

Belief-free Equilibria in Games with Incomplete Information: Characterization and Existence*

Johannes Hörner[†], Stefano Lovo[‡], Tristan Tomala[§]

November 10, 2009

Abstract

We characterize belief-free equilibria in infinitely repeated games with incomplete information with $N \geq 2$ players and arbitrary information structures. This characterization involves a new type of individual rational constraint linking the lowest equilibrium payoffs across players. The 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. Further, we provide necessary conditions and sufficient conditions on the information structure for this set to be non-empty, both for the case of known-own payoffs, and for arbitrary payoffs.

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

JEL codes: C72, C73

*We thank Pierpaolo Battigalli, Drew Fudenberg, David Levine, George Mailath, Jérôme Renault, Ariel Rubinstein, Larry Samuelson for useful comments and seminar audiences at the GDR Conference in Luminy 2009, Stony Brook 2008 and 2009, SAET 2009, as well as at Arizona State University, Bocconi University, Boston University, Duke University, Ecole Polytechnique, Hong-Kong University, Paris 1 University, Princeton University, Warwick University, University of Rochester, Washington University at Saint Louis, Yale University.

[†]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA. 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. T. Tomala gratefully acknowledges the support of the ANR project “Croyances,” and of the Fondation du Risque, Chaire Groupama, “Les particuliers face au risque.” S. Lovo gratefully acknowledges financial support from the HEC Foundation.

1 Introduction

This paper characterizes the set of payoffs achieved by equilibria that are robust to the specification of beliefs, and provides necessary and sufficient conditions for its non-emptiness. We consider n -player repeated games with incomplete information and low discounting. This class of equilibria has been introduced by Hörner and Lovo (2009) 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 complete 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 complete information.

In Hörner and Lovo (2009), the analysis is restricted to two-player games, and the players' private information has a "product" structure. That is, the information structure can be represented as a matrix. Each state of nature corresponds to a cell in this matrix. Player 1 is informed of the true row, while player 2 is informed of the true column. This paper generalizes these results to the most general setting:

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 state of nature. That is, the state of the world need not be distributed knowledge.

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 instead that each player has a subjective probability distribution over those states of nature that the players' combined information cannot distinguish, and use this distribution to treat each such set as a singleton. 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 less restrictive definitions, by replacing for instance such col-

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

lections of states by a single state, and payoffs in that state by the relevant expectations.² From a theoretical point of view, it is unclear to us why an optimality criterion used by a single decision-maker should depend on whether those states that he cannot distinguish can be distinguished collectively or not.

The focus of the analysis is on the set of belief-free equilibrium payoff vectors as the discount factor tends to one. We provide a set of necessary conditions defining a closed, convex, and possibly empty set. These necessary conditions have simple interpretations in terms of incentive compatibility, individual rationality in every state, and *joint rationality*, an additional requirement absent from the earlier analysis for two-player games, and that is related to the fact that, because strategies depend on private information, there might be histories after which it is not possible to uniquely identify the deviator. Conversely, we prove that every payoff vector in the interior of this set is a belief-free equilibrium payoff vector provided that 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 necessary and sufficient conditions on the information structure for non-emptiness of this set for different classes of payoff functions. With two players, for instance, non-emptiness was already known to obtain if each player knows his own payoff, and one player is informed of the state. For general payoff functions, the necessary and sufficient condition is that no two players are *essential* (as defined in Section five) in distinguishing between any two states. This result is due to Renault and Tomala (2004) for undiscounted games and we adapt it to our setup. Our main result provides both a necessary and a sufficient condition for the important case of known-own payoffs (KOP). In that case, non-emptiness obtains for all payoff functions satisfying KOP only if a given information structure satisfies the following. Divide the states into

²Note that in this case the payoff function will depend on the beliefs used to compute such expectations.

the finest partition with the property that for any two states lying in distinct cells of this partition, at least three players distinguish them (i.e. get different signals for those two states), and restrict attention to the projection of the information partition on any given cell. Then for each state k , there must exist a player i who is as well informed as all others at that state. Further, either no player can distinguish any two states for which he is not the best informed player (if he ever is), or there is a second player $j \neq i$ who is as well informed as all players but i at that state. This latter case is shown to be sufficient. Our next result states that, 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. Finally, for the class of payoff function that satisfy both KOP and the bad outcome property, we show that there must be at most one essential player per state.

A special class of games covered by these conditions is the class of “reputation” games in which there is exactly 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.

The set of belief-free equilibrium payoffs has already appeared in the literature, most notably (but not only) for two players, 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). For more than two players, Renault (2001) studies three-player games with two informed players and one uninformed player, and introduces the joint

rationality condition in this context. Renault and Tomala (2004a) study existence for all payoff functions in the n -player case.

Our work is also related to the literature on existence of equilibria for non-zero-sum undiscounted games with incomplete information. It is known since Aumann and Maschler (1995) that some conditions on information structures are required to get existence. Sorin (1983) shows existence of belief-based equilibrium in two-player games with one-sided incomplete information and two states of nature. Simon, Spież and Toruńczyk (1995) extend this result to an arbitrary number of states. For more than two players, no general result is known. See for instance Renault (2001) for 3-player games with lack of information on one-side.

