Int J Game Theory
DOI 10.1007/s00182-013-0374-x

Sending information to interactive receivers playing a generalized prisoners' dilemma

Kfir Eliaz · Roberto Serrano

Accepted: 3 April 2013

Abstract Consider the problem of information disclosure for a planner who faces two agents interacting in a state-dependent multi-action prisoners' dilemma. We find conditions under which the planner can make use of his superior information by disclosing some of it to the agents, and conditions under which such information leakage is not possible. Although the problem is entirely symmetric, the planner's only way to reveal part of the information is based on creating asymmetries between the two agents by giving them different pieces of information. We also find conditions under which such partially informative equilibria are the planner's best equilibria.

Keywords Information disclosure · Generalized prisoners' dilemma · Uninformative equilibria · Partially or fully informative equilibria

JEL Classification C72 · D82 · D83

1 Introduction

The problem studied in this paper considers an informed planner and two uninformed agents. Thus, with respect to implementation theory under complete information,the

We thank the Editor, the Associate Editor and the referees for their helpful comments, which helped improve the paper.

K. Eliaz (✉)
Eitan Berglas School of Economics, Tel-Aviv University, Tel Aviv, Israel
e-mail: kfire@post.tau.ac.il

R. Serrano
Economics Department, Brown University, Providence, RI, USA
e-mail: Roberto_Serrano@brown.edu

Published online: 08 May 2013

informational roles of planner and agents are reversed. Furthermore, 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. On the other hand, both approaches allow strategic interaction among agents. Finally, while implementation theory has achieved an impressive degree of generality, for the moment we are just considering a specific example in the current paper. We proceed to detail.

Two agents play a fixed normal form game where the set of actions available to each player is $A = \{1, 2, \ldots, K\}$ with $K > 2$. There is a set $\Omega = \{1, 2, \ldots, K\}$ of (equally likely) states of nature.[1] Actions are ordered in each state: specifically, each state ω defines a linear ordering $\rhd_\omega$ over the actions, such that

$$
\begin{aligned}
&1 \lhd_1 2 \lhd_1 \ldots \lhd_1 K-1 \lhd_1 K\\
&2 \lhd_2 3 \lhd_2 \ldots \lhd_2 K \lhd_2 1\\
&\vdots\\
&K \lhd_K 1 \lhd_K \ldots \lhd_K K-2 \lhd_K K-1
\end{aligned}
$$

If $a \lhd_\omega b$ for some actions $a, b \in A$, we say that action "b is higher than a at ω" . We interpret $\lhd_\omega$ as a state-dependent ranking of the "intensity" or "strength" of the actions. A "high" action in one state may be considered a "low" action in another state. Each player's payoff depends on the "intensity/strength" of the players' pair of actions, and its relationship with the state. For instance, the outcome of the action pair $(1, 2)$ in state 1 is equivalent to the outcome of the action pair $(2, 3)$ in state 2, and so on. Furthermore, actions are ex-ante symmetric, in the sense that for any given action there exists a state in which it is the highest, another state in which it is the lowest, and states in which it occupies any intermediate position. All of the above is common knowledge. However, the planner observes the state but the agents do not.

One may therefore think of the actions in A as if they are ordered clockwise on a circle from 1 to K. One of these actions, say a, is considered the "lowest" action and every action that follows it clockwise along the circle is considered to be a higher action (with the highest action being the action adjacent to a counter-clockwise). This linear ordering determines the payoffs to the agents and to the planner, from every pair of actions. That is, when the agents play (a, b) the payoffs depend on where a and b are ranked on the linear ordering. To restate the knowledge assumptions, there is common knowledge regarding the payoffs that will be obtained when one player chooses the i-th highest action and the other player chooses the j-th highest action for every $i, j = 1, \ldots, K$. However, there is asymmetric information regarding the identity of the lowest action: this is known to the planner but not to the agents.

To have a concrete running example in mind that fits our model, consider the dilemma faced by a State Governor in the face of a major weather disaster. Suppose the states of the world are the zip codes in the State of Rhode Island, ordered clockwise

[1] We shall discuss the case of $K = 2$ in our concluding section. The uniform distribution assumption is made for convenience; obviously, the statemens and conditions would need to be adjusted accordingly were one to use a different distribution.

on the map. In the event of a major disaster, an emergency plan is activated with the aim of allocating resources to various parts of the State. One of the aspects of the fixed institution facing the Governor is that, in order to avoid disruption in cases of such emergencies, there is a fixed protocol of zip codes, which is imposed for the display of the relevant resources (firemen, police, army, national guard, etc.). It seems reasonable to think that such a display should take place in an "ordered" manner, say, always "around the circle in a clockwise direction." Thus, the state k is interpreted as the zip code to which the major Federal help will be taken. The Governor knows what that k is, but bad weather conditions or other sources of uncertainty may give him enough leeway to make announcements that contain k but add noise to it. The lowest action for each player (the players in the game are the agencies mentioned above) would be simply to perform their preparatory duties to help with the effectiveness of the Federal emergency help only in state k, i.e., only in the zip code where it will be dropped. The highest action would be to perform such prep exercises in all zip codes from k to $k-1$, i.e., all Rhode Island. Each agency (think of police and firemen, to match our two players in the model) cares in part about the social goal of helping out with the emergency, but also cares about saving effort. Indeed, the agencies are facing a generalized prisoner's dilemma situation.

In these circumstances, it is not clear how much information the planner should give out to the agents. The paper offers a set of negative results (general conditions under which no information disclosure is possible). It also reports (somewhat more special) conditions compatible with such information revelation. In addition, when this happens, the only way to do it is through creating asymmetries in the information provided to agents, in spite of the complete symmetry of the problem. We continue now with the description of the abstract model.

2 Additional preliminaries

As already described, each state ω induces a linear ordering $\lhd_\omega$ over the actions with $a = \omega$ being the "lowest" action. (as part of a generalized prisoner's dilemma, we will assume that this is also a dominant action; details are found in the next section). It is convenient to normalize actions with respect to state 1. For any $a \in A$ and $\omega \in \Omega$, define $s(a, \omega)$ as the number of elements in A that are ranked below a according to $\lhd_\omega$. That is,

$$s(a, \omega) = \#\{b \in A : b \lhd_\omega a\}$$

For any $a \in A$ and $\omega \in \Omega$, we define $x(a, \omega)$ to be

$$x(a, \omega) \equiv \{a^* \in A : s(a, \omega) = s(a^*, 1)\}$$

In other words, $x(a, \omega)$ measures the rank of a according to $\lhd_\omega$. As we shall see in the next section, this tells us how "close" a is to the dominant action in state ω.

Player i's payoff in state ω from the action pair (a_1, a_2) is denoted $u_i(a_1, a_2 \mid \omega)$. Given our normalization to actions in state 1, denote $u_i(a_1, a_2 \mid 1)$ by $u_i(a_1, a_2)$, and we assume that for $i = 1, 2$ and for all a_1, a_2 and ω, $u_i(a_1, a_2 \mid \omega) = u_i[x(a_1, \omega), x(a_2, \omega)]$. We assume the game is symmetric: for any $x, y \in \{1, \dots, K\}$, $u_i(x, y) = u_j(y, x)$. Hence, to simplify the exposition we write $u(x, y)$ to denote the payoff to a player who chooses x, while the other player chooses y.

