Belief-free equilibria in games with incomplete information: the N-player case

Johannes Hörner*, Stefano Lovo†, Tristan Tomala‡

March 11, 2009

Abstract

We characterize belief-free equilibria in infinitely repeated games with incomplete information with an arbitrary number of players and arbitrary information structures. This generalizes Hörner and Lovo (2008), which restrict attention to the two-player case and to information structures that have a product structure. Our characterization requires introducing a new type of individual rational constraint that links the lowest possible equilibrium payoffs across players. As in the two-player case, our characterization is tight: we define a set of payoffs that contains all the belief-free equilibrium payoffs; conversely, any point in the interior of this set is a belief-free equilibrium payoff vector when players are sufficiently patient.

Keywords: repeated game with incomplete information; Harsanyi doctrine; belief-free equilibria.

JEL codes: C72, C73

1 Introduction

This paper characterizes the set of payoffs achieved by equilibria that are robust to the specification of beliefs. We consider n-player repeated games with incomplete information and low discounting. This class of equilibria has been introduced by Hörner and Lovo

*Yale University. johannes.horner@yale.edu.
†HEC School of Management, Paris and GREGHEC, 78351 Jouy-en-Josas, France. lovo@hec.fr.
‡HEC School of Management, Paris and GREGHEC, 78351 Jouy-en-Josas, France. tomala@hec.fr.

(2008) in two-player games with incomplete information, as defined by Aumann and Maschler (1995). A strategy profile is a *belief-free equilibrium* if, after every history, every player's continuation strategy is optimal, given his information, and *independently* of the information held by the other players. That is, it must be a subgame-perfect equilibrium for every game of complete information that is consistent with the player's information.

Such equilibria offer several advantages. From a practical point of view, they do not require the specification of beliefs after all possible histories, and the verification of their consistency with Bayes' rule. From a theoretical point of view, they represent a stringent refinement, in the sense that such equilibrium outcomes are also equilibrium outcomes for every Bayesian solution concept, such as sequential equilibrium, for instance. But more importantly, these equilibria do not rely on the Bayesian paradigm. To predict behavior in environments with unknown parameters, a model typically includes a specification of the players' subjective probability distributions over these unknowns, following Harsanyi (1967-1968). Since beliefs are irrelevant here, belief-free equilibria do not require that players share a common prior, or that they update their beliefs according to Bayes' rule; and they remain equilibria even if players receive additional information as the game unfolds.

Nevertheless, as in the case of games with perfect information, players may randomize, and they maximize their expectation with respect to such lotteries.[1] Belief-free equilibria require precisely as much probabilistic sophistication as is usually assumed in games with perfect information.

In Hörner and Lovo (2008), the analysis is restricted to two-player games, and information has a product structure. That is, the information structure can be represented as

[1]This is also the standard assumption used in the literature on 'non-Bayesian' equilibria (see, for instance, Monderer and Tennenholtz, 1999).

a matrix. Each state of nature corresponds to a cell in this matrix. Player 1 knows the row, and Player 2 knows the column. This paper generalizes the results along those two dimensions:

1. There are $N \geq 2$ players, rather than only two players.

2. Arbitrary finite information structures are considered. In particular, the players' combined information may not pin down the (payoff-relevant) state of nature.

This latter generalization requires an appropriate extension of the definition of belief-free equilibrium. We choose the most restrictive version, and require players to use strategies that are best-replies independently of the state of nature, even for those states that cannot be identified by the players' combined information. Clearly, such an equilibrium remains an equilibrium for weaker versions of this definition. For instance, one may wish to assume that each player has a subjective probability distribution over those states of nature that the players' combined information cannot distinguish. We do so for both practical and theoretical reasons. From a practical point of view, it is immediate to modify our results to cope with weaker definitions, by replacing for instance such collections of states by a single state, and payoffs in that state by the relevant expectations. From a theoretical point of view, it is unclear to us for what reason the optimality criterion used by an agent should distinguish between uncertainty that can or cannot be identified collectively.

The focus of the analysis is on the set of belief-free equilibrium payoffs as the discount factor tends to one. We provide a set of necessary conditions that defines a closed, convex, and possibly empty set. These necessary conditions have simple interpretations in terms of incentive compatibility, individual rationality, and *joint rationality*, an additional requirement absent from the earlier analysis for two-player games, and that is related to the fact that, when play is supposed to reveal the players' private information, inconsistencies might arise for which it is not possible to identify a single deviator. Conversely, we prove that every payoff vector in the interior of this set is a belief-free equilibrium payoff provided the discount factor is sufficiently close to one.

As mentioned, this set of payoffs might be empty, and therefore, belief-free equilibria need not exist. We provide sufficient conditions for non-emptiness. With two players, for instance, this was known to be the case if there are two states only, or if each player knows his own payoff, and only one player has private information. We provide necessary and sufficient conditions on the information structure for non-emptiness, as a function of the types of payoffs that are considered. For general payoff functions, it must be that no two players are essential to distinguish between any two states. If the payoff functions are such that some action profile yields a payoff no larger than the individually rational payoff (the *bad outcome* property), for all players and for all states simultaneously, then it must be that no single player is essential to distinguish between any two states. With known-payoffs, there is a slight gap between our necessary and sufficient conditions. Non-emptiness requires that no player be essential, and this is also sufficient if payoffs further satisfy the bad outcome property. Without this further requirement, a sufficient condition is that information be embedded, in the sense that one player has superior information to another player, who in turn has superior information to all other players. This generalizes the one-sided information condition with two players only. We do not know whether this sufficient condition can be weakened further.

A special class of games covered by this sufficient condition is the class of 'reputation' games in which there is one player whose payoff type is unknown. We identify the value of reputation for such games. Consider the lowest belief-free equilibrium payoff that this player can guarantee for a given set of alternative payoff types he might be. We identify the highest such payoff, across all sets of alternative types, and identify a set of types achieving this maximum.

In the case of two players, the set of belief-free equilibrium payoffs had already appeared in the literature, most notably (but not only) in the context of undiscounted Nash equilibrium payoffs for games with one-sided incomplete information. See, among others, Cripps and Thomas (2003), Forges and Minelli (1997), Koren (1992) and Shalev (1994). The most general characterization of Nash equilibrium payoffs is obtained by Hart (1985) for the case of one-sided incomplete information. A survey is provided by Forges (1992). Israeli (1999) provides an analysis of reputation in two-player undiscounted games, from which our proofs in the section on reputations are inspired. Further references to non-Bayesian studies can be found in Hörner and Lovo (2008).

The concept of belief-free equilibrium is also related to the ex post equilibrium that is used in mechanism design (see Crémer and McLean, 1985) as well as in large games (see Kalai, 2004). A recent study of ex post equilibria and related belief-free solution concepts in the context of static games of incomplete information is provided by Bergemann and Morris (2007).

The concept of belief-free equilibria has been introduced in games with imperfect monitoring. See Piccione (2002) and Ely and Välimäki (2002) and Ely, Hörner and Olszewski (2005), among others. In this literature, belief-free equilibria are defined as equilibria for which continuation strategies are optimal independently of the private history observed by the other players, and has allowed the construction of equilibria in cases in which only trivial equilibria were known so far.

Section two introduces the notation and defines belief-free equilibria. Section three gives necessary conditions that belief-free equilibrium payoffs must satisfy. Section four shows that every payoff vector in the interior of the set defined by the necessary conditions is indeed a belief-free equilibrium payoff vector for low enough discounting. Section five provides necessary and sufficient conditions for non-emptiness of this set. Section six applies the previous results to games of reputation with one informed player.

2 Notations

The set of players is $N = \{1, \ldots, N\}$. Player i chooses action a_i from a finite set A_i, and $a \in A := \prod_i A_i$ is an action profile. The finite state space is $K = \{1, \ldots, K\}$. Given a set S, let $\triangle S$ denote the probability simplex over S, $1\{S\}$ the indicator function of S, $|S|$ the cardinality of S, $\operatorname{int} S$ the interior of S, and $\operatorname{co} S$ the convex hull of S. To avoid trivialities, assume that $|A_i| \geq 2$, all $i \in N$.

Player i's reward function is a map $u_i : K \times A \to \mathbb{R}$. Let $M := \max_{i \in N, k \in K, a \in A} |u_i(k, a)|$. A reward profile is denoted $u := (u_1, \ldots, u_N)$. Mixed actions of player i are denoted α_i. The definition of rewards is extended to mixed, possibly correlated, action profiles $\mu \in \triangle A$ in the usual way.

At the beginning of the game, each player receives once and for all a signal that allows him to narrow down the set of possible states of nature. This can be represented by an information structure is $\mathcal{I} = (\mathcal{I}_1, \ldots, \mathcal{I}_N)$, where $\mathcal{I}_i$ denotes player i's information partition of K. We let $I_i(k)$ denote the element of $\mathcal{I}_i$ containing k. We refer to $I_i(k) =: \theta_i \in \Theta_i$ as player i's *type*, and write $\Theta := \prod_i \Theta_i$, and $\Theta_{-i} := \prod_{j \neq i} \Theta_j$. Given $\theta \in \Theta$, $\kappa(\theta) := \bigcap_{i \in N} \theta_i$ denote the set of states that are consistent with type profile θ. Also, for $\theta_{-i} \in \Theta_{-i}$, we write $\kappa(\theta_{-i}) := \bigcap_{j \neq i} \theta_j$ for the set of states that are consistent with a type profile of the set of players different from i. We do not require that $\kappa(\theta) \neq \varnothing$: it might be that some type profile cannot arise. Similarly, we do not require that the join of the players' information partitions need reveal the state i.e., we allow for $|\kappa(\theta)| > 1$. The information partitions are common knowledge, but the realized signal is private information.

The game is infinitely repeated, with periods $t = 0, 1, 2, \ldots$ A history of length t is a vector $h^t \in H^t := A^t$ ($H^0 := \{\varnothing\}$). An outcome is an infinite history $h \in H := A^\infty$. A behavior strategy for player i's type θ_i is a mapping $\sigma_{i,\theta_i} : \cup_{t\in\mathbb{N}} H^t \to \triangle A_i$. We write $\sigma_i := \{\sigma_{i,\theta_i}\}_{\theta_i\in\Theta_i}$ for player i's strategy, and $\sigma := (\sigma_1, \ldots, \sigma_N)$ for a strategy profile.

Conditional on a state, players maximize their *payoff*, namely the expected average discounted sum of rewards, where the expectation is taken with respect to mixed action profiles. Players use a common discount factor $\delta < 1$. That is, given some outcome $\{a_t\}_{t\in\mathbb{N}_0}$, player i's payoff in state k is

$$\sum\nolimits_{t\geq 0}(1-\delta)\delta^t u_i(k, a_t).$$

As usual, the domain of rewards is extended to strategy profiles. Neither mixed actions nor realized payoffs are observed. On the other hand, realized actions are perfectly observed. Given a strategy profile σ, let $\mu_k \in \triangle A$ denote the occupation measure over action profiles induced by σ when the state is k, that is

$$\mu_k(a) := (1-\delta)\,\mathbb{E}_\sigma\left[\sum\nolimits_{t\geq 0}\delta^t 1\{a_t = a\}\right]\ \forall a \in A.$$

Let $u(k, \mu_k) \in \mathbb{R}^N$ denote the players' payoffs in state k under the occupation measure μ_k:

$$u(k, \mu_k) := \sum\nolimits_{a\in A}\mu_k(a)u(k, a).$$

Definition: A belief-free equilibrium (hereafter, an equilibrium) is a strategy profile σ such that, for every state k, σ is a subgame-perfect Nash equilibrium of the game with rewards $u(k,\cdot)$. A vector $v \in \mathbb{R}^{NK}$ is an equilibrium payoff vector if there exists a equilibrium σ such that $v = u(\sigma)$.

In what follows, we write v^k for the payoff vector in state k. Let B_δ be the set of belief free equilibrium (BFE) payoff vectors of the δ-discounted game. The purpose of this paper is to characterize $\lim_{\delta\to 1} B_\delta$.

3 Necessary Conditions

We first derive necessary conditions for a vector $v \in \mathbb{R}^{NK}$ to be an equilibrium payoff vector. These conditions can be divided into three categories: feasibility, individual rationality, and incentive compatibility.

3.1 Feasibility

The payoff vector $v \in \mathbb{R}^{NK}$ is *feasible* if there exists $(\mu_k)_{k\in K} \in (\triangle A)^K$ such that

1. $\forall k \in K : v^k = u(k, \mu_k)$;

2. $\forall k, k'$: $I_i(k) = I_i(k')\ \forall i \in N \Rightarrow \mu_k = \mu_{k'}$.

The first condition states that an equilibrium payoff must be feasible. That is, there exists an occupation measure μ_k that yields the payoff v^k.

The second condition states that if the join of the players' information partitions does not distinguish between two states, then the equilibrium strategies σ cannot either, and the equilibrium occupation measures over action profiles that determines the payoff vector must be the same for both states. Given the second condition, we may alternatively write μ_θ for the occupation measure. Conversely, throughout the paper, the notation $(\mu_\theta)_{\theta \in \Theta}$ implies that the set $(\mu_k)_{k \in K}$ satisfies the second condition.

3.2 Incentive Compatibility

If two signals θ_i and θ_i' are both consistent with a signal profile θ_{-i} of the other players, it must be the case that player i weakly prefers the occupation measure $\mu_{\theta_i,\theta_{-i}}$ to $\mu_{\theta_i',\theta_{-i}}$ for every state that is possible given (θ_i, θ_{-i}). Therefore, if v is an equilibrium payoff vector, then it must be feasible for some probability distributions satisfying a set of incentive compatibility conditions.