Israeli (1999) provides an analysis of reputation in two-player undiscounted games, to which our own analysis of reputation owes a great deal. Further references to non-Bayesian studies can be found in Hörner and Lovo (2009). Finally, Peški (2008) considers discounted games with known-own payoffs, two states of the world, and one informed player. He defines the set of payoffs that satisfy both individually rationality after every history, and incentive compatibility, and shows that its closure is equal to the limit set (as the discount factor tends to one) of the Nash equilibrium payoffs, under full dimensionality. Therefore, his result shows that, at least in his set-up, the notion of individual rationality that captures Nash equilibrium is expected individual rationality after every history (where the expectation is, for the uninformed player, with respect to his beliefs about the state). In contrast, the notion of individual rationality that captures belief-free equilibrium is individual rationality for every state (what he calls *IR-in-every-state*.) The equivalence of those two notions of individual rationality in the case of undiscounted payoffs is the main reason why the characterization of belief-free equilibrium payoffs is reminiscent of some of the results in the literature on Nash equilibrium payoffs of undiscounted games. Understanding the relationship between the two payoff sets in general environments is an important open question.

Belief-free equilibrium is also related to ex post equilibrium, 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 notion of belief-free equilibria has been introduced in games with imperfect monitoring. See Piccione (2002), 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.

The most closely related papers are Hörner and Lovo (2009), already discussed, and Fudenberg and Yamamoto (2009a, 2009b). Fudenberg and Yamamoto (2009b), which itself generalizes Fudenberg and Yamamoto (2009a), is complementary to this paper. By combining belief-free equilibrium with perfect public equilibrium, they extend the analysis to the case of repeated games with incomplete information, and imperfect and unknown monitoring. That is, players receive imperfect public signals and the map from actions into signal distributions is itself unknown. Their contribution is two-fold. First, they develop linear algebraic techniques to study the limit payoff set, whose usefulness is illustrated via examples. Second, they use these techniques to provide sufficient conditions for the folk theorem to hold. The latter contribution is especially important, as it provides conditions under which, as far as limit payoffs are concerned, the restriction to these equilibria is without loss of generality.

The paper is related more broadly to the literature on the robustness of equilibrium in repeated games. Miller (2009) develops a related notion, in which the ex post requirement is imposed in each period, but players' continuation payoffs are evaluated according to their beliefs. Chassang and Takahashi (2009) examine the robustness of equilibria to incomplete information that is modelled

by payoff shocks that are independent across periods. Wiseman (2008) considers the case in which the payoff matrix is unknown, but players learn over time, and provides conditions under which a folk theorem obtains.

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 finite set of players is $N := \{1, \dots, N\}$. 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 $K := \{1, \dots, K\}$. Given a set S , let ΔS denote the probability simplex over S , $1\{S\}$ the indicator function of S , $|S|$ the cardinality of S , $\text{int } S$ the interior of S , and $\text{co } S$ 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 $u_i : K \times A \rightarrow \mathbb{R}$. Let $M := \max_{i \in N, k \in K, a \in A} |u_i(k, a)|$. A reward profile is denoted $u := (u_1, \dots, u_N)$. Mixed actions of player i are denoted α_i . The definition of rewards is extended to mixed, possibly correlated, action profiles $\mu \in \Delta A$ 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. Without loss of generality (see Aumann, 1976), this process can be represented by an information structure $\mathcal{I} := (\mathcal{I}_1, \dots, \mathcal{I}_N)$, 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 all players but i . We do not require that $\kappa(\theta) \neq \emptyset$: it might be that some type profile cannot arise. Similarly, it might be that $|\kappa(\theta)| > 1$: the join of the players' information partitions need not reduce to the state. The information partitions are common knowledge, but the realized signal is private information.

The game is infinitely repeated, with periods $t = 0, 1, 2, \dots$. A history of length t is a vector $h^t \in H^t := A^t$ ($H^0 := \{\emptyset\}$). An outcome is an infinite history $h \in H := A^\infty$. Neither mixed actions nor realized payoffs are observed. On the other hand, realized actions are perfectly observed. A behavior strategy for player i 's type θ_i is a mapping $\sigma_{i,\theta_i} : \cup_{t \in \mathbb{N}} H^t \rightarrow \Delta A_i$. We write $\sigma_i := \{\sigma_{i,\theta_i}\}_{\theta_i \in \Theta_i}$ for player i 's strategy, and $\sigma := (\sigma_1, \dots, \sigma_N)$ for a strategy profile.

Players use a common discount factor $\delta < 1$. The *payoff* of player i in state k is the expected average discounted sum of rewards, where the expectation is taken with respect to mixed action profiles. That is, given some outcome $h = (a_0, \dots, a_t, \dots)$, player i 's payoff in state k is

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

As usual, the domain of rewards is extended to mixed action profiles and strategy profiles. Given a strategy profile σ , let $\mu_k \in \Delta A$ denote the *occupation measure* over action profiles induced by σ when the state is k , that is, for every $a \in A$,

$$\mu_k(a) := (1 - \delta) \mathbb{E}_\sigma \left[\sum_{t \geq 0} \delta^t 1\{a_t = a\} \right].$$

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

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

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 an 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 $\lim_{\delta \rightarrow 1} B_\delta$ (a limit that is shown to be well-defined) and establish conditions under which this limit set is non-empty.