The planner, who knows that the state is ω, sends each agent a private message, which consists of any subset of Ω that contains the true ω. Such a subset is interpreted as the set of states that are possible. This captures situations where the true state will eventually be revealed, and it is prohibitively costly for the planner to lie. However, the planner can manipulate the precision of the evidence he discloses via the cardinality of the set he reports.

If in each state, the planner could send *either* private messages *or* a public message, then none of our results would change. The reason is, that the planner would not prefer to send a public message over private messages.[2] If, however, the planner could send *only* public messages, then all the results up to Lemma 2 (our impossibility results) would continue to hold.

The payoff to the planner is given by a function V. That is, the planner who knows that the true state is ω, when the agents choose actions a_1 and a_2 receives a payoff of $V(x(a_1, \omega), x(a_2, \omega))$. One leading specification of the function V, although not assumed in the formal results, is that it is the sum of players' payoffs, i.e.:

$$V(x(a_1, \omega), x(a_2, \omega)) = u[x(a_1, \omega), x(a_2, \omega)] + u[x(a_2, \omega), x(a_1, \omega)]$$

Given that the game is symmetric, we assume that V is symmetric so that $V(x, y) = V(y, x)$.

3 Assumptions on payoffs

The game played in each state is interpreted as a generalized multi-action Prisoners' Dilemma. To capture this interpretation we make a number of assumptions on payoffs. We group them in assumptions on the players' payoffs (denoted as a group by the letter A) and on the planner's payoffs (denoted by the letter B).

The first assumption on players' payoffs one could start with is that the lowest action in each state is strictly dominant: for any $x > 1$ and for any y,

$$u(1, y) > u(x, y)$$

However, because of the multi-action framework, this will not suffice for most of our results. Hence, we strengthen it as follows:

- (A1) Each player has a weak incentive to lower his action: for all $x > x' \geq 1$ and y,

  $$u(x', y) \geq u(x, y)$$

  with a strict inequality for $x' = 1$.

[2] The assumption is that a player can tell whether a message he receives is private or public (e.g., a public message is posted ona billboard while a private message is sent by mail).

Assumption (A1) captures the basic intuition in a generalized prisoners' dilemma logic, that each player has a preference for lowering his action/effort, ceteris paribus. In terms of the running example given above, each agency (police, firemen) has an incentive to "save effort" going around the circle. For some of our last results, we shall add another assumption in addition to (A1), which we introduce as follows.

First, consider the change in an agent's payoff when he lowers his action to the dominant one, in a state where his opponent chooses any action other than the dominant: For any $k > 1$,

$$\Delta_k(a) \equiv u(1, k) - u(a, k)$$

Consider next the change in the player's payoff when he lowers his action from the highest action in a state where his opponent chooses the dominant action:

$$\Delta_1(a) \equiv u(a, 1) - u(K, 1)$$

By (A1), while $\Delta_k(a)$ is increasing in a, $\Delta_1(a)$ is decreasing. Consider the following condition,

- (A2) For any $k > 1$, $\min_{a>1} \Delta_k(a) > \max_a \Delta_1(a)$ or given (A1): $\Delta_k(2) > \Delta_1(1)$.

To interpret condition (A2), one should think of the gains associated with lowering one's action. The assumption says that any such gain when the opponent is not playing the dominant action exceeds any such gain when he is. The individual gains from exerting less effort are very small if the other agent is already exerting the smallest amount of effort. In terms of our example, (A2) is a hint that the payoff function of firemen and police also contains social aspects: in particular, if the other agencies are "shirking" by only prepping in one zip code, the zip code where the help will be dropped, the payoff savings from (A1) are offset by the "bad feeling" of considering shirking given that the other agency is already shirking the most.

Next, we move to assumptions on the planner's payoffs. Two basic features, consistent with a prisoners' dilemma scenario, that we would like to capture are the following. On the one hand, coordination on the lowest action is the worst outcome for the planner: for any $x \geq 1$, $y \geq 1$ with at least one strict inequality,

$$V(1, 1) < V(x, y)$$

On the other hand, coordination on the highest action is the best outcome for the planner: for any $x \leq K$, $y \leq K$ with at least one strict inequality,

$$V(K, K) > V(x, y)$$

As before, we need to strengthen these basic features in order to obtain results in our multi-action model. That is, to capture a sense in which coordination on an intermediate action—i.e., coordination on some $a < K$—is also beneficial for the planner, we begin by considering the following:

- (B1) The planner prefers coordination on any action $1 < x \leq K$ to any pair of lower actions: for any $y \leq z \leq x$,

$$V(x, x) \geq V(y, z)$$

A stronger version of (B1), which we shall use in one of the results, offers a complete conflict of interest between planner and players. That is, while (A1) stipulates that players' payoffs are decreasing in the player's action, assumption (B2) poses that the planner's payoffs are increasing in actions:

- (B2) the planner's payoff increases with the players' actions: if $x > x'$, then for every y, $V(x, y) > V(x', y)$.

Assumptions (A1) and (B2) together may be interpreted as a public good problem in which the agents face uncertainty regarding the cost of investing in the public good. Assumption (A1) can be viewed as saying that the more an action contributes to the optimal public good, the more costly it is, so that agents have a strict incentive to free-ride. Assumption (B2) means that the planner is interested in the highest amount of contribution. In terms of our disaster relief example, (B1) and (B2) seem easy to justify, simply by covering "more of the circle" with the actions of police and firemen.

In contrast to Assumption (B2), Assumption (B2') below implies that, aside maximal cooperation, the planner's payoff increases when players' difference in actions rises. Namely, we assume that if for some reason the planner knows that the agents will not cooperate in the highest action, then he prefers that one agent undercuts the other with the highest margin (note that while this is a departure from (B2), it is still compatible with (B1)). As will be seen, whether one assumes (B2) or (B2'), the information disclosure happening in equilibrium will be significantly different. Assumption (B2') follows:

- (B2') For any $x < y < z$,

$$V(x, z) > \max\{V(x, y), V(y, z), V(y, y)\}. \tag{1}$$

This assumption may be viewed as capturing a trade-off that the planner has between the long-term (delayed) benefit of high effort and the short-term (immediate) benefit of low effort. It can also be viewed as capturing the idea that the planner cares about the sum of agents' payoffs in a situation where the payoff from one-sided free-riding is sufficiently large.

To have a concrete interpretation of this assumption, consider our disaster example once again. Short of providing full prep services to the entire State (the "top outcome") which would be saluted very favorably by the Federal Administration and recognized with strong praise by the national press and the voters, the extremely tight resource constraint facing the Governor makes him prefer that only one of the agencies (say, only the firemen, or only the police) go around the State doing the prep work, while the other stays put in the base zip code.

4 Analysis and results

The planner's information disclosure strategy can be described as follows: in each state ω he sends a pair of private messages (m_1, m_2), where $\omega \in m_i \subseteq \Omega$. That is, the planner announces a list of possible states, but one of the states he announces must be the true state.