To introduce those, define UD_i (for unilateral deviation) as the set of triples $(\theta_i, \theta_i', \theta_{-i}) \in \Theta_i \times \Theta_i \times \Theta_{-i}$ such that $\kappa(\theta_i, \theta_{-i}) \neq \varnothing$ and $\kappa(\theta_i', \theta_{-i}) \neq \varnothing$. The incentive compatibility conditions can be written as

$$\forall i, (\theta_i, \theta_i', \theta_{-i}) \in UD_i, k \in \kappa\,(\theta_i, \theta_{-i}) : u_i(k, \mu_{\theta_i,\theta_{-i}}) \geq u_i(k, \mu_{\theta_i',\theta_{-i}}). \quad (IC(i, \theta_i, \theta_i', \theta_{-i}))$$

Lemma 3.1 *If $v \in B_\delta$, then v is feasible for some $(\mu_\theta)_{\theta \in \Theta}$ that satisfy $IC(i, \theta_i, \theta_i', \theta_{-i})$ for all $i \in N$ and $(\theta_i, \theta_i', \theta_{-i}) \in UD_i$.*

Proof: Suppose for the sake of contradiction that for some $i \in N$ and $(\theta_i, \theta_i', \theta_{-i}) \in UD_i$, the reverse inequality holds. Suppose now that the state is k and consider player i of type θ_i. By playing as if his type were θ_i', player i can guarantee $u_i(k, \mu_{\theta_i',\theta_{-i}})$, which exceeds his equilibrium payoff $u_i(k, \mu_{\theta_i,\theta_{-i}})$. This is a profitable deviation. $\square$

3.3 Individual Rationality

A deviating player might be easy to identify or not. For instance, if player i chooses an action that is inconsistent with all his types' equilibrium strategies, then it is immediately common knowledge among players that i deviated.

Define, for $\theta_{-i} \in \Theta_{-i}$,

$$\varphi_{i,\theta}(q) := \min_{\alpha_{-i} \in \prod_{j \neq i} \triangle A_j} \max_{a_i \in A_i} \sum\nolimits_{k \in \kappa(\theta_{-i})} q(k) u_i(k, \alpha_{-i}, a_i).$$

For each player i and each $\theta_{-i} \in \Theta_{-i}$, consider the set of inequalities

$$\forall q \in \triangle\kappa(\theta_{-i}) : \ \sum\nolimits_{k \in \kappa(\theta_{-i})} q(k) v_i^k \geq \varphi_{i,\theta}(q). \qquad (IR(i, \theta_{-i}))$$

Note that if $\kappa(\theta_{-i}) = \varnothing$, the inequality is vacuously satisfied. These inequalities are the immediate generalizations of the individual rationality conditions for the two-player case. In the definition of $\varphi_{i,\theta}$, note that players' actions are statistically independent.

Lemma 3.2 *If $v \in B_\delta$, it satisfies the inequalities $(IR(i, \theta_{-i}))$ for each player i and θ_{-i}.*

Proof: If these conditions are violated, there necessarily exists one player, a type profile θ_{-i} and $q \in \triangle\kappa(\theta_{-i})$ such that the reverse inequality holds. This implies that for every α_{-i}, there exists $a_i(\alpha_{-i})$ such that

$$\sum\nolimits_{k \in \kappa(\theta_{-i})} q(k) u_i(k, \alpha_{-i}, a_i(\alpha_{-i})) > \sum\nolimits_{k \in \kappa(\theta_{-i})} q(k) v_i^k. \tag{1}$$

Assume by contradiction that v is in B_δ and let σ be the corresponding equilibrium. Note that players $-i$ play the same strategy in each state in $k \in \kappa(\theta_{-i})$. Consider thus the strategy τ_i of player i that plays $a_i(\alpha_{-i})$ after a history h such that $\sigma_{-i}(h) = \alpha_{-i}$. The stage expected payoff of player i under (τ_i, σ_{-i}) satisfies the inequality (1) and therefore, so does the discounted payoff. It follows that there exists a state $k \in \kappa(\theta_{-i})$ at which τ is a profitable deviation. $\square$

Under these conditions, following Blackwell (1956), players $-i$ can devise a punishing strategy against player i. The reasoning is exactly the same as in the two-player case. Given θ_{-i}, and any payoff vector v that satisfies these inequalities strictly, there exists $\varepsilon > 0$ and a strategy profile $\widehat{s}_{-i}^{\theta_{-i}}$ for players $-i$ such that, if players $-i$ use $\widehat{s}_{-i}^{\theta_{-i}}$, then player i's undiscounted payoff in any state k that is consistent with θ_{-i} is less than $v_i^k - \varepsilon$ in any sufficiently long finite-horizon version of the game, no matter i's strategy. By continuity, this also holds true for sufficiently long finite-horizon versions of the game when payoffs are discounted, provided the discount factor is high enough, fixing the length of the game. When players $-i$ use $\widehat{s}_{-i}^{\theta_{-i}}$, players $-i$ are said to *minmax* player i. Player i is the *punished* player, and players $-i$ are the *punishing* players.

However, it might be that i's action is consistent with some of his types' strategies, and so is player j's action, but no pair of types for which both actions would be simultaneously consistent exists. Then it is common knowledge that some player deviated, but not whether it is player i or j. With at least three players, because of this identifiability issue, we have another condition that links the payoffs of different players. Let D be the set of type profiles that are impossible, but consistent with some state of nature and a unilateral deviation. That is, θ is in D if $\kappa(\theta) = \varnothing$ and $\Omega_\theta := \{(i, \theta_i') \mid i \in N, \, \kappa(\theta_i', \theta_{-i}) \neq \emptyset\} \neq \varnothing$. In other words, if players were to report their types, and the reported profile was in D, all players would know that some player must have lied. Further, the deviating player and the true state of nature must be such that $\kappa(\theta_i', \theta_{-i}) \neq \varnothing$. The set Ω_θ is the set of pairs (player, type) that could have caused the problematic announcement θ.

For each $\theta \in D$, consider the condition

$$\exists \mu \in \triangle A, \forall (i, \theta_i') \in \Omega_\theta, \forall k \in \kappa(\theta_i', \theta_{-i}) : \; v_i^k \geq u_i(k, \mu). \qquad (JR(\theta))$$

These inequalities are called Joint Rationality (JR), since they provide a link across payoffs of different players.[2]

Lemma 3.3 *Every $v \in B_\delta$ satisfies all constraints $(JR(\theta))_{\theta \in D}$.*

Proof: Let $v \in B_\delta$ be an equilibrium payoff vector and σ be the corresponding equilibrium. Let $\theta = (\theta_i)_i \in D$ and consider for each $(i, \theta_i') \in \Omega_\theta$ the deviation τ^i of player i such that, if his type is θ_i', player i plays as if he were of type θ_i, i.e. $\tau_{i,\theta_i'} = \sigma_{i,\theta_i}$, and which coincides with σ_i for all other types. Take two elements (i, θ_i') and (j, θ_j') in Ω_θ. The distribution over plays under $(\tau_{i,\theta_i'}, \sigma_{-i,\theta_{-i}})$ and $(\tau_{j,\theta_j'}, \sigma_{-j,\theta_{-j}})$ are the same, i.e. this is the distribution under $\sigma_\theta = (\sigma_{l,\theta_l})_{l \in N}$. In words, there is no way to distinguish the situation in which player i consistently mimicks type θ_i and the one in which player j consistently mimicks type θ_j. Let $\mu \in \triangle A$ denote the occupation measure generated by σ_θ. If $JR(\theta)$ is violated, there exists a player i and a state $k \in \kappa(I_i(k), \theta_{-i})$ such that player i's equilibrium payoff in state k, v_i^k, is strictly lower than his payoff if he were to follow σ_{θ_i}, a contradiction. □

To conclude this section, we note that the conditions $JR(\theta)$ are closely related to the conditions $IR(i, \theta)$. Indeed, using the minmax theorem, we may write those inequalities in the following alternative and compact way

$$\forall q \in \triangle\{(i,k) : k \in \kappa(\theta_{-i})\} \; : \; \sum_{i,k} q(i,k) v_i^k \geq \min_{a \in A} \sum_{i,k} q(i,k)\, u_i(k,a).$$

[2]Joint Rationality has been first introduced in Renault (2001) in a three-player setup.

For the sake of brevity, we often omit arguments and refer to each type of condition simply as IC, IR, or JR.

4 Sufficient Conditions

Let $V^* \subset \mathbb{R}^{KN}$ denote the set of feasible payoff vectors that satisfy IC, IR, and JR. The set V^* may be empty. Nevertheless, we show that this set characterizes the set of belief-free equilibria, up to its boundary.

Let $\hat{K} := \left\{k : \bigcap_{i \in N} I_i(k) \neq \{k\}\right\}$ be the set of states that cannot be distinguished by the join of the players' information partitions. Let $\hat{u}$ be the matrix $(u_i^k(a))$ with $N \times |\hat{K}|$ rows and $|A|$ where k belongs to $\hat{K}$. The reward u is *generic* if the matrix $\hat{u}$ has rank $N \times |\hat{K}|$. Indeed, viewing any such matrix as an element of $\mathbb{R}^{N|\hat{K}||A|}$, this condition is generically satisfied whenever $|A| \geq N|\hat{K}|$. The main result of this paper is the following.

Theorem 4.1 *If $v \in \text{int}V^*$ and u is generic, there exists $\bar{\delta} < 1$, $\forall \delta > \bar{\delta}$, $v \in B_\delta$.*

The interiority assumption is rather standard in the literature on repeated games with discounting, and has been first introduced by Fudenberg and Maskin (1986). In the next subsection, we provide a proof under the additional assumptions that there exists a public randomization device in every period (an independent draw from the uniform distribution on the unit interval), and that players can send costless messages, or *reports*, at the end of every period, as well as before the first period of the game. The proof in the appendix dispenses with these assumptions. (The proof of the dispensability of the public randomization follows ideas of Fudenberg and Maskin (1991) and Sorin (1986) and is only sketched.)

If $\mathcal{I}$ and $\mathcal{I}'$ are two different information structures for the same game, and V^*, V'^* are the corresponding sets of feasible, incentive compatible, individually and jointly rational payoff sets, observe that $V^* \subseteq V'^*$ if $\mathcal{I}'_i$ is finer than $\mathcal{I}_i$ for all $i \in N$. That is, the limit set of belief-free equilibrium payoffs is monotonic with respect to the information structure.

Simplified Proof

Player i's message set is Θ_i. The timing in a given period is as follows.

1. A draw from the uniform distribution on $[0, 1]$ is publicly observed;

2. Actions are simultaneously chosen;

3. Messages are simultaneously chosen.

As far as messages go, players always report their types truthfully in equilibrium. We refer to the event in which one player does not report truthfully as *misreporting* by this player. A type profile is *inconsistent* if $\kappa(\theta) = \varnothing$, and it is *consistent* otherwise.

As far as actions go, equilibrium play can be divided into three phases: *regular* phases, *penitence* phases and *punishment* phases. Regular and penitence phases last one period. Punishment phases last T period, for some $T \in \mathbb{N}$ to be defined.

In regular and penitence phases, players use an action profile that is coordinated by the public randomization device. In a punishment phase, a player is minmaxed by his opponents, in the sense of Blackwell described above.

To ensure that the strategy profile is belief-free, we must make sure that the punished player is playing the same way independently of the state, and that the punishing players have incentives to carry out the minmax strategy, even when this strategy calls for mixed actions. This complicates somewhat the description of the equilibrium strategies.

There are two kinds of deviations. The punishment phase is triggered if a player deviates in his choice of an action ('deviation in action'), and deters him from making such deviations. The penitence phase is triggered only if an inconsistent type profile is observed, and deters players from misreporting ('deviation in report') to induce an inconsistent type profile. Incentive compatibility of payoffs deters players from misreporting to induce a false but consistent type profile.

The equilibrium path consists of an infinite repetition of the regular phases.

Regular phases are denoted $R^{\theta}(\varepsilon)$, with $\kappa(\theta) \neq \varnothing$ and $\varepsilon \in \mathbb{R}^{N|\kappa(\theta)|}$. Penitence phases are denoted $E^{\theta}(\varepsilon)$, where $\kappa(\theta) = \varnothing$ and $\varepsilon \in \mathbb{R}^{NK}$. Punishment phases are denoted $P^{\theta_{-i}}$, with $\kappa(\theta_{-i}) \neq \varnothing$.

Actions and Messages

(i) *Regular phase*: In a regular phase, actions are determined by the outcome of the public randomization device. In phase $R^{\theta}(\varepsilon)$, action profiles are selected according to a probability distribution $\mu_{\theta}(\varepsilon)$ in such a way that

$$u_i(k, \mu_{\theta}(\varepsilon)) = v_i^k + \varepsilon_i$$

for $k \in \kappa(\theta_i, \theta_{-i})$, and

$$u_i(k, \mu_{\theta_i,\theta_{-i}}(\varepsilon)) > u_i(k, \mu_{\theta_i',\theta_{-i}}(\varepsilon')) \tag{2}$$

for all i, all $\varepsilon_i \in [-\overline{\varepsilon}, \overline{\varepsilon}]$, all $\varepsilon_i' \in [-\overline{\varepsilon}, \overline{\varepsilon}]$, all (θ_i, θ_{-i}) and (θ_i', θ_{-i}) such that $\kappa(\theta_i, \theta_{-i}) \neq \varnothing$ and $\kappa(\theta_i', \theta_{-i}) \neq \varnothing$. Such a distribution exists for sufficiently small $\overline{\varepsilon} > 0$ given that $v \in \operatorname{int} V^*$ is strictly incentive compatible.