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, incentive compatibility, and (individual and joint) rationality.

3.1 Feasibility

Definition: The payoff vector $v \in \mathbb{R}^{NK}$ is *feasible* if there exists $(\mu_k)_{k \in K} \in (\Delta A)^K$ 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 is the obvious feasibility condition. That is, there exists an occupation measure μ_k that yields the payoff vector v^k .

The second condition is rather a measurability restriction. It states that, if players cannot collectively distinguish two states, then the equilibrium occupation measures over action profiles must be the same in 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 $\mu_{\theta'_i, \theta_{-i}}$ in 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 $\kappa(\theta_i, \theta_{-i}) \neq \emptyset$ and $\kappa(\theta'_i, \theta_{-i}) \neq \emptyset$. 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. Consider now the game of complete information in which 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 and Joint 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. Since we seek to identify here a necessary condition that player i 's equilibrium payoff vector must satisfy, the more effective the punishment, the weaker the condition. Therefore, we may start by assuming that, if player i deviates, all other players commonly know the information that is distributed among them, as these are the most favorable conditions for a punishment. This is also the reason why we may assume that player i 's deviation is common knowledge, even if, for some deviations by i , this need not be.

Still, if the set of states $\kappa(\theta_{-i})$ is not a singleton, players $-i$ cannot tailor the punishment strategy to the actual state of the world. Suppose, for instance, that $\kappa(\theta_{-i}) = \{1, 2\}$, as illustrated in Figure 1. Because player $-i$'s strategy, after such a deviation, must be effective in both games of complete information simultaneously, it must guarantee that player i 's payoff is lower than v_i in both its coordinates, independently of what strategy player i uses. Note that it is irrelevant whether player i can distinguish these two states himself.

Determining for which values of the vector v_i players $-i$ have such a strategy available may appear a formidable task, but as is well-known, this is by definition equivalent (at least in the undiscounted case) to the orthant $W := \{v_i\} - \mathbb{R}_+^2$ being an approachable set, and necessary and sufficient conditions for this are provided by Blackwell (1956).




Figure 1: Players $-i$ must have a strategy that guarantees that i 's payoff lies in W .

To this end, define, for $\theta_{-i} \in \Theta_{-i}$,

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

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

$$\forall q \in \Delta \kappa(\theta_{-i}) : \sum_{k \in \kappa(\theta_{-i})} q(k) v_i^k \geq \varphi_{i,\theta}(q). \quad (IR(i, \theta_{-i}))$$

These inequalities are the immediate generalizations of the individual rationality conditions for the two-player case. Note that if $\kappa(\theta_{-i}) = \emptyset$, the inequality is vacuously satisfied. If $\kappa(\theta_{-i})$ is a singleton set $\{k\}$, the inequality reduces to the familiar definition of individual rationality under complete information, i.e. $v_i^k \geq \text{val } u_i(k, \cdot)$, where $\text{val } u_i(k, \cdot)$ denotes player i 's minmax payoff in state k . In the definition of $\varphi_{i,\theta}$, note that the action of players $-i$ 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 one of these conditions is violated, there necessarily exists one player, a type profile

θ_{-i} and $q \in \Delta\kappa(\theta_{-i})$ such that the reverse inequality holds. This implies that for every α_{-i} , there exists $a_i(\alpha_{-i}) \in A_i$ such that

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

Assume instead that v is in B_δ and let σ be the corresponding equilibrium. Note that players $-i$ play the same strategy in each state $k \in \kappa(\theta_{-i})$. Consider thus the strategy τ_i of player i that plays $a_i(\alpha_{-i})$ after a history h^t such that $\sigma_{-i}(h^t) = \alpha_{-i}$. The reward of player i under (τ_i, σ_{-i}) satisfies the inequality (1) and therefore, so does the 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 . 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.

While individual rationality is a necessary condition, it is not the only one. There are other conceivable deviations, leading to an additional necessary condition. In particular, even if a deviation gets detected, it might not be possible to identify the deviator. 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 among all players that some player deviated, but not necessarily whether it is player i or j . With two players, of course, the identity of the deviator is always common knowledge.

To be more formal, let D be the set of type profiles that are inconsistent, but could arise if there was a unilateral deviation. That is, θ is in D if $\kappa(\theta) = \emptyset$ and $\Omega_\theta := \{(i, \theta'_i) \mid i \in N, \kappa(\theta'_i, \theta_{-i}) \neq \emptyset\} \neq \emptyset$. In other words, if players were to report their types, and the reported profile was in D , all players would know that one player must have lied. 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

$$\exists \mu \in \Delta A, \forall (i, \theta'_i) \in \Omega_\theta, \forall k \in \kappa(\theta'_i, \theta_{-i}) : v_i^k \geq u_i(k, \mu). \quad (JR(\theta))$$

These inequalities are called Joint Rationality (JR), since they involve payoffs of different players simultaneously.³ Note that joint rationality does not imply individual rationality (there is no requirement that player i 's action be a best-reply), nor is it implied by it.