We focus on pure-strategy PBNE using the following restrictions on out-of-equilibrium beliefs:

β_1 Thanks to the private communication assumption, one can think of the planner's move as two separate decision nodes, subject to independent trembles. Consistent with the idea of independent trembles, a player who finds himself off the equilibrium path believes that only he received an out-of-equilibrium message (i.e., he believes the other player received a message according to equilibrium).[3]

β_2 Consistent with the logic of some refinements (e.g., the intuitive criterion), a player assigns a positive probability to the event that a planner of type ω deviated from equilibrium only if this planner type has an incentive to deviate in this state.

We shall use these restrictions for the construction of the equilibria in Propositions 1, 4, and 5. The rest of results do not rely on these restrictions.

The following lemma is instrumental in our analysis.

Lemma 1 *If (A1) holds, then any PBNE has the following property. If an agent believes, both on or off the equilibrium path, that the set of all possible states is S, then he must choose an action, which is dominant in one of these states.*

Proof Consider a PBNE, a player, say 1, and a planner type ω. Let $S \subseteq \Omega$ be the message that the planner sends to player 1 in state ω. If $S = \Omega$ then the Lemma is trivially true. Suppose therefore that $S \subset \Omega$ and player 1 responds by choosing an action a, which is not dominant in any of the states in S. Then, there exists a state $\omega' \in S$ satisfying $x(a, \omega') < x(a, \omega'')$ for all $\omega'' \in S$. Since a is not dominant in ω', $x(a, \omega') > 1$. By our assumption on how the ranking of actions changes across states and by our choice of ω', any action b with $x(b, \omega') < x(a, \omega')$ satisfies $x(b, \omega'') < x(a, \omega'')$ for all $\omega'' \in S$. Let b^* be the dominant action in state ω', i.e., $x(b^*, \omega') = 1$. By (A1), if player 1 deviated to b^* his expected payoff would strictly increase as his payoff would increase in every state in S, a contradiction. □

To illustrate our model and some of our results, consider the following simple example. Suppose $A = \{1, 2, 3\}$ such that in state 1 the actions are ordered $1 - 2 - 3$, in state 2 they are ordered $2 - 3 - 1$ and in state 3 they are ordered $3 - 1 - 2$. This

[3] This restriction corresponds to what some of the literature has referred to as "passive beliefs" (see, e.g., McAfee and Schwartz (1994)). Our results would be robust to allowing some form of "active beliefs" , i.e., allowing a player who receives an out-of-equilibrium message to believe that the other player also received an out-of-equilibrium message. This is because in our proofs we assume that when a player gets a message out-of-equilibrium he puts probability one on one of the states in that message. So we can replace β_1 with the following assumption: when a player gets an out-of-equilibrium message, he believes that the planner revealed the true state—according to the player's belief—to the other player, i.e., he believes the planner sent a singleton message to the other player which contains the state that the player believes is true. Such beliefs may be viewed as being "paranoid"—a player thinks the planner favors the other side.

situation may be depicted graphically as a triangle in which the vertices are the actions 1, 2 and 3. One of the vertices corresponds to the lowest ranked action and the ranking of actions increases in a clockwise fashion:

$$\begin{tikzcd} & 1 \arrow[dr] & \\ 3 \arrow[ur] & & 2 \arrow[ll] \end{tikzcd}$$

Let $a_i(\omega)$ denote the i-th lowest action in state ω (e.g., $a_1(1) = 1$ and $a_1(3) = 3$). The left matrix below displays the mapping between the ranks of the chosen actions (i.e., $(a_i(\omega), a_j(\omega))$) and the payoffs to the agents. The right matrix displays the "average normal form game" (i.e., the normal form of the game when both players are uninformed).

$$\begin{array}{cccc} & a_1 & a_2 & a_3 \\ a_1 & 0,0 & 2,-1 & 5,-1 \\ a_2 & -1,2 & 1,1 & 2,-1 \\ a_3 & -1,5 & -1,2 & 3,3 \end{array} \qquad \begin{array}{cccc} & a_1 & a_2 & a_3 \\ a_1 & \frac{4}{3},\frac{4}{3} & 1,1 & 1,1 \\ a_2 & 1,1 & \frac{4}{3},\frac{4}{3} & 1,1 \\ a_3 & 1,1 & 1,1 & \frac{4}{3},\frac{4}{3} \end{array}$$

Thus, in state $\omega = 1, 2, 3$, action ω is dominant for each agent. However, the efficient outcome corresponds to the action profile $(\omega - 1, \omega - 1)[\text{mod}3]$ with payoffs $(3, 3)$. Suppose the planner's payoff is given by the sum of the agents' payoffs.

What information should the planner send to the agents? Put differently, what messages does the planner send in the equilibrium with the highest expected payoff to the planner?[4] There exists a pure-strategy Perfect Bayesian Nash Equilibrium (PBNE) in which the planner discloses no information by sending each agent the message $\{1, 2, 3\}$. In this uninformative equilibrium, the agents choose the same action in all states and the planner's ex-ante expected payoff is $\frac{8}{3}$.

However, the planner can do better than this by discriminating between the two agents. That is, even though the game is totally symmetric, there exists an equilibrium in which the planner earns an expected payoff of 3 by sending some information to one player and no information to the other player. Specifically, in this equilibrium there exists one player, say the row player, and one state, say $\omega = 3$, such that in state 3 the planner sends the message $\{3\}$ to the column player and the message $\{1, 2, 3\}$ to the row player, while in each of the other two states, the planner sends the message $\{1, 2\}$ to the column player and the message $\{1, 2, 3\}$ to the row player. The row player responds to any message containing state 2 by choosing $a = 2$, and he responds to the message $\{1, 3\}$ by choosing $a = 1$. The column player, on the other hand, plays $a = 3$ in response to the message $\{2, 3\}$, and he plays $a = 1$ in response to any message containing state 1. Both players respond to any message with only a single state by playing the dominant action in that state.

[4] The equilibria that we describe here do not rely on allowing any out-of-equilibrium beliefs. The particular restriction that we imposed will be described in Sect. 5. While it will play no role in our negative results, it will be used in the general construction of the equilibria of Propositions 1, 4 and 5.

To see that this is indeed an equilibrium, consider the row player first. His expected payoff in the proposed equilibrium is $(-1+5-1)/3 = 1$. If he chose a_1 in response to $\{1, 2, 3\}$, his payoff would be $(0+3-1)/3 = 2/3$. If he chose a_3 in response to $\{1, 2, 3\}$, his payoff would be $(-1+2+0)/3 = 1/3$. So he has no incentive to deviate from his response to the equilibrium message $\{1, 2, 3\}$. Next consider the column player. His expected payoff from playing a_1 in response to $\{1, 2\}$ is $(2-1)/2 = 1/2$. If he chose a_2 instead, his payoff would be $(1+0)/2 = 1/2$. If he chose a_3 instead, his payoff would be -1 for sure. So one of his best responses to $\{1, 2\}$ is to play a_1. Clearly, his best response to the singleton message $\{3\}$ is to play a_3.