At the end of a regular phase, all players truthfully report their types.

(ii) *Penitence phase*: In a penitence phase, actions are determined by the outcome of the public randomization device. Consider penitence phase $E^{\theta}(\varepsilon)$. Recall that $\kappa(\theta) = \varnothing$. We distinguish two cases.

1. $\theta \in D$: by definition, there exist a set Ω_θ of players and types $(i, \theta_i^{'})$ such that $\kappa\left(\theta_i^{'}, \theta_{-i}\right) \neq \varnothing$. Action profiles are selected according to a probability distribution $\mu_\theta(\varepsilon)$ in such a way that
$$u_i(k, \mu_\theta(\varepsilon)) < v_i^k + \varepsilon_i \tag{3}$$
for all $(i, \theta_i^{'}) \in \Omega_\theta$, $k \in \kappa\left(\theta_i^{'}, \theta_{-i}\right)$ and all $\varepsilon_i \in [-\overline{\varepsilon}, \overline{\varepsilon}]$. Such a distribution exists for sufficiently small $\overline{\varepsilon} > 0$ given that $v \in \operatorname{int} V^*$ satisfies (JR) with strict inequality.

2. $\theta \notin D$ (i.e., at least two players misreported): Players use some fixed, but arbitrary action profile $\underline{a} := \{\underline{a}_i\}_{i=1}^N \in A$.

At the end of a penitence phase, all players truthfully report their types.

(iii) *Punishment phase*: A punishment phase lasts T periods. In $P^{\theta_{-i}}$, players $-i$ use $\widehat{s}_{-i}^{\theta_{-i}}$. For some action $\underline{a}_i \in A_i$, let $s_i^{\underline{a}_i}$ denote the strategy of playing $\underline{a}_i$ after all histories within the punishment phase.[3] Player i plays $s_i^{\underline{a}_i}$ throughout the phase.

We pick $T \in \mathbb{N}, \overline{\delta} < 1$ and $\overline{\varepsilon} > 0$ such that, for all $\delta > \overline{\delta}$ and all $k \in \kappa(\theta_{-i})$, player i's average discounted payoff over the T periods is no larger than $v_i^k - 2\overline{\varepsilon}$. This is possible since v satisfies (IR) with strict inequality.

At the end of each period of a punishment phase, all players truthfully report their types.

[3] To avoid introducing additional notation, we have used here the same notation (i.e., $\underline{a}_i$) than in one of the specifications for the penitence phase. It is irrelevant whether these are the same actions or not.

Initial phase

All players truthfully report their types at the beginning of the game. Given report profile θ, the initial phase is $R^{\theta}(0)$.

Transitions

(i) *From a regular phase* $R^{\theta}(\varepsilon)$*:* Let a denote the (pure) action profile determined by the public randomization device, a' the realized action profile, and θ' the report of types at the end of the phase.

1. (Unilateral deviation) $a'_i \neq a_i$ for some $i \in N$ and $a'_{-i} = a_{-i}$:
   (a) $\kappa(\theta'_{-i}) \neq \varnothing$: the next phase is $P^{\theta'_{-i}}$;
   (b) $\kappa(\theta'_{-i}) = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_j = -\overline{\varepsilon}$ if $(j, \theta''_j) \in \Omega_{\theta'}$ for some $\theta''_j \in \Theta_j$, and $\varepsilon'_j = \varepsilon_j$ otherwise.

2. (Multilateral deviations, or no deviation) $a'_i \neq a_i$ for some $i \in N$ and $a'_{-i} \neq a_{-i}$, or $a' = a$:
   (a) $\kappa(\theta') \neq \varnothing$:
      i. $\theta' = (\theta_{-i}, \theta'_i)$ for some $i \in N$ and $\theta'_i \neq \theta_i$: the next phase is $R^{\theta'}(-\overline{\varepsilon}, \varepsilon_{-i})$;
      ii. otherwise, the next phase is $R^{\theta}(\varepsilon)$;
   (b) $\kappa(\theta') = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_i = -\overline{\varepsilon}$ if $(i, \theta''_i) \in \Omega_{\theta'}$ for some $\theta''_i \in \Theta_i$, and $\varepsilon'_i = \varepsilon_i$ otherwise.

(ii) *From a penitence phase* $E^{\theta}(\varepsilon)$: Let a denote the (pure) action profile determined by the public randomization device, a' the realized action profile, and θ' the report of types at the end of the phase.

1. (Unilateral deviations) $a'_i \neq a_i$ for some $i \in N$ and $a'_{-i} = a_{-i}$:

   (a) $\kappa(\theta'_{-i}) \neq \varnothing$: the next phase is $P^{\theta'_{-i}}$;

   (b) $\kappa(\theta'_{-i}) = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_j = -\overline{\varepsilon}$ if $(j, \theta''_j) \in \Omega_{\theta'}$ for some $\theta''_j \in \Theta_j$, and $\varepsilon'_j = \varepsilon_j$ otherwise.

2. (Multilateral deviations, or no deviation) $a'_i \neq a_i$ for some $i \in N$ and $a'_{-i} \neq a_{-i}$, or $a' = a$:

   (a) $\kappa(\theta') \neq \varnothing$: the next phase is $R^{\theta}(\varepsilon)$;

   (b) $\kappa(\theta') = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_i = -\overline{\varepsilon}$ if $(i, \theta''_i) \in \Omega_{\theta'}$ for some $\theta''_i \in \Theta_i$, and $\varepsilon'_i = \varepsilon_i$ otherwise.

(iii) *From a punishment phase* $P^{\theta_{-i}}$: The punishment phase lasts T periods. Let h^T denote an arbitrary history of length T. Let θ' denote the reported type profile in the T-th period. Then

1. (a) $\kappa(\theta') = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_i = -\overline{\varepsilon}$ if $(i, \theta''_i) \in \Omega_{\theta'}$ for some $\theta''_i \in \Theta_i$, and $\varepsilon'_i = \varepsilon_i$ otherwise;

   (b) $\kappa(\theta') \neq \varnothing$: the next phase is $R^{\theta'}(\varepsilon_i(h; P^{\theta_{-i}}), \varepsilon_{-i}(h; P^{\theta_{-i}}))$, with $\varepsilon_j(h; P^{\theta_{-i}}) \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, all j. The values $\varepsilon_j(h; P^{\theta_{-i}})$ are such that:

(4) for all $k \in \kappa(\theta')$, and conditional on any history $h \in H^T$, playing $s_i^{a_i}$ in the punishment phase is an optimal continuation strategy for player i, given $\widehat{s}_{-i}^{\theta_{-i}}$; further, if $\theta'_{-i} = \theta_{-i}$, player i's expected payoff, evaluated at the beginning of the punishment phase, from playing $s_i^{a_i}$ given $\widehat{s}_{-i}^{\theta_{-i}}$ (and given that θ' is truthfully reported), is equal to $\left(1-\delta^T\right)\left(v_j^k - 2\overline{\varepsilon}\right) + \delta^T\left(v_j^k - \overline{\varepsilon}\right)$, for all $k \in \kappa(\theta')$. That this is possible follows from inequality (6) below.

(5) for all $k \in \kappa(\theta')$, and conditional on any history $h \in H^T$, playing $\widehat{s}_j^{\theta_{-i}}$ is an optimal continuation strategy for player $j \neq i$, given $(s_i^{a_i}, (\widehat{s}_{j'}^{\theta_{-i}})_{j' \neq j})$; In addition $\varepsilon_j\left(\cdot; P^{\theta_{-i}}\right)$ is in $[\overline{\varepsilon}/3, \overline{\varepsilon}]$ if $\theta'_j = \theta_j$, and it is in $[-\overline{\varepsilon}, -\overline{\varepsilon}/3]$ otherwise (recall that h specifies θ'). That this is possible follows from inequality (6) below.

It is clear that these strategies do not depend on players' beliefs, but only on past history.

Optimality Verification

Given $v \in \operatorname{int} V^*$, we now pick $\overline{\varepsilon} > 0$ small to ensure that the probability distributions introduced above exist, and $\overline{\delta}$, and T such that the payoff of a punished player is low enough, as specified above for the punishment phase (see 'Actions and Messages'). In addition, we take these values to satisfy

$$-\left(1-\delta^T\right) M + \delta^T\left(v_j^k + \overline{\varepsilon}/3\right) > \left(1-\delta^T\right) M + \delta^T\left(v_j^k - \overline{\varepsilon}/3\right), \quad (6)$$

$$-\left(1-\delta\right) M + \delta\left(v_j^k - \overline{\varepsilon}\right) > \left(1-\delta\right) M + \delta\left(\left(1-\delta^T\right)\left(v_j^k - 2\overline{\varepsilon}\right) + \delta^T\left(v_j^k - \overline{\varepsilon}\right)\right). \quad (7)$$

Given v and $\overline{\varepsilon} > 0$, these are all satisfied as $\delta^T \to 1$ and $T \to \infty$, so they are also satisfied for values of T and δ that are large enough. Inequality (6) guarantees that a variation of $2\overline{\varepsilon}/3$ in continuation payoffs at the end of a punishment phase dominates any gains/losses that could be incurred during such a phase. Inequality (7) guarantees that the punishment phase is long enough to deter deviations in action.

Regular Phase: $R^\theta(\varepsilon)$ *and penitence phases* $E^\theta(\varepsilon)$: Let a denote the (pure) action profile determined by the public randomization device, a' the realized action profile, and θ' the report of types at the end of the phase.

Actions: Suppose that $a' = (a_{-i}, a'_i)$ for some i and $a'_i \neq a_i$, i.e., player i unilaterally deviates from the prescribed action profile. Then, provided players $-i$ truthfully report, the punishment phase $P^{\theta'_{-i}}$ starts. The maximum that player i can obtain by deviating is the right-hand side of (7), while by conforming to the prescribed action he gets at least as much as the left-hand side of (7).

Messages: let θ_i be player i's type. We distinguish two cases.

1. Either no or more than one player deviated in action:

   If player i reports truthfully, he gets at least $v_i^k - \overline{\varepsilon}$, where $k \in \kappa(\theta')$. If he misreports, we further distinguish two cases:

   (a) $\kappa(\theta') = \varnothing$: assuming the other players report truthfully, the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\overline{\varepsilon}$. So player i's payoff is at most $\max_{\theta'_i \neq \theta_i}(1-\delta)\, u_i(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)) + \delta\left(v_i^k - \overline{\varepsilon}\right)$, which is less than $v_i^k - \overline{\varepsilon}$, because of (3).

   (b) $\kappa(\theta') \neq \varnothing$: Player i gets at most $\max_{\theta'_i \neq \theta_i}(1-\delta)\, u_i\left(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)\right) + \delta\left(v_i^k - \overline{\varepsilon}\right)$, which is less than $\left(v_i^k - \overline{\varepsilon}\right)$, because of (2).

2. $a' = (a_{-j}, a'_j)$ for some j and $a'_j \neq a_j$ (i.e., player j deviated in action):

   Player j's report is irrelevant and he can as well report truthfully.

   If player $i \neq j$ reports truthfully his type, he gets at least $-\left(1-\delta^T\right) M + \delta^T \left(v_i^k + \overline{\varepsilon}/3\right)$. If he misreports, there are two cases:

   (a) $\kappa(\theta') = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\overline{\varepsilon}$, so his payoff is smaller than $(1-\delta) M + \delta\left(v_i^k - \overline{\varepsilon}\right) < \left(1-\delta^T\right) M + \delta^T \left(v_i^k - \overline{\varepsilon}/3\right)$, which is less than $-\left(1-\delta^T\right) M + \delta^T \left(v_j^k + \overline{\varepsilon}/3\right)$ because of (6).

   (b) $\kappa(\theta'_i, \theta_{-i}) \neq \varnothing$: Player i gets at most $\left(1-\delta^T\right) M + \delta^T \left(v_i^k - \overline{\varepsilon}/3\right)$ (assuming he reports truthfully at the end), which is less than $-\left(1-\delta^T\right) M + \delta^T \left(v_i^k + \overline{\varepsilon}/3\right)$ because of (6).

Punishment phase $P^{\theta_{-i}}$: Let θ' denote the reported type profile in the T-th period. *Actions*: We consider first player i, then Player $j \neq i$.

1. Player i: as mentioned, inequality (6) guarantees that we can specify $\varepsilon_i\left(h; P^{\theta_{-i}}\right)$ such that $s_i^{a_i}$ is optimal after every history in the punishment phase, given $\widehat{s}_{j\neq i}^{\theta_{-i}}$.

2. Player $j \neq i$: similarly, inequality (6) guarantees that we can specify $\varepsilon_j\left(h; P^{\theta_{-i}}\right)$ such that $\widehat{s}_j^{\theta_{-i}}$ is optimal after every history in the punishment phase, given $\widehat{s}_{j'\neq i,j}^{\theta_{-i}}$.

Messages: The only payoff relevant message is the one at the end of the punishment phase. Let θ' denote the reported type profile in the T-th period. If player $i \in N$ reports truthfully his type, he gets at least $v_i^k - \overline{\varepsilon}$. If he misreports, we distinguish two cases:

1. (a) $\kappa(\theta') = \varnothing$: the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\overline{\varepsilon}$, so player i's payoff is at most $\max_{\theta'_i \neq \theta_i} (1-\delta) u_i(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)) + \delta \left(v_i^k - \overline{\varepsilon}\right)$, which is less than $v_i^k - \overline{\varepsilon}$ because of (3).