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 outcomes 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 mimics type θ_i and the one in which player j consistently mimics type θ_j . Let $\mu \in \Delta A$ denote the occupation

³Joint Rationality has been first introduced in Renault (2001) in a three-player setup.

measure generated by σ_θ . If $JR(\theta)$ is violated, there exists a player i and a state $k \in \kappa(\theta_{-i})$ 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. \square

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

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

which suggests interpreting the identity of the deviator as part of the uncertainty itself. 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 . We show that this set characterizes the set of belief-free equilibrium payoff vectors, up to its boundary points.

Let $\hat{K} := \{k \in K : \bigcap_{i \in N} I_i(k) \neq \{k\}\}$ 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 $|A|$ columns, where k belongs to \hat{K} . The reward function 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 first 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 appendix, 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 without such a device or communication is rather standard but very lengthy, and can be found in the working paper (Hörner, Lovo and Tomala, 2009).

The rank assumption serves a similar purpose, as it allows players to provide appropriate incentives in states that cannot be distinguished.

It is worth making the following two remarks. First, 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 vectors, 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, under the natural ordering on such structures. Second, note that the *IC*, *IR* and *JR* conditions remain necessary even if we drop the sequential rationality constraint imposed by subgame-perfection. That is, the same characterization would hold if belief-free equilibria was defined with respect to Nash equilibria of the underlying complete information game.

5 Existence

Our main theorem states that, given $V^* \neq \emptyset$, 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 necessary and sufficient conditions for non-emptiness of V^* . We shall not address the issue of whether boundary points of V^* are themselves equilibria or not. Even in the case of complete information, it is not known under which conditions minmax payoffs are equilibrium payoffs themselves (this is the case, generically, when attention is restricted to pure strategies and there exist points in the feasible payoff set that give each player his minmax payoff (Thomas, 1995)), and such conditions appear all the more elusive here given that both IR and JR are multi-dimensional versions of individual rationality. Incentive compatibility, however, is an additional condition, and we will comment on when it can be made strict (this is the case, for instance, for our first set of results). As a practical matter, it is immediate to apply the characterization of V^* to verify that the set has non-empty interior. Note that Fudenberg and Yamamoto (2009b) provide useful sufficient conditions for this to be the case. Note also that, as mentioned, V^* has been shown to play an important role in the study of Nash equilibria in repeated games without discounting, for those special cases in which such a characterization has been obtained so far.

More precisely, we consider different classes of games each 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 payoff vectors satisfying IC , IR and JR , and provide counter-examples within those classes for the necessity part. 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 reward 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 .

We might wish to examine for which information structures non-emptiness obtains for all reward functions, or for all reward functions within some class $\mathcal{S} \subseteq \mathcal{U}$. We shall consider this first. Second, we examine for which reward functions non-emptiness obtains independently of the

information structure. This, in particular, will ensure existence for the applications in which the assumption that the information partitions are common knowledge appears exorbitant. We shall address this next. Proofs are outlined in the text and, when necessary, detailed in the appendices B–E.

5.1 Majority Components

It is useful to identify the information that can be readily disclosed either because it is shared by sufficiently many players or, for 2-player games, because it is common knowledge. For instance, if three (or more) players know the state of nature, it is straightforward to provide those players with strict incentives to disclose it: each informed player reports the true state (through an appropriate choice of actions); under any unilateral deviation, there are still at least two players (a *majority*) among informed players who report it truthfully. Truth-telling is thus optimal, and the state is revealed.

More generally, we shall make precise the information about the state that can be made common knowledge among players even under unilateral deviations. This will define a partition over the set of states K . An element of this partition is a *majority component*. That is, if the true state k belongs to the majority component A , then under strategies that ask players to report whether the state is in A or not, it becomes common knowledge that the true state lies in A once the reports are made, and even if a player unilaterally deviates.

This requires that, for every $k'' \in K \setminus A$, at least three players know that the state is not k'' , so that, even if one of them deviates, at least two players' reports rule out k'' . Conversely, if two states k and k' belong to the same majority component A , then, for some report of some player, there are no two other players who could, by reporting truthfully, distinguish between k and k' .

To define a majority component formally, we introduce the following equivalence relation.

Definition 5.1

- For each pair of states k, k' , let $\nu(k, k')$ be the number of players who distinguish k from k' . Define the binary relation R by kRk' iff $\nu(k, k') \leq \min\{2, N - 1\}$.
- Let $k \sim k'$ iff there is a chain of states $k = k_1, k_2, \dots, k_n = k'$ such that $k_m R k_{m+1}$ for each m . A majority component of K is an equivalence class of this relation.

Note that R is symmetric but not necessarily transitive, and \sim is the transitive closure of R (i.e. the smallest transitive extension of R), thus it is an equivalence relation.

If A, B are two distinct majority components of K , then for each $k \in A$ and each k' in B , $\nu(k, k') \geq 3$. Otherwise, there would exist a link (for the relation R) between some point in A and some point in B , and thus a chain linking any point in A to any point in B . Note that for 2-player games two states belong to the same majority component only if they can be distinguished by at most one player.