Finally, consider the planner. First consider Type 1 planner. Following any deviation from his strategy, given the responses from row and column players, only action profiles (a_1, a_1) and (a_2, a_1) are possible, neither of which improves his payoff. Next consider Type 2 planner. If he sent the row player any message, that message would have to contain state 2 in which case the row player would keep on choosing a_2. If he sent the column player any message other than $\{1, 2\}$, then that message would either be $\{2\}$ or $\{2, 3\}$ or $\{1, 2, 3\}$. If he sends $\{2\}$, the column player would choose a_2 and the payoff to the planner would go down from 4 to 0. If the planner sent $\{2, 3\}$, the column player would choose a_3 and the planner's payoff would go down from 4 to 1. If the planner sent $\{1, 2, 3\}$, the column player would still choose a_1 and the planner's payoff would not change. Finally, consider Type 3 planner. This planner cannot change the action of the column player. The only way he can change the action of the row player is to either sent him the message $\{3\}$ or the message $\{1, 3\}$. In response to $\{3\}$, the row player would choose a_3 and the planner's payoff would go down from 4 to 0. If the planner sent the row player the message $\{1, 3\}$, the row player would choose a_1 and the planner's payoff would go down from 4 to 1.

It can be shown that this asymmetric equilibrium is the best equilibrium for the planner. In fact, there is no other equilibrium in which the planner sends some information to at least one of the players, and where at least one of the players responds to information by choosing different actions in at least two states (see the "Appendix").

Note that this example satisfies Assumptions (A1), (A2) (this is satisfied with a weak inequality that could be made strict without altering the example), (B1) and (B2'). However, it violates Assumption (B2).

In what follows we analyze the model presented in Sects. 1 and 2, which generalizes the above example. Our objective is to understand the nature of the best equilibrium for the planner, i.e., how much information does the planner want to disclose and does he treat the two agents symmetrically? A comparison with the cheap-talk approach is provided after presenting our results.

4.1 Completely uninformative equilibria

We can now present our first main result, which states a necessary and sufficient condition for the existence of a *completely uninformative equilibrium*. First, we present the following definition:

We shall say that a PBNE is *symmetric* when: (i) for all ω, the type ω planner sends the same message $S(\omega)$ to player 1 and player 2, and (ii) both players 1 and 2 use the same strategy: for all messages $S(\omega)$, choose action $a(S(\omega))$.

Proposition 1 *Assume (A1) and (B1) hold. Then there exists a symmetric PBNE in which the planner sends the uninformative message* Ω *in every state if and only if*

$$\sum_{k=1}^{K} u(k,k) \geq \sum_{k=1}^{K} u((k+\delta)_{\text{mod} K}, k) \tag{2}$$

for any integer δ.

Remark We can term this condition "dominant main diagonal." Indeed, if one writes the matrix of payoffs for player 1 and one copies the first $K-1$ rows below that matrix, the condition states that the sum of elements along the main diagonal exceeds the sum of all diagonals (of the expanded matrix) placed below it. Given that it is necessary for the existence of the uninformative equilibrium, this is the condition we shall focus on in the general model, if we wish to study when a planner can improve his payoff by revealing some information.

Proof of Proposition 1 Consider the following profile of strategies. Every planner type sends the message Ω to each of the players, and each player responds to Ω by choosing the same action. Assume, without loss of generality, that this action is K. For any message $S \neq \Omega$, call $\omega \in S$ the state such that $\omega' < \omega$ for all $\omega' \in S \setminus \{\omega\}$, let a player assign probability 1 to state ω and choose the dominant action in state ω. Note that this covers the case of singleton messages and also the case of any message containing state K, to which the player would respond by continuing to choose action K. We now show that this strategy profile is a PBNE under inequality (2).

Consider each player first. By (2), no player has an incentive to play $a < K$ in response to the message Ω, given that the other player responds with K. In addition, in every state $\omega > 1$ the planner has an incentive to try and change the players' actions. Hence, in particular, the players' out-of-equilibrium beliefs are consistent with $[\beta_1]$–$[\beta_2]$.

Consider the planner next. In each state ω the proposed strategy profile generates a payoff of $V(1+K-\omega, 1+K-\omega)$. The planner has no incentive to send both players the same message because they will respond to it by choosing the same lower action in that state: indeed, if the type ω planner sent the message $S \neq \Omega$ to both players with k' being the maximal state in S, the planner's resulting payoff would be $V(1+k'-\omega, 1+k'-\omega) \leq V(1+K-\omega, 1+K-\omega)$ by (B1).

Suppose now the planner deviates by sending different messages to either player: S_1 with maximal element k_1 to player 1, and S_2 with maximal element k_2 to player 2. Given the players' beliefs, the type ω planner's payoff would be $V(1+k_1-\omega, 1+k_2-\omega)$. Note that, since the planner cannot lie, $k_1 \geq \omega$ and $k_2 \geq \omega$. Hence, by (B1), the planner weakly prefers that both players choose $1+K-\omega$.

Conversely, if there exists a symmetric uninformative PBNE, then there exists some action a that both players choose in every state. Since no player has an incentive to deviate, inequality (2) must hold. □

If the uninformative equilibrium exists, it generates the following ex-ante expected payoff to the planner:

$$\bar{V} \equiv \frac{1}{K} \sum_{k=1}^{K} V(k, k)$$

We shall call this "the uninformative payoff".

4.2 When even partial information disclosure is impossible

Next, we turn to the investigation of equilibria in which the planner gives out some information. A PBNE is said to be *partially informative* if for some player, at least two types choose different actions in equilibrium. The next result offers an impossibility, i.e., for the planner to give out some information, asymmetries in the equilibrium will be required.

Proposition 2 *Assume (A1) and (B1) hold. Then there is no symmetric PBNE that is partially informative.*

Proof of Proposition 2 Assume, by contradiction, that there exists a symmetric PBNE with this property. Let a be an action that is played in equilibrium. By Lemma 1, there must be a state ω such that a is played by both players in ω and a is dominant in ω. Suppose $b \neq a$ is each player's response to the uninformative message Ω. Then by (B1), a type ω planner can profitably deviate by sending both players the message Ω. It follows that a must be the players' response to the message Ω.

By assumption, there exists another state in which the players in equilibrium also choose the same action $a' \neq a$. Again, by Lemma 1, both players choose a' in a state ω' in which a' is the dominant action. But then by (B1), a type ω' planner can profitably deviate by sending both players the message Ω. Since it cannot be the case that the players respond to Ω with both a and a', it follows that there cannot be a PBNE with the stated properties. □

Proposition 2 implies the following:

Corollary 1 *Assume (A1) and (B1) hold. There exists no PBNE with full information.*

Proof This follows since the full information outcome would obtain from a symmetric profile, in which each player is told the true state by the planner in each state ω. □

Note that Proposition 2 and its corollary remain true for any linear ordering of actions that satisfies Lemma 1. Since there are many orderings with this property, these results do not hinge on the particular structure (i.e., clockwise ordering of actions) that we imposed. This will also be true for our next result, which explores the consequence of our first strengthening of the assumption on the planners's payoff, i.e., Assumption (B2):

Proposition 3 *Assume (A1) and (B2) hold. There is no partially informative PBNE.*

Proof of Proposition 3 Suppose there exists a partially informative PBNE. By assumption, there exist two states in which one of the players, say player 1, chooses different actions, say a and a', in equilibrium. By Lemma 1 there exist a pair of states, call them (ω, ω'), such that a is dominant in ω and a' is dominant in ω'. Hence, $x(b, \omega) > x(a, \omega)$ for all $b \neq a$ and $x(b', \omega') > x(a', \omega')$ for all $b' \neq a'$.