   (b) $\kappa(\theta') \neq \varnothing$: player i gets at most $\max_{\theta'_i \neq \theta_i} (1-\delta) u_i(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)) + \delta \left(v_i^k - \overline{\varepsilon}\right)$, which is less than $v_i^k - \overline{\varepsilon}$ because of (2).

5 Existence

Our main theorem states that, given $V^* \neq \varnothing$, all points in the interior of V^* are BFE payoffs if δ is large enough. However, achieving incentive compatibility together with individual rationality and joint rationality might not be possible, as is already known from the two-player case, and some conditions are required. In this section, we give conditions for existence of belief-free equilibria. More precisely, we consider different classes of games each one characterized by some properties of the reward functions and/or of the information structure. For each one of these classes we prove that V^* is not empty by identifying payoffs vectors that are IC, IR and JR. More precisely, given the set of players N, the set of states K and the set of actions profiles A, let $\mathcal{U} = (\mathbb{R}^{K \times A})^N$ be the set of all payoff functions and $\mathcal{Y}$ be the set of information structures. For an information structure $\mathcal{I}$ and a reward function u, we denote by $V^*(\mathcal{I}, u)$ the set of payoff vectors that satisfy IC, IR and JR.

5.1 Known punishments

Our first objective is to find conditions on the payoff function u such that V^* is nonempty independently of the information structure. Note first that $V^*(\mathcal{I}, u)$ is nonempty for all information structure $\mathcal{I} \in \mathcal{Y}$ if and only if $V^*(\mathcal{I}, u)$ is nonempty for the coarser information structure $\mathcal{I}$ i.e. for $I_i(k) = K$ for all $i \in N$ and all $k \in K$. Necessity is trivial. Sufficiency follows from the fact that, for any pair of comparable information structures $\mathcal{I}$ and $\mathcal{I}'$, with $\mathcal{I}'$ finer than $\mathcal{I}$ (i.e., $\mathcal{I}_i$ finer than $\mathcal{I}'_i$ for all i), if $V^*(\mathcal{I}, u)$ is nonempty, then $V^*(\mathcal{I}', u)$ is so. A belief-free equilibrium for a coarse information structure remains a belief-free equilibrium for any finer information structure. In order to provide necessary and sufficient condition on $\mathcal{U}$ such that $V^*(\mathcal{I}, u)$ is nonempty independently of the information structure $\mathcal{I}$, let

$$\varphi_i(q) := \min_{\alpha_{-i} \in \prod_{j \neq i} \triangle A_j} \max_{a_i \in A_i} \sum\nolimits_{k \in K} q(k) u_i(k, \alpha_{-i}, a_i).$$

Proposition 5.1 *The set $V^*(\mathcal{I}, u)$ is nonempty for all information structure $\mathcal{I}$ if and only if there exists a distribution over action profile $\mu^* \in \triangle A$, such that for each $i \in N$,*

$$\forall q \in \triangle K: \ \sum\nolimits_{k \in K} q(k) u_i(k, \mu^*) \geq \varphi_i(q).$$

Proof. It is sufficient to show that when $\mathcal{I}$ satisfies $I_i(k) = K$ for all $i \in N$ and all $k \in K$ then the conditions of the proposition are necessary and sufficient to have $V^*(\mathcal{I}, u) \neq \varnothing$. Sufficiency: Consider the payoff vector v^* obtained by implementing the distribution μ^* independently of the state. This payoff is clearly IC and JR since it is achieved using a strategy that is independent of the state. This payoff vector satisfies IR since the condition on μ^* states that no player i in no state k can guarantee more than v_i^{k*} when the other players use the Blackwell punishment strategy corresponding to a situation in which player i knows the state and the other players do not. Necessity: note first that the equilibrium play must be independent of the state because of feasibility condition 2.

Second, suppose that there exists no μ^* satisfying the condition of the proposition. In other words for each $\mu \in \triangle A$ there exists a player i and $q^\mu \in \triangle K$ such that

$$\sum\nolimits_{k\in K} q^\mu(k)u_i(k,\mu) < \varphi_i(q^\mu).$$

This implies that for any candidate equilibrium payoff achieved with some distribution over action profiles μ that is independent of the state, there exists a player i that finds it profitable to deviate in some state. □

The condition of proposition 5.1 is trivially satisfied when it is possible to find a pooling equilibrium distribution μ^* and a state independent punishment strategy. This is the case, for instance, in most auction formats and oligopoly games.

When focusing on finer information structures where players have types, punishment strategies sustaining a pooling equilibrium can be type dependent as illustrated by Proposition 5.2. Let $\hat{D}$ be the set of type profiles which are either compatible with some state, or which are compatible with some state after deletion of one type. That is,

$$\hat{D} = \{\theta \in \prod\nolimits_{i\in N} \Theta_i : \exists i \in N, \kappa(\theta_{-i}) \neq \varnothing\}.$$

The following condition guarantees that V^* is non-empty.

Proposition 5.2 *If there exists a distribution over action profile $\mu^* \in \triangle A$, such that for each $\theta \in \hat{D}$, there exists $\mu^\theta \in \triangle A$, such that for each i in N and each $k \in \kappa(\theta_{-i})$,*

$$\max_{a_i\in A_i} u_i(k, a_i, \mu^\theta_{-i}) \leq u_i(k, \mu^*),$$

then V^* is non-empty.

Proof. It is sufficient to show that that $v := (u_i(k, \mu^*))_{i \in N, k \in K}$ is in V^*. IC: The payoff vector v can be achieved by implementing the occupation measure μ^* irrespective of the announcements, hence it is incentive compatibility. IR and JR: the condition on μ^θ implies that when distribution over action profile μ^θ is implemented, then in all possible states a player cannot gain more than v even if he unilaterally deviates or misreport generating an inconsistent report profile. Thus, μ^θ can be used to deter unilateral deviation or misreport guaranteeing that v is individually and jointly rational. □

Proposition 5.2 is the equivalent for the N-player case of condition 4 in Hörner and Lovo (2008) Still, this condition might not be met in general.

In the remainder of this section, we focus on information structure such that for each k, $\cap_{i \in N} I_i(k) = \{k\}$. In this instance, $\hat{k} = \varnothing$ and the reward function u trivially satisfy the genericity condition of Theorem 4.1. [4] Given a set of payoff functions $\mathcal{S} \subseteq \mathcal{U}$, our aim is to provide necessary and sufficient conditions on the information structure $\mathcal{I}$ in order to have: $V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{S}$.

5.2 No restriction on payoffs: $\mathcal{S} = \mathcal{U}$

The result follows from Renault and Tomala (2004) who study the existence of completely revealing equilibria in infinitely repeated undiscounted games. They show that a fully revealing equilibrium exists for any payoff function if and only if a simple majority rule allows players to learn the true state during the announcement phase, even under some unilateral misannouncement. To formalize this idea, we give some definitions.

For each player i and state k, let $I_{-i}(k) = \cap_{l \neq i} I_l(k)$ denote the meet of the information

[4] This is without loss of generality when players have known-own payoff, but not always innocuous otherwise. For example, if each player's payoff function depends only on his own action and on the state, and the optimal action is not the same in two states that no player distinguishes, then BFE do not exist.

partitions of the other players at k. For each pair of players i, j, let $I_{-ij}(k) = \cap_{l \neq i, l \neq j} I_l(k)$ denote the meet of the information partitions of players other than i and j.

Definition 5.3 *1. Player i is essential at k if $I_{-i}(k) \neq \{k\}$.*

2. The pair of players i, j is essential at k if $I_{-ij}(k) \neq \{k\}$.

If player i (or the pair i, j) is essential at k, then his information is needed in order to learn the state k.

Definition 5.4 *1. The information structure $\mathcal{I}$ has no essential player (NEP), if for each state k, no player is essential at k.*

2. The information structure $\mathcal{I}$ has the majority property (MAJ) if it has no essential player and for each state k, no pair of players is essential at k.

In two-player games, NEP and MAJ both require that each player knows the state: $I_i(k) = \{k\}$ for each $i = 1, 2$ and each k. With at least three players, MAJ requires that for every pair of states k, k', there exists at least three players i for which $I_i(k) \neq I_i(k')$. Thus, MAJ clearly implies NEP. The converse is not true, as shown by the following example.

Example 5.5 *There are 4 players and 4 states, $K = \{k_1, k_2, k_3, k_4\}$. Let I_i be the information partition of player i such that player i distinguishes state k_i from the other states, but cannot tell these other states apart: $I_i(k_i) = \{k_i\}$, $I_i(k) = K \backslash \{k_i\}$ for $k \neq k_i$. This is represented by the following matrix, where the i-th line gives the information of player i ($*$ stands for $K \backslash \{k_i\}$).*

	k_1	k_2	k_3	k_4
1	k_1	$*$	$*$	$*$
2	$*$	k_2	$*$	$*$
3	$*$	$*$	k_3	$*$
4	$*$	$*$	$*$	k_4

$\mathcal{I}$

Note that this information structure $\mathcal{I}$ satisfies NEP but not MAJ. On one hand, no player is essential, but on the other hand, without the information of players 3 and 4, it is impossible to tell apart k_3 from k_4. The following information structure $\mathcal{J}$ satisfies MAJ.

	k_1	k_2	k_3	k_4
1	k_1	k_2	$*$	$*$
2	$*$	k_2	k_3	$*$
3	$*$	$*$	k_3	k_4
4	k_1	$*$	$*$	k_4

$\mathcal{J}$

For instance, there are three players (1,2 and 4) who distinguish k_1 from k_2.

We have the following result:

Proposition 5.6 *$V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{U}$, if and only if $\mathcal{I}$ has the majority property.*

Proof. The proof is straightforward and follows the arguments of Renault and Tomala (2004). To see that the majority property is a sufficient condition, assume that the information structure has the majority property and for each state k fix an payoff vector v^k that is individually rational in the full complete information game corresponding to state k. If the profile of announcements is such that, deleting the announcements of some pair of players i, j, a state k is identified by the announcements, then the action profile implementing v^k is played (if there are several such pair of players, pick one at random). Otherwise, some fixed action profile is played. This clearly defines a point in $V^*(\mathcal{I}, u)$: it satisfies individual rationality by construction, and no player has an incentive to unilaterally misreport since the true state is revealed anyway.

The necessity part is as follows. Suppose that MAJ is not satisfied, but that NEP is. This means that there exists two players i, j, say players 1 and 2, and two states k, k' such that $I_1(k) \neq I_1(k')$, $I_2(k) \neq I_2(k')$ and for each player $l \neq 1, 2$, $I_l(k) = I_l(k')$. We are thus in a situation where only two players are informed. The following example is due to Renault (2001) and shows that individual rationality and joint rationality might be mutually exclusive.

Example 5.7 *Each player has two actions. Player 1 chooses the row, player 2 chooses the column and player 3 chooses the matrix. There are two states k, ℓ. Players 1 and 2 are informed of the state, player 3 is not. The payoff matrix in state k is the following:*

	L	R
T	$1,1,0$	$1,1,0$
B	$1,1,0$	$1,1,0$

W

	L	R
T	$0,0,1$	$0,0,1$
B	$0,0,1$	$0,0,1$

E

The payoff matrix in state k' is:

	L	R
T	$0,0,1$	$0,0,1$
B	$0,0,1$	$0,0,1$

W

	L	R
T	$1,1,0$	$1,1,0$
B	$1,1,0$	$1,1,0$

E

Suppose that there exists a payoff vector in $V^*(\mathcal{I}, u)$. If players 1 and 2 both announce k, individual rationality implies that player 3 plays E. The payoff vector in state k is thus $(0, 0, 1)$. Similarly, if players 1 and 2 announce k', player 3 plays W and the payoff vector in state k' is $(0, 0, 1)$.

Now, suppose that player 1 announces k and player 2 announces k': either the true state is k and player 2 is misreporting, or the true state is k' and player 1 is misreporting. Joint rationality implies that there exists an occupation measure μ such that $u_1(k', \mu) \leq 0$ and $u_2(k, \mu) \leq 0$. This is impossible, since for each action profile a, $u_1(k', a) + u_2(k, a) = 1$.

If neither MAJ nor NEP are satisfied, then there exists a player i, and two states k, k' such that $I_i(k) \neq I_i(k')$ and for each $j \neq i$, $I_j(k) = I_j(k')$. This is the case, for instance, in two-player games with one sided incomplete information for which it is easy determine the payoff function u such that $V^*(\mathcal{I}, u) = \varnothing$. As an illustration, see example 5.10 below. This concludes the proof of Proposition 5.6. $\Box$.

In the sequel, we show how, by restricting the possible payoff functions, the requirements on the information structure ensuring existence can be relaxed.

5.3 Bad outcome

Definition 5.8 *The payoff function has a Bad Outcome (BO) if there exists a possibly correlated distribution over action profiles that provides each player with no more than his minmax payoff in each state:*

$$\exists \mu^o \in \triangle A, \forall i \in N, \forall k \in K, \, u_i(k, \mu^o) \leq \underline{u}_i^k,$$

with $\underline{u}_i^k = \min_{\alpha_{-i} \in \prod_{j \neq i} \triangle A_j} \max_{a_i \in A_i} u_i(k, a_i, \alpha_{-i})$*. Let* $\mathcal{B}$ *be the set of payoff functions that have a bad outcome.*

Theorem 5.9 $V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{B}$, *if and only if* $\mathcal{I}$ *has no essential player.*

Proof. Sufficiency. For each state k, fix a vector v^k that is individually rational in the full information game corresponding to state k, i.e., $v^k \geq \underline{u}^k$. We show that $v := \left\{v^k\right\}$ is in V^*. IC and JR: when there is no essential player, the information held by players other than i is sufficient to reveal the state. Thus, Player i has no other choice than letting the state be revealed or being inconsistent with the other players. The distribution corresponding to the bad outcome can be used to deter a player from being inconsistent. IR: The distribution corresponding to the bad outcome can be used to deter unilateral deviations. Thus, no player has an incentive to deviate nor of being inconsistent with the play of other players as this would bring about the distribution over actions profile corresponding to the bad outcome.