The study of belief-free equilibria can be made on each majority component separately. Given $A \subseteq K$, let \mathcal{I}_A denote the information structure on A induced by \mathcal{I} :

$$I_{A,i}(k) = I_i(k) \cap A, \quad \forall i \in N, \forall k \in A.$$

Note that, by definition, a BFE given K and \mathcal{I} must induce a BFE given A and \mathcal{I}_A . If A is a majority component, the discussion above can be summarized in the following lemma.

Lemma 5.2 $V^*(u, \mathcal{I}) \neq \emptyset$ iff for each majority component A , $V^*(u, \mathcal{I}_A) \neq \emptyset$.

5.2 Existence for Various Reward Functions

In this subsection, we focus on information structures such that for each k , $\cap_{i \in N} I_i(k) = \{k\}$. In this instance, $\hat{K} = \emptyset$ and the reward function u trivially satisfy the genericity condition of Theorem 4.1.⁴

5.2.1 No restriction on rewards: $\mathcal{S} = \mathcal{U}$

The following result identifies the restriction on the information structure that ensures that BFE exists for all reward functions (see also Renault and Tomala, 2004a).

Theorem 5.3 $V^*(\mathcal{I}, u) \neq \emptyset, \forall u \in \mathcal{U}$, if and only if all majority components are singletons.

The proof is straightforward and follows the theorems 3.2. and 3.3. in Renault and Tomala (2004a). The condition is obviously sufficient. If all majority components are singletons, then the true state k can be identified by truthful announcements. Unilateral deviations are disregarded. Then a feasible and individually rational payoff vector in the revealed state k is implemented. For the necessity part we provide an example in the supplemental material (Appendix C).

This condition is obviously very demanding, although BFE might very well exist for a given reward function. The remainder of this section examines how the condition is relaxed once restrictions are imposed on the reward function. Without loss of generality, given Lemma 5.2, we assume hereafter that there is a single majority component, with at least two states (if there is a single state, existence is immediate).

⁴This is without loss of generality when players have known-own payoff. If no such restriction is imposed on rewards, then it is also necessary for non-emptiness of V^* . For example, if each player's reward 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.

5.2.2 Known-own payoffs

In this subsection we provide a condition on the information structure that is necessary to obtain $V^*(\mathcal{I}, u) \neq \emptyset$ in all games of known-own payoff, and a sufficient condition for non-emptiness.

Definition 5.4 *The game has known-own payoffs (KOP) if the reward function 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' :*

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

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

Note that the definition of known-own payoff implies that $\bigcap_{i \in N} I_i(k) = \{k\}$. In two-player games with KOP, existence obtains whenever information is one-sided, that is, whenever player 1 has more information than player 2 (Shalev, 1994). These conditions are also necessary in two-player games: Hörner and Lovo (2009) and Koren (1992) provide examples in which existence fails if information is two-sided. One might then expect that this result might generalize to N -player games with KOP. However, the following example shows that having one fully informed player is not sufficient to ensure existence.

Example 5.5 *There are three states k, k', k'' . The information of player 1 is $I_1(k) = \{k, k''\}$, $I_1(k') = \{k'\}$. The information of player 2 is $I_2(k) = \{k, k'\}$, $I_2(k'') = \{k''\}$. Player 3 knows the state. The payoff matrix is as follows.*

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

state k

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

state k''

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

state k'

In this game, V^* is empty. Assume for the sake of contradiction that there is a point v in V^* . Individual rationality of players 1 and 2 imply that in state k' , T is always played, and (T, R) is played with a (discounted) frequency no greater than $1/4$. The payoff of player 3 in state k' is thus $v_3^{k'} \leq 3/4$. Similarly, in state k'' , R is always played, and (T, R) with frequency no greater than $1/4$. The payoff of player 3 in state k'' is thus $v_3^{k''} \leq 3/4$.

Consider now the inconsistent reports in which player 1 claims that the state is k' , while player 3 claims that the state is k . Continuation play must “punish” player 1 in state k , and player 3 in state k' . Note that, for every action profile a , $u_1^k(a) + u_3^{k'}(a) \geq 3$. Now, assume that the payoff of player 1 in state k is such that: $v_1^k \leq \frac{11}{16}3$. It follows that

$$v_k^1 + v_{k'}^3 \leq \frac{11}{16}3 + 3/4 = 45/16 < 3.$$

This latter inequality is impossible. From JR, there must exist a distribution α of action profiles such that $v_1^k \geq u_1^k(\alpha)$ and $v_3^{k'} \geq u_3^{k'}(\alpha)$ and $u_1^k(\alpha) + u_3^{k'}(\alpha) \geq 3$. We conclude that $v_1^k > \frac{11}{16}3$. A similar argument (considering the inconsistent reports in which player 2 claims that the state is k'' and player 3 claims that the state is k) yields $v_2^k > \frac{11}{16}3$. Thus $v_k^1 + v_k^2 > 66/16 = 4 + 1/8$, which is impossible, since no action profile in state k yields $u_1^k + u_2^k > 4$.

In what follows we show that if V^* is nonempty in all games with KOP, then for each state k , first, there exists a player i who is as well informed as all others at that state, and second, either no player can distinguish any two states for which he is not the best informed player (if he ever is), or there is a second player $j \neq i$ who is as well informed as all players but i at that state. In this latter case, we show that V^* is nonempty in all games with KOP.