Let b^* denote player 1's equilibrium response to the uninformative message Ω. If $b^* \neq a$, then (B2) implies that in state ω the planner can profit by sending the message Ω only to player 1, regardless of player 2's action in that state. If $b^* = a$, then again by (B2), in state ω' the planner can profit by sending Ω only to player 1. This means that at least in one state, the planner has a profitable deviation, which is a contradiction. □

Proposition 3 then implies that in the types of public good problems fitting Assumptions (A1) and (B2), there is no pure-strategy equilibrium that is better for the planner than the uninformative equilibrium.

4.3 When some information disclosure is possible

We explore next some of the circumstances under which it is in the interest of the planner to disclose some information in equilibrium. In doing so, consider instead a different strengthening of (B1): we shall assume (B2′) instead of (B2). We shall also strengthen (A1) into (A2). We note that these are sufficient conditions, but not necessary (we know we can prove the next results under somewhat weaker conditions).

The following lemma gives a hint about a basic feature of the equilibria we are seeking:

Lemma 2 *Assume (B2′). Any PBNE with partial information has the property that in every state where players choose different actions, one player actually chooses the dominant action.*

Proof Let ω be a state in which the two agents choose different actions a and b with $x(a, \omega) < x(b, \omega)$. We argue by contradiction. Then $x(a, \omega) > 1$. By (B2′), the planner can profitably deviate by sending the message $\{\omega\}$ to the agent choosing a, which is a contradiction. □

We now turn to show that by creating asymmetries between the agents, the planner can sustain a PBNE where some information is revealed. To clarify the role of asymmetry in our positive result, note that since (A2) and (B2') imply (A1) and (B1), we know that under the former pair of assumptions, there is no symmetric PBNE with information revealed. Therefore, asymmetry has to be a necessary ingredient in the construction of the positive results. Also a symmetric equilibrium means both symmetry in information disclosure and in players' strategies; on the other hand, it is not really meaningful to talk about symmetry in strategies when receivers are getting different messages (see the equilibrium strategies in Propositions 4 and 5).

To simplify the exposition, we assume in our next result that K is even.

Proposition 4 *Assume (A2) and (B2′) hold. Let K be even. Then there exists a partially informative equilibrium, which has the following structure. The planner gives out information so as to create the following partitions:* $\{\{1, 2\}, \{3, 4\}, \ldots, \{K-1, K\}\}$ *for player 1, and* $\{\{2, 3\}, \{4, 5\}, \ldots, \{K, 1\}\}$ *for player 2. Then, on the equilibrium path, players 1 and 2 choose action k following the information* $\{k, k+1\}$ *(note that this implies that, in each state, one player chooses the dominant action and the other chooses the highest action).*

Proof of Proposition 4 Suppose K is even. The planner's strategy is implicit in the statement of the proposition. As for the players' strategies, they are as follows. On the equilibrium path, they are also described in the statement. Off the equilibrium path, given any singleton message $\{k\}$, they choose the dominant action in that state. For any other message, player 1 assigns probability one to the *lowest odd* state, while player 2 assigns probability one to the *highest even* state, and their response is to play the corresponding dominant action. If a message contains only even states, player 1 assigns probability one to the *lowest even* state and chooses the dominant action for that state. Similarly, if a message contains only odd states, then player 2 chooses the dominant action for the *highest odd* state.

To check that this is a PBNE of this game, we define a player's type space to be the set of messages he receives in the proposed equilibrium. The planner's type is naturally defined to be the state of nature. We begin by verifying that no player type has any incentive to deviate. Consider type $\{k, k+1\}$of player 1. By playing $a = k$, this type obtains an expected payoff of

$$\frac{1}{2}u(1, K) + \frac{1}{2}u(K, 1) \tag{3}$$

Playing any $a' > 1$ would yield an expected payoff of

$$\frac{1}{2}u(a', K) + \frac{1}{2}u(a'-1, 1). \tag{4}$$

By (A2) applied to $k = K$, $\Delta_K(a') > \Delta_1(a'-1)$, i.e., $u(1, K) - u(a', K) > u(a'-1, 1) - u(K, 1)$, which implies that no deviation to taking action a' is profitable. Of course, the argument for player 2 is identical.

Consider now the planner, who receives a payoff of $V(1, K)$. From (B2'), it follows that the only way a planner of type k can increase his payoff is to induce both players to choose the action $k-1$. With the strategies written, we show this is impossible. Indeed, suppose first k is odd. To induce both players to choose an action a, the message that each player receives must contain both k and a. If the message to player 2 also contains an even action, then players would not choose the same action, as player 2 would choose an even action while player 1 would choose an odd action. This implies that a must be odd. This also means that the message sent to player 2 must contain only odd actions. Since player 1 would choose the lowest odd action in his message, while player 2 would choose the highest odd action, and since the message to each player must contain both k and a, the two players would not be able to choose a. By a similar argument, the planner cannot induce both players to choose the same action when k is even. It follows that in no state can the planner send an admissible message (i.e., a message containing that state) that will induce both players to choose the same action, unless it is the dominant action in that state. Hence, no planner type has any incentive to deviate.[5] □

An analogous result can be obtained when K is odd. We state it next, but its proof is relegated to the appendix.

Proposition 5 *Assume (A2) and (B2′) hold. Let K be odd. There exists a partially informative equilibrium, which has the following structure. The planner gives out information so as to create the following partitions:* $\{\{1, 2\}, \{3, 4\}, \ldots, \{K-2, K-1\}, \{K\}\}$ *for player 1, and* $\{\{2, 3\}, \{4, 5\}, \ldots, \{K-1, K, 1\}\}$ *for player 2. Then, on the equilibrium path, players 1 and 2 choose action k following the announcements* $\{k\}$, $\{k, k+1\}$ *or* $\{k, k+1, k+2\}$ *(note that this implies that in all states but one a player chooses the dominant action and the other chooses the highest action. In state 1, in contrast, one player chooses the dominant action and the other chooses the second highest action).*

Propositions 4 and 5 show that under certain conditions there exist partially informative equilibria, which are asymmetric even though the underlying game in each state is completely symmetric. In addition, the informational structures induced in these equilibria are similar in nature to the typical information structures in global games (see Morris and Shin (2003)).

4.4 Social evaluation of the partially informative equilibria

A natural question that arises is under what conditions would these equilibria maximize the ex-ante expected payoff to the planner. One could rephrase this question as follows: under what conditions would the informational structures that are typically assumed in global games could be explained as being induced by an informed planner in the equilibrium that is best for him?