Necessity. Consider the following game that has a bad outcome and where player 1 is essential to learn the state. For this game, $V^*(\mathcal{I}, u) = \varnothing$.

Example 5.10 *(This example is adapted from Hörner and Lovo, 2008). There are two states* k, k'*, and two players. Players 1 is informed of the state, player 2 is not. The payoff matrix in state* k *is the following:*

	L	M	R
T	$10, -4$	$1, 1$	$10, -4$
B	$1, 1$	$0, 0$	$-1, -4$

The payoff matrix in state k' *is:*

	L	M	R
T	$0, 0$	$1, 1$	$10, -4$
B	$1, 1$	$10, -4$	$-1, -4$

Action profile $\{B, R\}$ is the bad outcome. Player 1 can guarantee a payoff of at least 3 *in one of the states by randomizing equally between U and D and player 2 can guarantee at least* 0 *in each state. This implies that the equilibrium distribution over action profiles cannot assign probability more than* $1/5$ *to action profiles yielding* -4 *to player 2. In turn, this implies that player 1's payoff is at most* $14/5$ *in each state, a contradiction.*

5.4 Known-own payoffs

In two-player games, existence obtains whenever players know their payoffs and information is one-sided; i.e., whenever player 1 has more information than player 2 (Shalev, 1994). These conditions are also necessary in two-player games: Hörner and Lovo (2008) and Koren (1992) provide examples where existence fails if information is two-sided. In this section, we show how these results extend to N-player games.

Definition 5.11 *The game has* known-own payoffs *(KOP) if the payoff of each player i depends only on the action profile and on her type. That is, for each action profile a and each pair of states k, k':*

$$I_i(k) = I_i(k') \Longrightarrow u_i(k, a) = u_i(k', a).$$

Let $\mathcal{S}_{\mathcal{I}}$ be the set of KOP payoff functions when the information structure is $\mathcal{I}$.

Note that in games with known own payoff implies that $\bigcap_{i \in N} I_i(k) = \{k\}$. First we provide a sufficient condition on the information structure to get $V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{S}_{\mathcal{I}}$.

We say that *player i has more information than player j* if player i can deduce player j's type from her own type, i.e. if player i's information partition is finer than player j's partition: $I_i(k) \subseteq I_j(k)$ for each k.

Definition 5.12 1. *The information structure is* embedded *if for each pair of players (i, j), i has more information than j or j has more information than i.*

2. *The information structure is* weakly embedded *(WE) if there exists a pair of players i, j, say players 1 and 2, such that player* 1 *has more information than any other player and player* 2 *has more information than any player other than* 1*.*

We have the following:

Theorem 5.13 *In a game with known-own payoffs and weakly embedded information, $V^*(\mathcal{I}, u)$ is non-empty.*

Proof. First, we may assume without loss of generality, that the types of players $3, 4, \dots, n$ are commonly known. Indeed, let players announce their types. For each $i \geq 3$, players 1, 2 and i know the type of player i. Each player can thus find out the true type of player i from the announcement profile by applying the majority rule: even if there is a one report that differs from all others, at least two players' reports reveal the type of player i. It is thus sufficient to prove the following statement.

Proposition 5.14 *Consider a game with known-own payoffs and an information structure such that: player 1 knows the state and players $3, \dots, n$ have no information (i.e. $\forall k$, $I_1(k) = \{k\}$, $I_3(k) = \cdots = I_n(k) = K$). Then, $V^*(\mathcal{I}, u)$ is non-empty.*

Proof. Denote by $\underline{u}_i$ the minmax level of player $i = 3, \ldots, n$, by $\underline{u}_1^k$ the minmax level of player 1 in state k and by $\underline{u}_2^\theta$ the minmax level of player 2 of type θ. For each type θ of player 2, consider the set $\mathcal{A}_\theta$ of mixed actions profiles α such that:

- For each $i = 3, \ldots, N$,

$$\forall a_2 \in A_2,\ u_i(\alpha_1, a_2, \alpha_3, \ldots, \alpha_N) \geq \underline{u}_i;$$

- α_2 is a best-reply of player 2 of type θ to $(\alpha_1, \alpha_3, \ldots, \alpha_N)$.

The set $\mathcal{A}_\theta$ is clearly compact.

Claim 5.15 $\mathcal{A}_\theta$ *is non-empty.*

Proof. We fix α_1. For $i \geq 3$ consider the correspondence $F_i(\alpha_1, \cdot) : \prod_{j \notin \{1,2,i\}} \triangle A_j \to \triangle A_i$ defined by

$$F_i(\alpha_1, \alpha_{-1-2-i}) = \{\alpha_i : \forall a_2, u_i(\alpha_1, a_2, \alpha_i, \alpha_{-1-2-i}) \geq \underline{u}_i\}.$$

This correspondence is convex and compact valued. Let us prove that this is also non-empty valued. For a given α_{-1-2-i}, player i has a mixed action that yields a payoff no less than

$$\max_{\alpha_i} \min_{a_2} u_i(\alpha_1, a_2, \alpha_i, \alpha_{-1-2-i}) = \min_{\alpha_2} \max_{a_i} u_i(\alpha_1, \alpha_2, a_i, \alpha_{-1-2-i}),$$

where the equality follows from the minmax theorem. Now,

$$\min_{\alpha_2} \max_{a_i} u_i(\alpha_1, \alpha_2, a_i, \alpha_{-1-2-i}) \geq \underline{u}_i,$$

and $F_i(\alpha_1, \alpha_{-1-2-i})$ is non-empty.

Denote by $BR_{2,\theta}$ the best-reply correspondence of player 2 of type θ and by $BR_{1,f}$ the best-reply correspondence of player 1 when his payoff function is $f : A \to \mathbb{R}$. Consider the correspondence $\Phi_{f,\theta}$ from $\prod_i \triangle A_i$ to itself defined by

$$\Phi_{f,\theta}(\alpha) = \{\beta : \beta_1 \in BR_{1,f}(\alpha_{-1}),\ \beta_2 \in BR_{2,\theta}(\alpha_{-2}),\ \forall i \geq 3,\ \beta_i \in F_i(\alpha_1, \alpha_{-1-2-i})\}.$$

The correspondence $\Phi_{f,\theta}$ has non-empty, convex and compact values and it is straightforward to check that it has a closed graph. Thus, it admits a fixed point $\bar{\alpha}$ by Kakutani's fixed point theorem. Clearly, $\bar{\alpha}$ is in $\mathcal{A}_\theta$. Note that this profile has the additional property to lie on the best-reply graph of player 1. We thus have some degrees of freedom as we can choose any payoff function for player 1. This ends the proof of the Claim.

Let α^k be a mixed action profile that maximizes $u_1(k, \alpha)$ over $\alpha \in \mathcal{A}_{I_2(k)}$. We claim that the payoff vector

$$(u_1(k, \alpha^k), u_2(I_2(k), \alpha^k), u_3(\alpha^k), \dots, u_N(\alpha^k))$$

is in V^*. Under the assumptions of Proposition 5.14, the constraints defining V^* are the following:

- *Individual rationality for player 1.*

  For each θ and each $q \in \triangle\theta$,

$$\sum_{k\in\theta} q_k u_1(k, \alpha^k) \geq \min_{\alpha_{-1}} \max_{\alpha_1} \sum_{k\in\theta} q_k u_1(k, \alpha_1, \alpha_{-1}).$$

- *Individual rationality for players 2, 3, ..., N:*

  For each k, $u_2(I_2(k), \alpha^k) \geq \underline{u}_2^{I_2(k)}$, for each $i \geq 3$, $u_i(\alpha^k) \geq \underline{u}_i$.

- *Incentive compatibility for player 1:*

  For each k, k' such that $I_2(k) = I_2(k')$, $u_1(k, \alpha^k) \geq u_1(k, \alpha^{k'})$.

- *Joint rationality for players 1 and 2:*

  For each announcement (k', θ) such that $\theta \neq I_2(k')$, set $\alpha^{k',\theta} = (\alpha_{-2}^{k'}, \alpha_2^{k',\theta})$, where $\alpha_2^{k',\theta}$ is a best-reply of player 2 of type θ to $\alpha^{k'}$. The true state is either k' (and player 2 is misreporting) or $k \in \theta$ (in which case player 1 is misreporting). The following must hold:

  $u_1(k, \alpha^k) \geq u_1(k, \alpha^{k',\theta})$ for $k \in \theta$, and $u_2(I_2(k'), \alpha^{k'}) \geq u_2(I_2(k'), \alpha^{k',\theta})$.

Let us check all these points.

Individual rationality for player 1. Fix θ and $q \in \triangle\theta$. It follows by construction that

$$\sum\nolimits_{k\in\theta} q_k u_1(k, \alpha^k) = \sum\nolimits_{k\in\theta} q_k \max_{\alpha\in\mathcal{A}_\theta} u_1(k, \alpha) \geq \max_{\alpha\in\mathcal{A}_\theta} \sum\nolimits_{k\in\theta} q_k u_1(k, \alpha).$$

Let $\bar{\alpha}$ be a fixed point of $\Phi_{f,\theta}$ where f is chosen to be $\sum_{k\in\theta} q_k u_1(k, \cdot)$. We get that

$$\max_{\alpha\in\mathcal{A}_\theta} \sum\nolimits_{k\in\theta} q_k u_1(k, \alpha) \geq \sum\nolimits_{k\in\theta} q_k u_1(k, \bar{\alpha}) = \max_{\alpha_1} \sum\nolimits_{k\in\theta} q_k u_1(k, \alpha_1, \bar{\alpha}_{-1}),$$

where the last equality holds since $\bar{\alpha}$ is on the graph of $BR_{1,f}$. The right-hand-side is no less than $\min_{\alpha_{-1}} \max_{\alpha_1} \sum_{k\in\theta} q_k u_1(k, \alpha_1, \alpha_{-1})$.

Individual rationality for players 2, 3, ..., N. Individual rationality for players $3, \dots, N$ holds by construction of $\mathcal{A}_\theta$. Individual rationality for player 2 holds since she plays a best-reply to some mixed action profile.

Incentive compatibility for player 1. Suppose that the true state is k and let $\theta = I_2(k)$. Player 1 gets the payoff $\max_{\alpha \in \mathcal{A}_\theta} u_1(k, \alpha)$. If player 1 reports instead that the state is k' with $I_2(k') = \theta$, the induced action profile $\alpha^{k'}$ also belongs to $\mathcal{A}_\theta$, and the resulting payoff for player 1 is at most $\max_{\alpha \in \mathcal{A}_\theta} u_1(k, \alpha)$.

Joint rationality. As above, suppose that the true state is $k \in \theta$, but player 1 reports that the state is k' with $\theta' = I_2(k') \neq \theta$. Still, the action of player 2 is dictated by her type θ, so that the induced action profile belongs to $\mathcal{A}_\theta$. Thus, player 1 does not increase her payoff by this deviation.

Suppose now that the true state is k' but player 2 pretends that her type is θ. Players other than 2 play $\alpha^{k'}_{-2}$ and the best-reply of player 2 of type θ' is $\alpha^{k'}_{-2}$ by construction. Player 2 has thus no incentive to misreport. □

Now, we provide a necessary condition to get existence in all KOP games.

Definition 5.16 *The information structure has the 1EP property if the information structure is such that there exists at most one essential player at every state. That is, for every k, there exists i, such that, for every $j \neq i$, $I_{-j}(k) = \{k\}$.*

Note that NEP implies 1EP and WE implies 1EP. A typical instance in which 1EP holds, but not NEP, is when the state is the list of types for all players, player 1 knows the state, and all other players know their own type only. The next theorem states that

1EP is necessary for existence in games with known-own payoffs, and it is necessary and sufficient for existence in games with a bad outcome and known-own payoffs.

Theorem 5.17 *1. If $V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{S}_{\mathcal{I}}$, then $\mathcal{I}$ satisfies 1EP.*

2. $V^*(\mathcal{I}, u) \neq \varnothing$, $\forall u \in \mathcal{S}_{\mathcal{I}} \cap \mathcal{B}$ if and only if $\mathcal{I}$ satisfies 1EP.

Proof of 1. Assume for the sake of contradiction contradiction that $\mathcal{I}$ does not satisfy 1EP. This means that there exists a state k and two players i, j, say players 1 and 2, such that $I_{-1}(k) \neq \{k\}$ and $I_{-2}(k) \neq \{k\}$. This implies that there exists two states k' and k'' such that:

- $I_1(k) \neq I_1(k')$ and for each $l \neq 1$, $I_l(k) = I_l(k')$,
- $I_2(k) \neq I_2(k'')$ and for each $l \neq 2$, $I_l(k) = I_l(k'')$.

Thus, the situation is similar to two-sided incomplete information, where two players have private information. So we can find a payoff function with $V^*(\mathcal{I}, u) = \varnothing$ by building on the counter-example of Hörner and Lovo (2008), or Koren (1992), as follows.