More formally, we say that *player i has more information than player j* if player i can deduce player j 's type from his 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 5.6 1. *The information structure is locally weakly embedded (LWE) if for each state k , there exists a pair of players i, j , such that player i has more information than any other player, and player j has more information than any player other than i . Note that i, j may depend on the state⁵.*

2. *The information structure has the all-or-nothing property if there exists a partition of K , $K = \cup_{i=1, \dots, N} K_i$ with K_i possibly empty, such that for each i , $I_i(k) = \{k\}$ if $k \in K_i$, $I_i(k) = K \setminus K_i$ otherwise.*

We have the following result (recall that attention is restricted, without loss of generality, to a single component).

Theorem 5.7 *If $V^*(\mathcal{I}, u) \neq \emptyset$, $\forall u \in S_{\mathcal{I}}$, then the information structure is locally weakly embedded, or has the all-or-nothing property. Further, if the information structure is locally weakly embedded, then $V^*(\mathcal{I}, u) \neq \emptyset$, $\forall u \in S_{\mathcal{I}}$.*

The proof is rather involved and is deferred to Appendix B. In order to prove necessity, we establish a structural result on information structures with a single majority component. This reduces the number of configurations for which counter-examples (in which $V^*(\mathcal{I}, u) = \emptyset$ for some $u \in S_{\mathcal{I}}$) must be provided whenever the information structure is neither LWE nor has the all-or-nothing property. The sufficiency part relies on the following lemma.

Lemma 5.8 *Consider a N -player finite game with actions sets A_i and payoff functions u_i and let $\underline{u}_i = \min_{\alpha_{-i} \in \times_{j \neq i} \Delta A_j} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i})$ be players i 's individual rationality level. There exists*

⁵It is not difficult to check that the pair (i, j) is the same for all states in the same majority component.

α_{-1}^* , such that

$$\forall i \neq 1, \forall \alpha_1, u_i(\alpha_1, \alpha_{-1}^*) \geq \underline{u}_i$$

This states that $N - 1$ players can play cooperatively in order to secure their minmax level, irrespective of the behavior of player 1. A more general statement is proved in Appendix B. With known-own payoffs, one can easily deduce non-emptiness of V^* , if player 1 is informed of the state and the other players have no information. They just have to play such a profile α_{-1}^* , and player 1 takes a best-reply given his actual reward function. If there are two (partially) informed players 1 and 2, we use a sequential construction where player 1 announces his mixed action, allowing the other players to secure their individually rational levels, irrespective of the action of player 2.

Unfortunately, we were unable to prove or disprove existence in the remaining case of information structures satisfying the all-or-nothing property. Countless numerical simulations suggest the following conjecture.

Conjecture 5.9 *The set $V^*(\mathcal{I}, u)$ is non-empty for all $u \in S_{\mathcal{I}}$ if and only if the information structure is locally weakly embedded, or has the all-or-nothing property.*

5.2.3 Bad outcome

In this subsection, we consider a class of reward functions in which there is a distribution of action profiles which yields a low payoff to all players simultaneously. This encompasses many economic settings, e.g. environments with quasi-linear utilities.

Definition 5.10 *The reward function has a bad outcome if there exists a distribution over action profiles that provides each player with no more than his minmax payoff in each state:*

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

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

For each player i and state k , denote by $I_{-i}(k) := \cap_{l \neq i} I_l(k)$ the combined information of the other players at k . We say that player i is *essential* at k if $I_{-i}(k) \neq \{k\}$. The information structure \mathcal{I} has no essential player if, for each state k , no player is essential at k .

Theorem 5.11 $V^*(I, u) \neq \emptyset, \forall u \in \mathcal{B}$, if and only if \mathcal{I} has no essential player.

The proof is straightforward and the intuition is as follows. Let players report their type. Then either a state is identified, or there is an inconsistency in the reports. In that case, the bad outcome is played long enough to deter such deviations. Details are provided in the supplemental material (Appendix D).

5.2.4 Known-own-payoffs and bad outcome

Assuming both known-own-payoffs and bad outcome yields existence for a broader set of information structures.

Theorem 5.12 $V^*(\mathcal{I}, u) \neq \emptyset, \forall u \in \mathcal{S}_{\mathcal{I}} \cap \mathcal{B}$, if and only if \mathcal{I} has at most one essential player in each state.

The proof of this result can be found in the supplemental material (Appendix E).

5.3 Existence for all Information Structures

Our objective is to find conditions on the reward function u such that V^* is non-empty independently of the information structure. Note first that $V^*(\mathcal{I}, u)$ is non-empty for all information structure $\mathcal{I} \in \mathcal{Y}$ if and only if $V^*(\mathcal{I}, u)$ is non-empty 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 our earlier observation 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 non-empty, then $V^*(\mathcal{I}', u)$ is also non-empty. Let

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

Proposition 5.13 *The set $V^*(\mathcal{I}, u)$ is non-empty for all \mathcal{I} if and only if there exists a distribution over action profile $\mu^* \in \Delta A$ such that, for each $i \in N$,*

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

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 for $V^*(\mathcal{I}, u) \neq \emptyset$. 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 $\mu \in \Delta A$ there exists a player i and $q^\mu \in \Delta K$ such that

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

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. \square

The condition of proposition 5.13 is trivially satisfied when it is possible to find a pooling equilibrium distribution μ^* and a punishment strategy that is independent of the state. This is the case, for instance, in most auction formats and oligopoly games (take a very high and a very low price, or quantity).