To address this question we introduce the following notation. For any $M \subseteq \{2, \ldots, K-1\}$ let

$$W(M) = \frac{1}{|M|+2}\left[\sum_{m\in M} V(m, m) + V(1, 1) + V(K, K)\right].$$

Proposition 6 *Assume (B2′) holds. If K is even and $V(1, K) \geq W(M)$ for all $M \subseteq \{2, \ldots, K-1\}$, then there is no PBNE with a higher ex-ante payoff to the planner than the equilibrium described in Proposition 4. If K is odd and $V(1, K-1) \geq W(M)$ for*

[5] Note that this proof would also hold if we replaced the "passive beliefs" restriction β_1 with the alternative assumption of "paranoid beliefs" discussed in footnote. To see why, suppose player 1 gets an out-of-equilibrium message containing both odd and even states. Since he believes that the true state is the lowest odd state, he thinks this is the state that was revealed to the other player. He then best responds by choosing the dominant action.

all $M \subseteq \{2, \ldots, K-1\}$, *then there is no PBNE with a higher ex-ante payoff to the planner than the equilibrium described in Proposition 5.*

Proof of Proposition 6 Assume K is even and suppose, by contradiction, that there is an equilibrium with a higher ex-ante payoff to the planner. In that equilibrium, partition the set of states into two categories, those states in which players mis-coordinate (i.e., choose different actions) and those in which they coordinate (i.e., choose the same action). In the equilibrium of Proposition 4, the planner obtains a payoff of $V(1, K)$ in each state. Consider some state in the mis-coordination category and let $V(x, y)$ be the payoff to the planner in that state, where $x \geq 1$ and $y \leq K$. By (B2'), $V(1, K) \geq V(x, y)$. Let M be the set of states in the coordination category. Note that for every state in which the players coordinate on some action $x > 1$, there exists a state in which they coordinate on $x = 1$. To see why, assume, without loss of generality, that in state 1 the players play action $k > 1$. Obviously, the planner in state 1 must be sending a message to each player that contains both state 1 and state k, and hence, for each state in which the planner is receiving a payoff $V(k, k)$, there exists another state in which the planner is receiving $V(1, 1)$. Furthermore, if there exists a pair of states in which the players coordinate on different actions, then there exists a pair of other states where the players play the dominant action for those states. By assumption, $V(1, K) \geq W(M)$. Since each state is equally likely, the ex-ante payoff to the planner in the hypothesized equilibrium cannot be higher than $V(1, K)$. Essentially the same argument applies for the case of odd K. □

5 Verifiable versus cheap-talk messages

To better understand the role of the verifiability story that underlies our "hard facts" assumption for the planner's messages, and to contrast it with cheap-talk messages, we revisit the three-state example from Sect. 4, but perturb the payoffs slightly. Different perturbations show that the two approaches yield very different results, and we would expect a larger multiplicity of equilibria under cheap-talk, as is standard.

For instance, suppose that when a player chooses the *second lowest* action in a state, while the other player chooses the *highest* action in a state, the payoff to the former is 3 rather than 2. Then our original equilibrium is sustained under verifiable information, but *not* under cheap-talk. In state 1 the payoffs are as follows:

	a_1	a_2	a_3
a_1	$0, 0$	$2, -1$	$5, -1$
a_2	$-1, 2$	$1, 1$	$3, -1$
a_3	$-1, 5$	$-1, 3$	$3, 3$

Consider first the verifiable information case and the equilibrium strategies we described. The row player plays a_2 while the column player plays a_1, and the planner gets 1. In order to move the column player to a different action he must send him a message that does not contain the state 1, which is impossible. Given that the column player must choose a_1, the planner would like to move the row player to his third action. The only way to achieve this is to send him the singleton message $\{\omega_3\}$, which is impossible.

Consider next the case of cheap-talk. To sustain the equilibrium outcome described in the example, there must be two messages m_1 and m_3 such that the column player chooses a_3 for m_3 and a_1 for m_1. But this means that in state 1 the planner can profitably deviate by sending the column player the message m_3.

Since we have established that in state 1 the planner has a profitable deviation under cheap-talk but not under verifiability, it remains to show that with verifiable information the planner has no profitable deviation in states 2 and 3. Suppose the state is 2. Then the payoff matrix is

	a_1	a_2	a_3
a_1	3, 3	−1, 5	−1, 3
a_2	5, −1	0, 0	2, −1
a_3	3, −1	−1, 2	1, 1

In the verifiable case, our equilibrium calls for the row player to choose a_2 and for the column player to choose a_1. The planner is then getting a payoff of 4, and the only way he can improve is to have the row player choose a_1. But this is impossible, since by verifiability the planner must send a message containing the state 2, and any such message will induce the row player to choose a_2.

Suppose the state is 3. Then the payoff matrix is

	a_1	a_2	a_3
a_1	1, 1	3, −1	−1, 2
a_2	−1, 3	3, 3	−1, 5
a_3	2, −1	5, −1	0, 0

In the verifiable case, our equilibrium calls for the row player to choose a_2 and for the column player to choose a_3. The planner is then getting a payoff of 4, and the only way he can improve is to have the column player choose a_2. By verifiability, the planner must send the column player a message containing the state 3. However, any such message induces either a_1 or a_3.

Consider now a different perturbation. Indeed, one can also perturb the example so that an uninformative equilibrium is sustainable in cheap-talk but *not* with verifiable information. To see this, note first that under cheap-talk, an uninformative equilibrium is sustained by having both players choose the same action (say, action a_2) regardless of the message received. Suppose that the payoffs in state 1 are:

	a_1	a_2	a_3
a_1	0, 0	4, −1	5, −1
a_2	−1, 4	1, 1	3, −1
a_3	−1, 5	−1, 3	3, 3

Suppose there exists an uninformative equilibrium in which both players choose a_2 in every state. Then in state 1 the planner gets a payoff of 2. Under verifiable information, the planner can reveal to the column player that the state is 1 (by sending him the message $\{\omega_1\}$). In any PBNE, the column player must best-respond by choosing a_1. But this gives the planner a payoff of 3, which is higher than what he would receive in the uninformative equilibrium.

6 Related literature

Our paper is closely related to the cheap-talk literature that studies games between an informed "sender" and uninformed "receivers" .[6] 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.[7] 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.[8]

Another related literature analyzes voluntary disclosure of verifiable information (e.g., Milgrom and Roberts (1986); Okuno-Fujiwara et al. 1990). 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.[9]

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; 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

[6] For a comprehensive survey on cheap-talk games see Krishna and Morgan (2007).

[7] 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.

[8] Their paper also illustrates the optimality of asymmetric treatment of the agents.

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

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

Since we model the interaction between the receivers as a generalized prisoners' dilemma, our paper is naturally related to the literature on public goods games. When there is no threshold or provision point, voluntary contribution games typically reduce to a prisoners' dilemma in which zero contribution is a dominant strategy. In particular, our paper is closely related to Teoh (1997), which examines a public goods game in which the payoffs depend on a state of nature. That paper compares the equilibrium outcomes under two extreme regimes, one in which the players are perfectly informed of the state and one in which they are completely uninformed. Teoh (1997) shows that under certain conditions, the equilibrium with uninformed players is Pareto superior to the equilibrium with informed players. However, in contrast to us, Teoh (1997) does not examine what information structures would arise in equilibrium when the informed party is a player.