Example 5.18 *Payoffs depend on the actions players 1 and 2 only and payoffs in state k, k', k'' are as follows. Player 1 chooses rows and player 2 chooses columns.*

	L	R
T	$3, 1$	$0, 0$
B	$0, 0$	$1, 3$

state k

	L	R
T	$3, 0$	$0, 1$
B	$0, 0$	$1, 1 + \varepsilon$

state k''

	L	R
T	$1+\varepsilon, 1$	$1, 0$
B	$0, 0$	$0, 3$

state k'

Here, $0 < \varepsilon < 1/36$. The proof that $V^*(\mathcal{I}, u) = \varnothing$ for this game is in Hörner and Lovo (2008). The argument is as follows. In state k', player 1 has a dominant strategy, and individual rationality requires T to be played with probability at least $1 - \varepsilon$ in that state. Now, in state k, player 1 may claim that the state is k'. Incentive compatibility requires thus (T, L) to be played with probability at least $3/4 - 3\varepsilon$ in state k. A symmetric argument for player 2 shows that (B, R) must be played with probability at least $3/4 - 3\varepsilon$ in state k. These two requirements are mutually inconsistent for small enough ε.

Proof of 2. Necessity follows from observing that (B, L) is a bad outcome in the example above. Consider now a game with known-own payoffs and a bad outcome, and an information structure which satisfies 1EP. We may partition the set of states as follows,

$$K = K_0 \cup K_1 \cup \cdots \cup K_S,$$

where for each $k \in K_0$, there is no essential player at k, and for each $s = 1, \ldots, S$, there exists a unique player i_s such that:

a) for all k, k' in K_s, $I_{i_s}(k) \neq I_{i_s}(k')$,

b) for all k, k' in K_s and all players $j \neq i_s$, $I_j(k) = I_j(k')$,

c) for all $k \in K_s$, $k' \notin K_s$, there exists $j \neq i_s$ such that $I_j(k) \neq I_j(k')$.

To construct a cell K_s of this partition, consider a state k such that some player i is essential at this state. This means that $I_{-i}(k) \neq \{k\}$. Set then $K_s = I_{-i}(k)$ and $i_s = i$. Property b) is clearly satisfied. Property a) holds since $I_{-i}(k) \cap I_i(k) = \{k\}$. Property c) holds since if $k' \notin K_s = I_{-i}(k)$, then there must exist $j \neq i_s$ such that $I_j(k) \neq I_j(k')$.

Choose, for each $k \in K_0$, an individually rational payoff v^k in state k. For each $s = 1, \ldots, S$, consider the game with incomplete information Γ_s where:

- It is common knowledge that the state belongs to K_s,
- Player i_s knows the state and other players have no information.

Let V_s^* be the set of IC, IR and JR payoffs of this game. The information structure of Γ_s is weakly embedded. Thus, from Theorem 5.13, V_s^* is non-empty. Pick a payoff vector in this set for each s. We construct the overall equilibrium as follows. Let players announce their information.

- If the announcements reveal a state $k \in K_0$, the action profile yielding v^k is played.
- If, given the announcements, the set K_s is common knowledge, the chosen equilibrium of Γ_s is played.
- If the announcements are inconsistent, the bad outcome is chosen.

The induced payoff vector is individually rational. We argue now that no player has an incentive to misreport. Player i who is not essential at state k has no choice but letting the state be revealed or being inconsistent with the other players' reports. The bad outcome ensures that he weakly prefers to tell the truth. Consider player i_s at some state $k \in K_s$. If he announces $I_{i_s}(k')$ for some $k' \in K_s$, the announcements are consistent. Each player knows now that the state may be any k in K_s and the equilibrium of Γ_s can be played. If player i_s announces $I_{i_s}(k')$ for some $k' \notin K_s$, property c) above implies that this announcement is inconsistent with some other player's announcement. Player i_s has thus no choice but letting K_s be revealed, or inducing the bad outcome. This ensures truth-telling.

6 Reputations

It follows from the previous section that V^* is non-empty when players know their own payoffs, and the incomplete information concerns one player's payoff only, so that the payoffs of all players but one are commonly known. Formally, for every player i, $u_i(k, \cdot) = u_i(\theta_i, \cdot)$, and for all $i \neq 1$, $|\Theta_i| = 1$. This environment with one-sided incomplete information is the focus of a large literature on 'reputations,' starting with Fudenberg and Levine (1989), and is assumed throughout this section. In Hörner and Lovo (2008), it was shown how results by Israeli (1999) for the set of undiscounted Nash equilibrium payoffs in two-player games with such information structures could be applied with hardly any change to the set of belief-free equilibrium payoffs as the discount factor tends to one. In this section, the generalization of those results to n players is presented. Proofs are straightforward generalizations of those by Israeli.

Fix one (payoff) type of player 1, the *rational* type. The purpose of this section is to identify how much the rational type is guaranteed to get in equilibrium, as the discount factor tends to one, as a function of his other possible payoff types. The rational type's payoff is denoted u_1, while his other possible payoff types are denoted u_1^k, $k = 2, \ldots, K$. We fix throughout the payoff functions $(u_2, \ldots, u_N)$ of players $i = 2, \ldots, N$. Given some payoff function u_1^k, u_i, let $\underline{u}_1^k$, $\underline{u}_i$ denote the corresponding minmax payoffs $\text{val}\, u_1^k$ and $\text{val}\, u_i$.

Given any vector $u^K := (u_1^2, \ldots, u_1^K)$ such that V^* is non-empty, let $v_1(u^K)$ be the infimum of the payoff of player 1's rational type over V^*. We define the *reputation payoff* of player 1's rational type as

$$u_1^* := \sup_{\{u^K : K \geq 2\}} v_1(u^K).$$

Observe that the rational type's equilibrium payoff must be at least equal to

$$\min_{\mu \in \triangle A} u_1(\mu) \text{ such that } u_1^k(\mu) \geq \underline{u}_1^k, u_i(\mu) \geq \underline{u}_i, \forall i, k \geq 2.$$

Indeed, if the state is k, the play specified by the equilibrium strategies must be an equilibrium of the game with complete information in state k, and therefore this play must be such that all players get at least their minmax payoff in that state. Since player 1's rational type can always follow the strategy of player 1's type k, he must receive at least as much as he would get from following this play. Therefore, it must be that

$$u_1^* \geq \sup_{\{u^K : K \geq 2\}} \left\{ \min_{\mu \in \triangle A} u_1(\mu) : u_1^k(\mu) \geq \underline{u}_1^k, u_i(\mu) \geq \underline{u}_i, \forall i, k \geq 2 \right\}.$$

Focusing on $K = 2$, the dual problem is

$$\sup_{u_1^2} \max_{\{p_i \geq 0 : i = 1, \ldots, N\}} p_1 \underline{u}_1^2 + \sum\nolimits_{i=2}^N p_i \underline{u}_i \text{ such that } p_1 u_1^2 + \sum\nolimits_{i=2}^N p_i u_i \leq u_1.$$

Since the constraint must bind, the reputation payoff is at least

$$\sup_{\{p_i \geq 0 : i = 2, \ldots, N\}} \text{val} \left(u_1 - \sum\nolimits_{i=2}^N p_i (u_i - \underline{u}_i 1) \right),$$

where 1 is a vector in $\mathbb{R}^{|A|}$ with all entries equal to one. Note that this lower bound is always larger than $\underline{u}_1$ (take $(p_2, \ldots, p_N) = 0$). The following theorem shows that this lower bound is actually achieved, and provides an alternative characterization of it.

Theorem 6.1 *The reputation payoff is equal to*

$$u_1^* = \sup_{\{p_i \geq 0 : i=2,\ldots,N\}} \mathrm{val}\,(u_1 - \sum_{i=2}^{N} p_i(u_i - \underline{u}_i 1)) = \sup_{\alpha_1 \in \triangle A_1} \min_{\alpha_{-1} \in Y(\alpha_1)} u_1(\alpha_1, \alpha_{-1}),$$

where $Y(\alpha_1) := \{\alpha_{-1} \in \triangle A_{-1} : u_i(\alpha_1, \alpha_{-1}) \geq \underline{u}_i, \forall i = 2, \ldots, N\}$. *The reputation payoff is achieved if* $K = N$ *and* $u_1^k = -u_k$, $\forall k = 2, \ldots, N$:

$$u_1^* = v_1(-u_2, \ldots, -u_N).$$

As is clear from the alternative characterization, the reputation payoff is lower than the usual *Stackelberg payoff*

$$\sup_{\alpha_1 \in \triangle A_1} \min_{\alpha_{-1} \in B(\alpha_1)} u_1(\alpha_1, \alpha_{-1}),$$

where $B(\alpha_1)$ is the set of Nash equilibria in the one-shot game between players $i = 2, \ldots, N$, given α_1. A *Stackelberg sequence* is any sequence $\{a_1^n\}_{n \in \mathbb{N}}$ achieving the supremum.

A game has *conflicting interest* if, for some Stackelberg sequence $\{a_1^n\}_{n \in \mathbb{N}}$, all Nash equilibria in $B(a_1^n)$ yield players $i \neq 1$ exactly their minmax payoff, for all $n \in \mathbb{N}$. It follows immediately from the theorem that player 1 can secure the Stackelberg payoff in all games of conflicting interest.

We further discuss the result after the proof.

Proof: Define

$$u_1' := \sup_{\{p_i \geq 0: i=2,\ldots,N\}} \operatorname{val}\left(u_1 - \sum\nolimits_{i=2}^N p_i(u_i - \underline{u}_i 1)\right)$$

and

$$u_1'' := \sup_{\alpha_1 \in \triangle A_1} \min_{\alpha_{-1} \in Y(\alpha_1)} u_1(\alpha_1, \alpha_{-1}).$$

We have already argued that $u_1^* \geq u_1'$. Let us first show that $u_1' \geq u_1''$. By definition, for all $\varepsilon > 0$, there exists $(p_2, \ldots, p_N) \geq 0$ and $\alpha_1 \in \triangle A_1$ such that

$$\begin{aligned}
u_1' - \varepsilon &\leq \operatorname{val}\left(u_1 - \sum\nolimits_{i=2}^N p_i(u_i - \underline{u}_i 1)\right) \\
&\leq \min_{\alpha_{-1}} \{u_1(\alpha_1, \alpha_{-1}) - \sum\nolimits_{i=2}^N p_i(u_i(\alpha_1, \alpha_{-1}) - \underline{u}_i)\} \\
&\leq \min_{\alpha_{-1}} \{u_1(\alpha_1, \alpha_{-1}) - \sum\nolimits_{i=2}^N p_i(u_i(\alpha_1, \alpha_{-1}) - \underline{u}_i 1) : \alpha_{-1} \in Y(\alpha_1)\} \\
&\leq \min_{\alpha_{-1}} \{u_1(\alpha_1, \alpha_{-1}) : \alpha_{-1} \in Y(\alpha_1)\} \leq u_1''.
\end{aligned}$$

Conversely, for every $\varepsilon > 0$, there exists $\alpha_1 \in \triangle A_1$ such that $\min_{\alpha_{-1} \in Y(\alpha_1)} u_1(\alpha_1, \alpha_{-1}) \geq u_1'' - \varepsilon$. Therefore, fixing $\alpha_1 \in \triangle A_1$, for every $\alpha_{-1} \in \mathbb{R}_+^{|A_{-1}|}$,

$$(u_i(\alpha_1, \alpha_{-1}) - \underline{u}_i \sum\nolimits_{a=1}^{|A_{-1}|} \alpha_{-1,a})_{i \neq 1} \geq 0 \Rightarrow u_1(\alpha_1, \alpha_{-1}) - (u_1'' - \varepsilon) \sum\nolimits_{a=1}^{|A_{-1}|} \alpha_{-1,a} \geq 0.$$

By Farkas' Lemma, there exists $(p_2, \ldots, p_N) \geq 0$ and a constant $\gamma \in \mathbb{R}_+^{|A_{-1}|}$ such that, for every $\alpha_{-1} \in \triangle A_{-1}$,

$$\begin{aligned}
u_1(\alpha_1, \alpha_{-1}) - u_1'' + \varepsilon &= \sum\nolimits_{i=2}^N p_i(u_i(\alpha_1, \alpha_{-1}) - \underline{u}_i) + \gamma \cdot \alpha_{-1} \\
&\geq \sum\nolimits_{i=2}^N p_i(u_i(\alpha_1, \alpha_{-1}) - \underline{u}_i).
\end{aligned}$$

Therefore,

$$u_1' + \varepsilon \geq \text{val}\,(u_1 - \sum_{i=2}^{N} p_i(u_i - \underline{u}_i 1)) + \varepsilon \geq u_1''.$$

We now show that the bound is attained by $u_1^k = -u_k$, $\forall k = 2, \dots, N$. Given some equilibrium, let $\mu^i \in \triangle A$ be the occupation measure when player 1 is of type i (the rational type is type 1). Player i's individual rationality is equivalent to, for all i, $u_i(\mu^i) \geq \underline{u}_i$. Further, player 1's individuality rationality condition states that, for every $p \in \triangle\{1, \dots, N\}$,

$$p_1 u_1(\mu^1) + \sum_{i=2}^{N} p_i(-u_i(\mu^i)) \geq \text{val}\,(p_1 u_1 - \sum_{i=2}^{N} p_i u_i),$$

and therefore, for the choice $p_i = 1, p_j = 0$, all $j \neq i$, it follows that $-u_i(\mu^i) \geq \text{val}\,(-u_i) = -\underline{u}_i$. Hence, $u_i(\mu^i) = \underline{u}_i$. Thus, we can rewrite the individual rationality condition as

$$u_1(\mu^1) \geq \text{val}\,(u_1 - \sum_{i=2}^{N} \frac{p_i}{p_1}(u_i - \underline{u}_i 1)),$$

i.e. $u_1(\mu^1) \geq u_1'$. Incentive compatibility of $(\mu^i)_i$ is obvious.