When focusing on finer information structures in which players have non-degenerate types, punishment strategies sustaining an equilibrium can depend on types. There are some obvious properties of the reward functions ensuring existence, if one gives up the requirement that existence obtains for all information structures. Proposition 5.14 provides a useful criterion, which is the N -player counterpart of condition 4 in Hörner and Lovo (2009). Let \hat{D} be the set of type profiles that are consistent with some state after deletion of some player's type. That is,

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

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

Proposition 5.14 *If there exists a distribution over action profile $\mu^* \in \Delta A$, and for all $\theta \in \hat{D}$, a profile $\mu^\theta \in \Delta A$ such that for all $i, k \in \kappa(\theta_{-i})$,*

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

then V^ is non-empty.*

Proof. It is sufficient to show 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 compatible. IR and JR: the condition on μ^θ implies that when the distribution over action profile μ^θ is implemented, in all possible states a player cannot gain more than v even if he unilaterally deviates or makes a report leading to an inconsistent report profile. Thus, μ^θ can be used to deter unilateral deviations or misreports, guaranteeing that v is individually and jointly rational. \square

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. While there exists a large literature on reputation in two-player games, Fudenberg and Kreps (1987) and Ghosh (2007) are, to the best of our knowledge, the only other papers considering reputations when the informed player faces multiple opponents. In Hörner and Lovo (2009), 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 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 reward is denoted u_1 , while his other possible payoff types are denoted u_1^k , $k = 2, \dots, K$. We fix throughout the reward functions (u_2, \dots, u_N) of players $i = 2, \dots, N$. Given some reward 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 $u^K := (u_1^2, \dots, u_1^K)$ 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

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

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

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

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

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

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

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

where $\mathbf{1}$ 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 $(p_2, \dots, p_N) = 0$). The following theorem shows that this lower bound is actually achieved, and provides an alternative characterization of it. The proof of it can be found in the supplemental material (Appendix F).

Theorem 6.1 *The reputation payoff is equal to*

$$u_1^* = \sup_{\{p_i \geq 0: i=2, \dots, N\}} \text{val} \left(u_1 - \sum_{i=2}^N p_i (u_i - \underline{u}_i \mathbf{1}) \right) = \sup_{\alpha_1 \in \Delta A_1} \min_{\alpha_{-1} \in Y(\alpha_1)} u_1(\alpha_1, \alpha_{-1}),$$

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

$$u_1^* = v_1(-u_2, \dots, -u_N).^6$$

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

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

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

⁶Note that zero-sum games violate the interiority assumption. However, as in Hörner and Lovo (2009, online appendix), it is straightforward to approach this reputation payoff by considering payoff matrices satisfying the interiority assumption, which are arbitrarily close to the zero-sum game.

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's rational type can secure the Stackelberg payoff in all games of conflicting interest.

7 Conclusion

This paper provides a characterization of the set of belief-free equilibrium payoffs in games with perfect monitoring. Further, necessary and sufficient conditions on the information structure are identified for non-emptiness of this set.

As discussed, belief-free equilibria have appealing properties. However, because they do not rely on beliefs, they are silent on how beliefs actually shape play. Game theory has played an important role in providing insights about when and how agents learn, whether it is advantageous to hide or disclose private information, or how fast to reveal it. This provides a useful perspective on the existence or non-existence results of belief-free equilibria. In an environment in which such equilibria do not exist, play must necessarily reflect beliefs, and this opens the door for robust findings on this dependence. This is the case, for instance, in zero-sum games with incomplete information on one-side, in which the speed of convergence can be determined (Mertens, 1998). On the other hand, if one attempts to address such issues in an environment in which belief-free equilibria exist, it becomes more important to stress why the choice of the particular equilibrium is compelling. This could be, for instance, because the equilibrium that is considered is efficient (see, however, the folk theorems established by Fudenberg and Yamamoto, 1999b). Alternatively, one must invoke considerations that are external to the repeated game, such as those involving measures of complexity, for instance.

References

- [1] Aumann R. J., 1976. “Agreeing to Disagree,” *The Annals of Statistics*, **4**, 1236–1239.
- [2] Aumann R. J. and M. B. Maschler, 1995. *Repeated Games with Incomplete Information*. The MIT Press.
- [3] Bergemann, D. and S. Morris, 2007. “Belief Free Incomplete Information Games,” Cowles Foundation Discussion Paper No. 1629, Yale University.
- [4] Blackwell, D., 1956. “An Analog of the Minmax Theorem for Vector Payoffs,” *Pacific Journal of Mathematics*, **6**, 1–8.
- [5] Chassang, S. and S. Takahashi, 2009. “Robustness to Incomplete Information in Repeated Games,” working paper, Princeton University.
- [6] Crémer, J. and R. McLean, 1985. “Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent,” *Econometrica*, **53**, 345–361.
- [7] 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.
- [8] Ely, J. and J. Välimäki, 2002. “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, **102**, 84–105.
- [9] Ely, J., J. Hörner and W. Olszewski, 2005. “Belief-free Equilibria in Repeated Games,” *Econometrica*, **73**, 377–415.