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); 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 et al. (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.

7 Concluding remarks

We have studied the problem of information disclosure for a planner who faces two agents interacting in a state-dependent generalized prisoners' dilemma. We have found conditions under which the planner can make use of his superior information by giving some of it out to the agents, and conditions under which such information leakage is not possible. We remark that, although the problem is entirely symmetric, the planner's *only* way to reveal part of the information is based on creating asymmetries between the two agents by giving them different pieces of information. We have also found conditions under which such partially informative equilibria are the best equilibria from the planner's point of view.

In our study, we have assumed that the number of actions—and states– was at least three. The two-action case, which depicts the standard prisoners' dilemma, yields the following. Under the "dominant main diagonal" condition 2 in Proposition 1, the completely uninformative equilibrium exists, but no other equilibrium can be found to dominate it. If this condition is violated, there is no equilibrium in pure strategies. Note, though, that the violation of the dominant-diagonal condition is also a violation of (B1), implying that cooperation on the non-dominant action is not efficient, moving us far afield from a prisoners' dilemma. Thus, the two-action case gives one a vastly simplified picture of the problem: under the assumptions that fit a prisoners' dilemma scenario, information disclosure is impossible. In contrast, while the same negative message is generally found in the multi-action case, there are (somewhat more special) conditions under which partial disclosure may occur in equilibrium.

At the beginning of the introduction we compared our approach to implementation theory. Both approaches make polar opposite assumptions on the information held by the planner. While implementation theory is concerned with how to design mechanisms in order to elicit information from the agents towards a socially desirable goal, our paper is about how much information should the planner leak in order to maximize his payoff when the institution –strategic situation—is not something that he can change. Therefore, both can be viewed as complementary approaches. Of course, while implementation theory has achieved an impressive degree of generality, our approach is much more limited: we analyze a specific strategic situation—a class of generalized prisoners' dilemma games–; one desirable next step would be to study the planner's information leakage problem in any arbitrary game played by the agents, a step clearly beyond our scope here.

8 Appendix

Claim *In the example of Sect. 4, there is no other equilibrium in which the planner sends some information to at least one of the players, and where at least one of the players responds to information by choosing different actions in at least two states.*

Proof The proof proceeds by a series of steps.

Step 1. There exists no equilibrium in which some player has full information.

Assume there exists such an equilibrium. Without loss of generality assume it is the column player. Consider the column player's equilibrium response to the message $\{1, 2, 3\}$. Without loss of generality assume his response is a_1 (the left-most column). Then in state 2, the row player must be choosing a_3, or the bottom row. Otherwise, the planner would have an incentive to deviate only to the column player by sending him the message $\{1, 2, 3\}$. To see why, note that in state 2, the column player, who is fully informed, must be playing his dominant action, which is a_2 or the middle column. If the row player were to choose the middle or top row the payoff to the planner would be strictly lower than what he could get if only the column player deviated and choose the left-most column (the planner would get either 6 instead of 4 or 4 instead of 0).

By a similar argument, the planner does not have an incentive to send the column player the message $\{1, 2, 3\}$ in state 3 only if the row player chooses a_2 (or the middle row) in this state. But if the row player is choosing different actions in states 2 and 3, it must be that he has full information in one of these states. But if this was true, he would choose either dominant action in state 2 or the dominant action in state 3, a contradiction.

Step 2. There exists no equilibrium in which both players have the same partial information, and both respond to the information they have.

Assume there exists such an equilibrium. Without loss of generality assume that both players know only whether or not state 3 has occurred. Then both players choose their third action in state 3. By assumption, both players respond to the information they receive, hence neither player chooses his third action (the state 3 dominant action) in either state 1 or 2. Therefore, it cannot be the case that both players respond to the message $\{1, 2, 3\}$ by choosing the state 3 dominant action. Otherwise, in state 1 the planner could profitably deviate by sending both players the message $\{1, 2, 3\}$. But this means that the planner can profitably deviate in state 3. The reason is that in state 3 the planner is getting zero since both players are choosing their dominant action. If one of the players responds to $\{1, 2, 3\}$ with an action different than the dominant action, the planner would be able to get a strictly positive payoff by sending only to that player the message $\{1, 2, 3\}$.

Step 3. There exists no equilibrium in which both players have different partial information, and both respond to the information they have.

Assume there exists such an equilibrium. Assume without loss of generality that the row player knows only whether or not state 3 has occurred, while the column player knows only whether or not state 1 has occurred. It follows that the row player chooses the bottom row in state 3, while the column player chooses the left-most column in state 1. The proof follows from the following series of observations.

Observation 1 *No player to respond to the message* $\{1, 2, 3\}$ *with the dominant action of the state in which he is informed.*

Suppose the claim was false. Suppose first that the row player responds to $\{1, 2, 3\}$ by playing the third row. Since by assumption, the row player responds to the information he receives, he does not choose the third row in state 1. Because the column player chooses the left-most column in that state, it means that the planner is getting a payoff which is lower than 4. But by sending the message $\{1, 2, 3\}$ a planner of type 1 could profitably deviate. Suppose the column player responded to $\{1, 2, 3\}$ by choosing the left-most column. Since both players respond to their partial information, the outcome in state 2 must be in the sub-matrix formed by the middle and top row and the middle and right column. Regardless of what the outcome is in this sub-matrix, a type 2 planner can profitably deviate by sending only to the column player the message $\{1, 2, 3\}$.

Observation 2 *In state 1 the row player chooses the middle row.*

Assume not. Since the row player responds to his information, he does not choose the bottom row in state 1. If he were to choose the top row, the payoff to the planner in state 1 would have been zero. By Observation 1, the column player responds to the message $\{1, 2, 3\}$ by choosing either the middle or right column. But this means that the planner can profitably deviate in this state by sending $\{1, 2, 3\}$ only to the column player. It follows that the row player must be choosing the middle row in state 1.

Observation 3 *In state 2 the row player chooses the middle row.*

This follows from Observation 2 and the fact that the row player does not know whether state 1 or 2 has occurred, hence he must be choosing the same action in both states.

Observation 4 *In state 2 the column player chooses the right column.*

If this was false, then the column player must be choosing the middle column (because the column player responds to the partial information he has, he does not choose the left column in state 2). By Observation 3, the row player also chooses his middle action in this state. Hence, the planner's payoff in this state is zero. In order for a type 2 planner not to have an incentive to deviate and send the message $\{1, 2, 3\}$ to either player, it must be that both players respond to $\{1, 2, 3\}$ by choosing their middle action. But this means that a type 1 planner has an incentive to deviate. To see why, note that the column player chooses his dominant action in that state and by claim 2, the row player chooses his middle action in that state. This means that the planner's payoff in that state is 1. If a type 1 planner sent both players the message $\{1, 2, 3\}$, both would choose their middle action, and the planner would obtain a payoff of 2.

Observation 5 *In state 3 the column player chooses the right column.*

This follows from Observation 4 and the fact that the column player does not know whether state 2 or 3 has occurred, hence he must be choosing the same action in both states.

