949-987.

[23] Shavell, S. (1989): “A Note on the Incentive to Reveal Information,” *The Geneva Papers on Risk and Insurance* **14**, 66-74.

[24] Stein, J. C. (1989): “Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements,” *American Economic Review* **79**, 32-42.