It remains to show that, for every choice of K and u_K, there always exists an equilibrium in which player 1's rational type does not exceed u_1'. Pick any such game. Let

$$v_1^k := \max_{\mu \in \triangle A} \{u_1^k(\mu) : u_1(\mu) \leq u_1', u_i(\mu) \geq \underline{u}_i, \forall i \geq 2\},$$

for all $k = 1, \dots, K$, with $u_1^1 = u_1$. Since $u_1' \geq \underline{u}_1$, the folk theorem under complete information ensure that the set on the right-hand side is non-empty, so that v_1^k is well-defined. Clearly, the action profiles α^k are incentive compatible, and individually rational for all players $i \geq 2$. It remains to show that it is incentive compatible for player 1, i.e., that for all $p \in \triangle\{1,\ldots,K\}$,

$$\sum_{k=1}^{K} p_k v_1^k \geq \mathrm{val}\,(\sum_{k=1}^{K} p_k u_1^k).$$

From the definition of v_1^k, it follows that for every $k = 1,\ldots,K$ and $\alpha \in \mathbb{R}_+^{|A|}$,

$$u_i(\alpha) \geq \underline{u}_i 1 \cdot \alpha, u_1' 1 \cdot \alpha \geq u_1(\alpha) \Rightarrow v_1^k 1 \cdot \alpha \geq u_1^k(\alpha).$$

By Farkas' Lemma, for every $k = 1,\ldots,K$, there exists $\gamma^k \geq 0, \lambda_i^k \geq 0$ such that $v_1^k 1 - u_1^k \leq \gamma^k(u_1' 1 - u_1) + \sum_{i=2}^{N} \lambda_i^k(u_i - \underline{u}_i 1)$. Therefore, for all $p \in \triangle\{1,\ldots,K\}$,

$$\mathrm{val}\,(\sum_{k=1}^{K} p_k u_1^k) \leq \sum_{k=1}^{K} p_k v_1^k - \sum_{k=1}^{K} p_k \gamma^k u_1' + \mathrm{val}\,(\sum_{k=1}^{K} p_k(\gamma^k u_1 - \sum_{i=2}^{N} \lambda_i^k(u_i - \underline{u}_i 1))),$$

and so individual rationality for player 1 is satisfied if

$$\sum_{k=1}^{K} p_k \gamma^k u_1' \geq \mathrm{val}\,(\sum_{k=1}^{K} p_k(\gamma^k u_1 - \sum_{i=2}^{N} \lambda_i^k(u_i - \underline{u}_i 1))).$$

This is satisfied if $\sum_{k=1}^{K} p_k \gamma^k = 0$, and if not, defining

$$\nu_i := (\sum_{k=1}^{K} \lambda_i^k p_k)/(\sum_{k=1}^{K} p_k \gamma^k) \geq 0,$$

it is equivalent to

$$u_1' \geq \mathrm{val}\,(u_1 - \sum_{i=2}^{N} \nu_i(u_i - \underline{u}_i 1))),$$

which is satisfied by definition of u_1'. $\square$

While there exists a large literature on reputation in two-player games, Ghosh (2007) is, to the best of our knowledge, the only paper considering reputations when the informed player faces multiple opponents. He considers the case in which the set of possible types includes all strategies of bounded recall mapping the past history of his opponents' actions into a pure action (these are not payoff types). He shows that the informed player rational type payoff can be as low as

$$l := \max_{a_1 \in A_1} \inf_{\alpha_{-1} \in W(a_1)} u_1(a_1, \alpha_{-1}),$$

where

$$W(a_1) := \{\alpha_{-1} \in \triangle A_{-1} : u_i(a_1, \alpha_{-1}) \geq \max_{a_1' \in A_1} \min_{\alpha_{-1,-i}} \max_{\alpha_i} u_i(a_1', \alpha_i, \alpha_{-1,-i}), \forall i = 2, \ldots, N\}.$$

References

[1] Aumann R.J. and M.B. Maschler, 1995. *Repeated Games with Incomplete Information*. The MIT Press.

[2] Bergemann, D. and S. Morris, 2007. "Belief Free Incomplete Information Games," Cowles Foundation Discussion Paper No. 1629, Yale University.

[3] Blackwell, D., 1956. "An Analog of the Minmax Theorem for Vector Payoffs," *Pacific Journal of Mathematics*, **6**, 1-8.

[4] Crémer, J. and R. McLean, 1985. "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, **53**, 345-361.

[5] Cripps, M. and J. Thomas, 2003. "Some Asymptotic Results in Discounted Repeated Games of One-Sided Incomplete Information," *Mathematics of Operations Research*, **28**, 433-462.

[6] Ely, J. and J. Välimäki, 2002. "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, **102**, 84-105.

[7] Ely, J., J. Hörner and W. Olszewski, 2005. "Belief-free Equilibria in Repeated Games," *Econometrica*, **73**, 377-415.

[8] Forges, F. 1992. "Non-zero-sum Repeated Games of Incomplete Information," in R. J. Aumann, S. Hart, eds. *Handbook of Game Theory*, Vol. 1. North Holland, Amsterdam, The Netherlands.

[9] Forges, F. and E. Minelli, 1997. "A Property of Nash Equilibria in Repeated Games with Incomplete Information," *Games and Economic Behavior*, **18**, 159-175.

[10] Fudenberg, D. and D. Levine, 1989. "Reputation and Equilibrium Selection in Games with a Single Patient Player," *Econometrica*, **57**, 759-778.

[11] Fudenberg, D. and E. Maskin, 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, **54**, 533-554.

[12] Fudenberg, D. and E. Maskin, 1991. "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, **53**, 428-438.

[13] Ghosh, S., 2007. "Multiple Opponents and the Limits of Reputation," working paper.

[14] Harsanyi, J.C., 1967-1968. "Games with incomplete information played by Bayesian players," *Management Science*, **14**, 159-182, 320-334, 486-502.

[15] Hart, S., 1985. "Nonzero-sum Two-person Repeated Games with Incomplete Information," *Mathematics of Operations Research*, **10**, 117-153.

[16] Hörner, J. and S. Lovo, 2008. "Belief-free equilibria in games with incomplete information," *Econometrica*, forthcoming.

[17] Israeli, E.,1999. "Sowing Doubt Optimally in Two-Person Repeated games," *Games and Economic Behavior*, **28**, 203-216.

[18] Kalai, E., 2004. "Large Robust Games," *Econometrica*, **72**, 1631-1665.

[19] Koren, G. 1992. "Two-Person Repeated Games where Players Know Their Own Payoffs," working paper, New York University.

[20] Monderer, D. and M. Tennenholtz, 1999. "Dynamic Non-Bayesian Decision Making in Multi-Agent Systems," *Annals of Mathematics and Artificial Intelligence*, **25**, 91-106.

[21] Piccione, M., 2002. "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, **102**, 70-83.

[22] Renault, J., 2001. "3-player repeated games with lack of information on one side," *Games and Economic Behavior*, **30**, 221-245.

[23] Renault, J., and T. Tomala, 2004. "Learning the State of Nature in Repeated Games with Incomplete Information and Signals," *Games and Economic Behavior*, **47**, 124-156.

[24] Shalev, J. 1994. "Nonzero-sum two-person Repeated Games with Incomplete Information and Known-own Payoffs," *Games and Economic Behavior*, **7**, 246-259.

[25] Sorin, S., 1986. "On Repeated Games with Complete Information," *Mathematics of Operations Research*, **11**, 147-160.

Appendix: Proof without communication device

Actions are periodically used as messages. Because players might have as few as two actions, each such communication phase might require several periods. As the actions played during this phase affect payoffs, communication phases must be short relative to regular phases. We shall not dispense with the randomization device altogether, as this allows us to achieve *exactly* the desired continuation payoff. Details on how to eliminate the public randomization device might be omitted altogether are the same as in the two-player case, following ideas introduced by Sorin (1986) and Fudenberg and Maskin (1991), and we refer the reader to Hörner and Lovo (2008).

Because communication requires several periods, strategies must also specify how a player plays within a communication phase if his own previous action already precludes him from reporting correctly his private information, or if his opponent's action already precludes his opponent from doing so. The construction must ensure that the specification is belief-free in both cases, and this explains why the construction that follows is more involved than one might have guessed. (In particular, it is the cause for the different kinds of communication phase described below.)

Play is divided into phases (or classes of phases): Communication phases, regular phases, penitence phases, and punishment phases.

Actions

Communication Phase

The *communication phase* replaces the communication stage. There are different versions of communication phase, denoted C, C_i, or C_i^*. (Roughly, a phase is indexed by player i if i's report during this phase is essentially ignored.[5]) A communication phase lasts c periods, where

$$c \geq 1 + \max_{i \in N} \frac{\ln |\Theta_i|}{\ln |A_i|},$$

so that $|A_i|^{c-1} \geq |\Theta_i|$, all $i \in N$. We fix two arbitrary but distinct actions for each player, denoted U and B, and a mapping

$$m_i : \Theta_i \to A_i^{c-1},$$

from his set of types into sequences of actions of length $c - 1$. Player i (or his play) *reports* θ_i if his play in the communication phase is equal to $(m_i(\theta_i), B)$ (so B is the action that he takes in the last period of this phase.) For any other play, he *reports* (U, n_i^U) where n_i^U is the number of periods in the communication phase in which $a_i = U$. We also write U rather than (U, n_i^U) whenever convenient, and let

$$\overline{\theta} \in \prod\nolimits_{i \in N} \Theta_i \cup \cup_{l=0}^{c} (U, l)$$

denote a *report*, or *message* profile. For $k \in K$, let $u_i^C(k, \overline{\theta})$ denote player i's average payoff from the communication phase if the state is k and the report is $\overline{\theta}$.[6]

In a communication phase C, player j's type θ_j plays the sequence $m_i(\theta_j, B)$, as long as his previous play in the phase does not preclude him from doing so. In a communication phase C_i so does player $j \neq i$, while player i plays (U, c). If a player's

[5]It cannot be entirely ignored, since we must give i incentives that do not depend on his type.

[6]This is an abuse of terminology, as payoffs are not uniquely identified by the report profile whenever a player reports U, since there might be many sequences of actions corresponding to this report. What is meant is the payoff given the actual sequence of action profiles.

past play prevents him from reporting his type θ_i, he plays U in every remaining period of the phase.

Transitions are described below.

Regular Phase

A *regular phase* is denoted $R(\overline{\theta}, \varepsilon)$, where $\kappa(\overline{\theta}) \neq \varnothing$, and $\varepsilon \in [-\overline{\varepsilon}, \overline{\varepsilon}]^N$, for some $\overline{\varepsilon} > 0$ to be specified.

A regular phase lasts at most n periods (to be specified), where $n > c$. We fix a (possibly correlated) mixed action profile $\mu(\overline{\theta}, \varepsilon) \in \triangle A$ such that, $\forall k \in \kappa(\overline{\theta})$, $\forall i \in N$, $\forall \varepsilon, \varepsilon' \in [-\overline{\varepsilon}, \overline{\varepsilon}]^N$ and $\forall \theta_i' \in \Theta_i, \theta_i' \neq \overline{\theta}_i$, such that $\kappa(\theta_i', \overline{\theta}_{-i}) \neq \varnothing$,

$$u_i^R(k, \mu(\overline{\theta}, \varepsilon)) := (1 - \delta^n) u_i(k, \mu(\overline{\theta}, \varepsilon)) + \delta^n u_i^C(k, \overline{\theta}) = v_i^k + \varepsilon_i,$$

and

$$u_i^R(k, \mu(\overline{\theta}, \varepsilon)) > u_i^R(k, \mu(\theta_i', \overline{\theta}_{-i}, \varepsilon')),$$

and

$$u_i^R(k, \mu(\theta_i', \overline{\theta}_{-i}, \varepsilon')) \leq v_i^k - 2\overline{\varepsilon}.$$

The strict inequalities can be satisfied for δ close enough to 1 and $\overline{\varepsilon}$ close enough to 0, since v is strictly incentive compatible.

In any period of the regular phase, players play $\mu(\overline{\theta}, \varepsilon)$. The regular phase $R(\overline{\theta}, \varepsilon)$ stops immediately after a unilateral deviation from $\mu(\overline{\theta}, \varepsilon)$, or if not, after n periods.

Transitions are described below.

Penitence Phase

A *penitence phase* is denoted $E(\overline{\theta}, \varepsilon)$, where $\varepsilon \in [-\overline{\varepsilon}, \overline{\varepsilon}]^N$, $\overline{\theta} \in \Theta$, $\kappa(\overline{\theta}) = \varnothing$, and $\overline{\theta} \in D$. A penitence phase lasts at most n periods. We fix a sequence $a(\overline{\theta}, \varepsilon) \in A^n$ such that $\forall (i, \theta_i') \in \Omega_{\overline{\theta}}$, $k \in \kappa(\theta_i', \overline{\theta}_{-i}), \varepsilon \in [-\overline{\varepsilon}, \overline{\varepsilon}]^N$,

$$u_i^E(k, a(\overline{\theta}, \varepsilon)) := \frac{1-\delta}{1-\delta^n} \sum\nolimits_{t=0}^{n-1} \delta^t u_i(k, a_t(\overline{\theta}, \varepsilon)) < v_i^k - 2\overline{\varepsilon}.$$

Such a penitence phase $E(\overline{\theta}, \varepsilon)$ stops immediately after a unilateral deviation from the sequence $a(\overline{\theta}, \varepsilon)$, or if not, after n periods. In period t of the penitence phase, players play $a_t(\overline{\theta}, \varepsilon)$.

Transitions are described below.