Observation 6 *Type 3 planner has a profitable deviation.*

Since the row player is fully informed in state 3, he chooses the bottom row. By Observation 5, the column player chooses the right column in this state. Hence, type 3 planner receives a payoff of zero. By Observation 1, the row player responds to $\{1, 2, 3\}$ with either the top or middle row. This means that by sending $\{1, 2, 3\}$ only to the row player, type 3 planner can obtain a strictly positive payoff. □

Proof of Proposition 5 Suppose K is odd. The planner's strategy is implicit in the statement of the proposition. As for the players' strategies, they are as follows. On the equilibrium path, they are also described in the statement. For any other message, player 1 assigns probability one to the *lowest odd* state, while player 2 assigns probability one to the *highest even* state, and their response is to play the corresponding dominant action. If a message contains only even states, player 1 assigns probability one to the *lowest even* state and chooses the dominant action for that state. Similarly, if a message contains only odd states, then player 2 chooses the dominant action for the *highest odd* state, except for the case where the message contains both the lowest and the highest states, in which case player 2 chooses the dominant action for the second highest state.

To check that these strategies constitute a PBNE of this game, we define a player's type space to be the set of messages he receives in the proposed equilibrium. The planner's type is naturally defined to be the state of nature. We begin by verifying that no player type has any incentive to deviate. First, type $\{K\}$ of player 1 clearly has no incentive to deviate as he chooses a dominant action. Next, consider type $\{1, 2\}$ of player 1. By playing $a = 1$ type $\{1, 2\}$ obtains an expected payoff of

$$\frac{1}{2}u(1, K - 1) + \frac{1}{2}u(K, 1)$$

while playing any $a > 1$ would yield an expected payoff of

$$\frac{1}{2}u(a - 1, K - 1) + \frac{1}{2}u(a - 2, 1)$$

By (A2) applied to $k = K - 1$,

$$u(1, K - 1) - u(a - 1, K - 1) > u(a - 2, 1) - u(K, 1)$$

Finally, consider a type $\{k, k + 1\}$ of player 1. By playing $a = k$, this type obtains an expected payoff of

$$\frac{1}{2}u(1, K) + \frac{1}{2}u(K, 1)$$

Playing any $a' \neq k$ would yield an expected payoff of

$$\frac{1}{2}u(a', K) + \frac{1}{2}u(a' - 1, 1).$$

By (A2) applied to $k = K$, $u(1, K) - u(a', K) > u(a' - 1, 1) - u(K, 1)$, which implies that no deviation to taking action a' is profitable.

To see that player 2 has no incentive to deviate, note first that any type $\{k, k + 1\}$ of player 2 has no incentive to deviate for the same reason that a type $\{k, k + 1\}$ of player 1 has no incentive to deviate. Consider then type $\{K - 1, K, 1\}$ of player 2. By playing $K - 1$, this type obtains an expected payoff of

$$\frac{1}{3}u(1, K) + \frac{1}{3}u(K, 1) + \frac{1}{3}u(K - 1, 1)$$

Playing any $a \neq K - 1$ would yield an expected payoff of

$$\frac{1}{3}u(a, K) + \frac{1}{3}u(a - 1, 1) + \frac{1}{3}u(a - 2, 1)$$

By (A2) applied to $k = K$,

$$u(1, K) - u(a, K) > [u(a - 1, 1) - u(K, 1)] + [u(a - 2, 1) - u(K - 1, 1)]$$

for all $a \neq K - 1$.

Consider now the planner, who receives a payoff of $V(1, K)$ in every state except for state 1, where he receives a payoff of $V(1, K-1)$. In every state other than 1, the planner can increase his payoff by inducing the players to choose the same high enough action a. In state 1, the planner can also profit by inducing maximum separation (one player choosing the lowest and the other the highest action).

Consider first planners of type $\omega > 1$. The reason this planner cannot Get both players to choose the same action follows from essentially the same arguments given in the proof of Proposition 4. The only modification that needs to be made is when player 2 receives a message with only odd states, and which contains states 1 and K. If the message contains at least one other state, then the two players would choose different actions. If the message is $\{1, K\}$, then both players would choose action 1. However, for the planner types we are currently concerned with, this message can only be sent in state K, and then, by (B2'), the planner prefers the equilibrium actions over both players choosing action 1: $V(1, K) > V(2, 2)$.

Finally, consider the type 1 planner. Note that player 1 will always choose action 1, for any message sent by the planner. Getting the two players to choose the same action is impossible, unless player 2 is sent the message $\{1, K\}$, but then both would be choosing action 1, which is worse for the planner than the equilibrium. And getting player 2 to choose action K to induce maximum separation is also impossible, given the strategies written.

It follows that the planner has no profitable deviation.

References

Angeletos G-M, Pavan A (2007) Efficient use of information and social value of information. Econometrica 75:1103–1142

Azacis H, Vida P (2012) Collusive communication schemes in a first-price auction, Working paper.

Baliga S, Sjöström T (1999) Interactive implementation. Games Econ Behav 27:38–63

Baliga S, Corchón L, Sjöström T (1997) The theory of implementation when the planner is a player. J Econ Theory 77:15–33

Farrell J (1986) Voluntary disclosure: Robustness of the unraveling result, and comments on its importance, Chap. 4. In: Grieson R (ed) Antitrust and regulation, Lexington, Lexington (Mass.), pp 91–103

Farrell J, Gibbons R (1989) Cheap talk with two audiences. Am Econ Rev 79:1214–1223

Gentzkow M, Kamenica E (2011) Bayesian Persuasion. Am Econ Rev 101:2590–2615

Goltsman M, Pavlov G (2011) How to talk to multiple audiences. Games Econ Behav 72:100–122

Jewitt I, Li D (2012) Cheap-talk information disclosure in auctions, Working paper.

Hörner J, Skrzypacz A (2011) Selling information, Working paper.

Koessler F (2008) Lobbying with two audiences: public vs private certification. Math Soc Sci 55:305–314

Krishna V, Morgan J (2007) " Cheap Talk". In the new Palgrave dictionary of economics, 2nd Edn.

McAfee P, Schwartz M (1994) Opportunism in multilateral vertical contracting; nondiscrimination, exclusivity, and uniformity. Am Econ Rev 84:210–230

Morris S, Shin HS (2002) The social value of information. Am Econ Rev 92:1521–1534

Morris S, Shin HS (2003) Global games: theory and applications. In: Dewatripont M, Hansen L, Turnovsky S (eds) Advances in economics and econometrics (Proceedings of the Eighth world congress of the econometric society) , Cambridge University Press, Cambridge.

Okuno-Fujiwara M, Postlewaite A, Suzumura K (1990) Strategic information revelation. Rev Econ Stud 57:25–47

Rayo L, Segal I (2010) Optimal information disclosure. J Polit Econ 118:949–987

Shavell S (1989) A note on the incentive to reveal information. Geneva Pap Risk Insur 14:66–74

Stein JC (1989) Cheap talk and the fed: a theory of imprecise policy announcements. Am Econ Rev 79:32–42

Teoh SH (1997) Information disclosure and voluntary contributions to public goods. RAND J Econ 28: 385–406