Punishment Phase

A *punishment phase*, indexed by i, is denoted $P_i(\overline{\theta}_{-i}, t)$, where $\overline{\theta}_{-i} \in \Theta_{-i}$ is such that $\kappa(\overline{\theta}_{-i}) \neq \varnothing$ and $t = n$ or T (to be defined) denotes the length of the punishment phase.

As before, we fix an action $\underline{a}_i \in A_i$ and let $s_i^{\underline{a}_i}$ denote the strategy of playing $\underline{a}_i$ in every period, independently of the history. In the punishment phase, player i uses $s_i^{\underline{a}_i}$, and players $-i$ use $s_{-i}^{\overline{\theta}_{-i}}$.

We pick $n, T, \overline{\delta} < 1$ and $\overline{\varepsilon}$ such that, $\forall \delta > \overline{\delta}$, $\forall k \in \kappa(\overline{\theta}_{-i})$, player i's average discounted payoff over the t periods in state k is no larger than $v_i^k - 2\varepsilon$, and that it is sufficiently larger when $t = n$ than when $t = T$, as explained below. This is possible since v satisfies individual rationality strictly.

We shall write C, R, E, P for a communication, regular, penitence and punishment phase without further argument when there is no risk of confusion.

Transitions

Given any message $\overline{\theta}$, define

- whenever $\overline{\theta} \in \Theta$, $\forall \theta \in \Theta$, $\Delta_I(\theta, \overline{\theta}) := \{i \in N | \theta_i \neq \overline{\theta}_i\}$;

- whenever $\overline{\theta} \in \Theta$, $\overline{\theta} \in D$, $\Delta_D(\overline{\theta}) := \{i \in N | (i, \theta_i') \in \Omega_{\overline{\theta}}$ for some $\theta_i' \in \Theta_i\}$;

- whenever $\overline{\theta} \notin \Theta$, $\Delta_U(\overline{\theta}) := \{i \in N | \overline{\theta}_i \notin \Theta_i\}$.

Given a unilateral deviation from a sequence $a(\overline{\theta}, \varepsilon)$, or from a mixed action $\mu(\overline{\theta}, \varepsilon)$, let Δ_A denote the index of the player who deviated.[7] Finally, given a set $\Delta \subset N$, let $-\Delta := N \setminus \Delta$.

From a communication phase

The transition depends on the message $\overline{\theta}$ during C, the phase $\Phi \in \{R, P, E, C\}$ immediately preceding C, and the play during Φ. Roughly speaking, if there is no unilateral deviation during Φ, and if $\overline{\theta} \in \Theta$, a regular or a penitence phase follows, while if $\overline{\theta} \notin \Theta$, either a punishment or a communication phase follows. If there is a unilateral deviation during Φ by player i, then if $\overline{\theta}_{-i} \in \Theta_{-i}$, a punishment phase follows. More precisely, if there is a unilateral deviation from $\Phi = E, R$, with $\Delta_A = \{i\}$, then the next phase is

[7]Recall that there is a public randomization device, so that we always assume that players use a pure action profile, as a function of the realization of the public randomization device, so that the mixed action profile obtains in expectations.

1. if $\overline{\theta}_{-i} \in \Theta_{-i}$, $\kappa(\overline{\theta}_{-i}) \neq \varnothing$: $P_i(\overline{\theta}_{-i}, T)$;

2. otherwise, it is C.

On the other hand, if there is no unilateral deviation from Φ, or if $\Phi = P, C$, and

1. Φ equals $R(\theta, \varepsilon)$ or $E(\theta, \varepsilon)$, the next phase is:

   (a) if $\overline{\theta} \in \Theta$, $\kappa(\overline{\theta}) \neq \varnothing$: $R(\overline{\theta}, \varepsilon_{-\Delta_I(\theta,\overline{\theta})}, -\overline{\varepsilon}_{\Delta_I(\theta,\overline{\theta})})$;

   (b) if $\overline{\theta} \in \Theta$, $\overline{\theta} \in D$: $E(\overline{\theta}, \varepsilon_{-\Delta_D(\overline{\theta})}, -\overline{\varepsilon}_{\Delta_D(\overline{\theta})})$;

   (c) if $\Delta_U(\overline{\theta}) = \{i\}$, $\kappa(\overline{\theta}_{-i}) \neq \varnothing$: $P_i(\overline{\theta}_{-i}, n)$;

   (d) otherwise, C;

2. Φ equals $P_i(\theta_{-i}, t)$, $t = n, T$, the next phase is:

   (a) if $\overline{\theta} \in \Theta$, $\kappa(\overline{\theta}) \neq \varnothing$: $R(\overline{\theta}, \tilde{\varepsilon}(\theta, \overline{\theta}))$, where $\tilde{\varepsilon}_i(\theta, \overline{\theta}) \in [-\overline{\varepsilon}, \overline{\varepsilon}]$ is chosen so that, given $\overline{\theta}$ and $s_{-i}^{\theta_{-i}}$, using $s_i^{\underline{a}_i}$ is optimal in the punishment phase for player i; and further, if $\overline{\theta}_{-i} = \theta_{-i}$, player i's continuation payoff in the repeated game, evaluated at the beginning of the punishment phase, is equal to, for all $k \in \kappa(\overline{\theta})$,

   $$(1 - \delta^t)(v_i^k - 2\overline{\varepsilon}) + \delta^t(v_i^k - \overline{\varepsilon});$$

   and for $j \neq i$, $\tilde{\varepsilon}_j(\theta, \overline{\theta})$ is chosen so that, given $\overline{\theta}$, $s_{-j}^{\theta_{-j}}$ and $s_i^{\underline{a}_i}$, $s_j^{\theta_j}$ is optimal for player j in the punishment phase. Further $\tilde{\varepsilon}_j(\theta, \overline{\theta}) \in [\overline{\varepsilon}/4, 3\overline{\varepsilon}/4]$ if $\theta_j = \overline{\theta}_j$ and $\tilde{\varepsilon}_j(\theta, \overline{\theta}) \in [-3\overline{\varepsilon}/4, -\overline{\varepsilon}/4]$ otherwise;

   (b) if $\overline{\theta} \in \Theta$, $\overline{\theta} \in D$: $E(\overline{\theta}, 0_{-\Delta_D(\overline{\theta})}, -\overline{\varepsilon}_{\Delta_D(\overline{\theta})})$;

   (c) otherwise, C.

3. Φ equals C, or C_i and θ is the report during Φ, the next phase is:

(a) if $\overline{\theta} \in \Theta$, $\kappa(\overline{\theta}) \neq \varnothing$: $R(\overline{\theta}, \hat{\varepsilon}(\theta, \overline{\theta}))$, where, if $\Phi = C$, or $j \neq i$,

$$\hat{\varepsilon}_j(\theta, \overline{\theta}) = \begin{cases} 0 & : \quad \theta_j = \overline{\theta}_j, \\ -\overline{\varepsilon}/4 + \rho n_U & : \quad \theta_j = (U, n_U), \\ -\overline{\varepsilon} & : \quad \text{otherwise,} \end{cases}$$

and if $\Phi = C_i$,

$$\hat{\varepsilon}_i(\theta, \overline{\theta}) = \begin{cases} -\overline{\varepsilon} + \rho n_U & : \quad \theta_i = (U, n_U), \\ -\overline{\varepsilon} & : \quad \text{otherwise,} \end{cases}$$

for some $\rho > 0$ to be defined;

(b) if $\overline{\theta} \in \Theta$, $\overline{\theta} \in D$: $E(\overline{\theta}, 0_{-\Delta_D(\overline{\theta})}, -\overline{\varepsilon}_{\Delta_D(\overline{\theta})})$;

(c) if $\Delta_U(\overline{\theta}) = \{i\}$, $\kappa(\overline{\theta}_{-i}) \neq \varnothing$: $P_i(\overline{\theta}_{-i}, n)$;

(d) otherwise, C.

From any other phase

Any other phase is followed by a communication phase. If there is a unilateral deviation from a phase $\Phi = R, E$, with $\Delta_A = \{i\}$, it is a communication phase C_i; otherwise, it is a communication phase C.

Initial phase

The game starts with a communication phase, at the end of which transitions occur as if the previous phase had been C, with $\theta = \overline{\theta}$, and $\varepsilon \in [-\overline{\varepsilon}, \overline{\varepsilon}]$ is such that the payoff (inclusive of the initial communication phase) is equal to v.

Verification of optimality

Consider first the incentives of player i to deviate during a regular phase. If he does so, a punishment phase P_i will start after the communication phase. Player i expects the type profile by θ_{-i} reported after the deviation and before the punishment phase to be correct; since his payoff at the beginning of the punishment phase is then

$$(1-\delta^T)(v_i^k - 2\overline{\varepsilon}) + \delta^T(v_i^k - \overline{\varepsilon}),$$

he has no incentive to deviate in this case, as whether or not his own report was correct, his payoff from following the equilibrium strategies is higher.[8]

Consider next a punishment phase P_i. The definition of $\tilde{\varepsilon}_i$ guarantees that $s_i^{\underline{a}_i}$ is optimal for player i. Similarly, the definition of $\tilde{\varepsilon}_j$ ensures that player $j \neq i$ has no incentive to deviate. This is true whether the punishment phase lasts n or T periods.

Consider next a possible deviation during the penitence phase. While the average payoff from the penitence phase is low, observe that it lasts only n periods (and, given the equilibrium strategies, the ensuing communication phase will be followed by a regular phase if the player refrains from deviating, independently of the history up to the contemplated deviation), while the punishment phase that the deviation would trigger lasts T periods. We pick T and n so as to ensure that no such deviation is profitable.

Consider finally a possible deviation during a communication phase. Start with a communication phase C.

[8]Note that the situation where the reported type profile by $-i$ is incorrect is not relevant for verifying that player i does not deviate during the regular phase. This is because, at the time of the deviation, he expects the other player to report correctly their type during the communication phase.

1. Assume first that the history in the communication phase is consistent with (possibly, among others) some type profile $\theta \in \Theta$ (i.e., the history in the communication phase is an initial segment of $(m_1(\theta_1), \ldots, m_N(\theta_N)))$, and θ_i is indeed player i's type. If the true state is θ, then by reporting U, a punishment phase P_i of length n will be entered, the expected payoff of which ensures that it is better not to do so. If the true state is not θ, then according to the equilibrium strategies, some player $j \neq i$ will report U in this communication phase. If player i reports U, a communication phase C will be entered, at the end of which a regular phase will be started, for which $\varepsilon_i < 0$ (pick ρ such that $-\bar{\varepsilon}/4 + \rho c < 0$); by sticking to the report of θ_i, either a communication phase C will start (in case θ_j and $\theta_{j'}$ differ from the true state for two players j, j'), in which case, in the ensuing regular phase, player i's ε_i is zero, or a punishment phase of length n will start, at the end of which, in the ensuing regular phase, player i's ε_i is at least $\bar{\varepsilon}/4$; of course, i's payoff during the n periods can be very low, but we can deter such deviations by picking ρ sufficiently small (but not too small, see below).

2. Assume next that the history in the communication phase is consistent with some type profile $\theta \in \Theta$, but θ_i is not player i's type. Thus, the equilibrium strategy calls for player i to report U (if there is at least one period; otherwise, there is nothing to show). Suppose first that the other players' type profile is indeed θ_{-i}. By reporting U, player i triggers a punishment phase P_i of length n, but by failing to do so, it triggers the play of a regular phase for which the play does not correspond to the true type profile. We can pick n small enough to guarantee that, since the payoff during such a regular phase is less that $v_i - \bar{\varepsilon}$, player i prefers not to deviate. Suppose next that there exists exactly one other player j for which θ_j is not the true type. By

reporting U, a second communication phase starts, but player i is guaranteed at least a $\varepsilon_i \geq -\overline{\varepsilon}/4$ in the regular phase at the end of it; if player i persists in reporting the incorrect type, a punishment phase P_j of length n follows, at the end of which player i's ε is strictly less than $-\overline{\varepsilon}/4$; finally, if there are two or more other players for which θ_j is incorrect, and if player i reports U, he also guarantees that, in the regular phase that will follow the second communication phase, $\varepsilon_i \geq -\overline{\varepsilon}/4$; if he reports differently, in the regular phase that will follow the second communication phase, $\varepsilon_i = -\overline{\varepsilon}$.

3. Assume finally that the history in the communication phase is not consistent with some type profile $\theta \in \Theta$, i.e. some player reports U already. The same arguments as before apply almost *verbatim*, since in the previous arguments, if θ_j was not the true type for one or more players, those players j were about to report U anyway. Note that postponing a report of U by one or more periods within a communication phase is suboptimal, since the argument ε_i from the relevant ensuing regular phase is increasing in the number of times player i choose U. (This is where we need that ρ be not too small, more precisely, it must be at least $(1-\delta)M$).

These arguments are readily adapted to the case in which the communication phase is C_i. Consider first the case in which the previous phase was E or R (i.e., player i deviated in actions). Suppose first that the other players' type profile θ_{-i} is consistent with the history in the communication phase. Since the equilibrium calls for a punishment phase to follow, the specification of $\tilde{\varepsilon}_j$, $\tilde{\varepsilon}_i$ ensures that no player gains from deviating: i.e., player i benefits from playing U as often as possible, and other players gain by reporting their type truthfully. Suppose now that the history in the communication phase is not consistent with some type profile $\theta_i \in \Theta_i$, then some player $-i$ will play U and a new communication phase C will follow. Also in this case player i benefits from playing U since $\varepsilon_i = -\overline{\varepsilon} + \rho n_U$ in the regular phase that will follow the new communication C .