- [10] 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.
- [11] Forges, F. and E. Minelli, 1997. “A Property of Nash Equilibria in Repeated Games with Incomplete Information,” *Games and Economic Behavior*, **18**, 159–175.
- [12] Fudenberg, D. and D. Kreps, 1987. “Reputation in the Simultaneous Play of Multiple Opponents,” *Review of Economic Studies*, **54**, 541–568.
- [13] Fudenberg, D. and D. Levine, 1989. “Reputation and Equilibrium Selection in Games with a Single Patient Player,” *Econometrica*, **57**, 759–778.
- [14] Fudenberg, D. and E. Maskin, 1986. “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, **54**, 533–554.
- [15] Fudenberg, D. and E. Maskin, 1991. “On the Dispensability of Public Randomization in Discounted Repeated Games,” *Journal of Economic Theory*, **53**, 428–438.
- [16] Fudenberg, D. and Y. Yamamoto, 2009a. “Perfect Public Ex-Post Equilibria of Repeated Games with Uncertain Outcomes,” working paper, Harvard University.
- [17] Fudenberg, D. and Y. Yamamoto, 2009b. “Learning from Private Information in Noisy Repeated Games,” working paper, Harvard University.
- [18] Ghosh, S., 2007. “Multiple Opponents and the Limits of Reputation,” working paper, Boston University.
- [19] Harsanyi, J. C., 1967–1968. “Games with incomplete information played by Bayesian players,” *Management Science*, **14**, 159–182, 320–334, 486–502.

- [20] Hart, S., 1985. “Nonzero-sum Two-person Repeated Games with Incomplete Information,” *Mathematics of Operations Research*, **10**, 117–153.
- [21] Hart, S. and D. Schmeidler, 1989. “Existence of correlated equilibria,” *Mathematics of Operations Research*, **14**, 18—25.
- [22] Hörner, J. and S. Lovo, 2009. “Belief-free Equilibria in Games with Incomplete Information,” *Econometrica*, **77**, 453–487.
- [23] Hörner, J., S. Lovo and T. Tomala, 2009. “Belief-free equilibria in games with incomplete information: Characterization and Existence,” Cowles Discussion Paper, Yale.
- [24] Israeli, E., 1999. “Sowing Doubt Optimally in Two-Person Repeated games,” *Games and Economic Behavior*, **28**, 203–216.
- [25] Kalai, E., 2004. “Large Robust Games,” *Econometrica*, **72**, 1631–1665.
- [26] Kohlberg, E., 1975. “Optimal Strategies in Repeated Games with Incomplete Information,” *International Journal of Game Theory*, **4**, 7–24.
- [27] Koren, G., 1992. “Two-Person Repeated Games where Players Know Their Own Payoffs,” working paper, New York University.
- [28] Mertens, J.-F., 1998. “The speed of convergence in repeated games with incomplete information on one side,” *International Journal of Game Theory*, **27**, 343–357.
- [29] Miller, D. A., 2009. “Optimal ex post incentive compatible equilibria in repeated games of private information,” working paper, U.C. San Diego.

- [30] Monderer, D. and M. Tennenholtz, 1999. "Dynamic Non-Bayesian Decision Making in Multi-Agent Systems," *Annals of Mathematics and Artificial Intelligence*, **25**, 91–106.
- [31] Peški, M., 2008. "Repeated games with incomplete information on one side," *Theoretical Economics*, **3**, 29–84.
- [32] Piccione, M., 2002. "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, **102**, 70–83.
- [33] Renault, J., 2001. "3-player repeated games with lack of information on one side," *Games and Economic Behavior*, **30**, 221–245.
- [34] Renault, J., and T. Tomala, 2004a. "Learning the State of Nature in Repeated Games with Incomplete Information and Signals," *Games and Economic Behavior*, **47**, 124–156.
- [35] Renault, J., and T. Tomala, 2004b. "Communication equilibrium payoffs of repeated games with imperfect monitoring," *Games and Economic Behavior*, **49**, 313–344.
- [36] Shalev, J., 1994. "Nonzero-sum two-person Repeated Games with Incomplete Information and Known-own Payoffs," *Games and Economic Behavior*, **7**, 246–259.
- [37] Simon, R. S., S. Spież and H. Toruńczyk, 1995. "The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk-Ulam type," *Israel Journal of Mathematics*, **92**, 1–21.
- [38] Sorin, S, 1983. "Some results on the existence of Nash equilibria for non-zero sum games with incomplete information," *International Journal of Game Theory*, **12**, 193–205.
- [39] Sorin, S., 1986. "On Repeated Games with Complete Information," *Mathematics of Operations Research*, **11**, 147–160.

- [40] Thomas, J. P., 1995. "Subgame-perfect attainment of minimax punishments in discounted two-person games," *Economics Letters*, **47**, 1–4.
- [41] Wiseman, T., 2008. "A Partial Folk Theorem for Games with Private Learning," working paper, U.T. Austin.

APPENDIX A: PROOF OF THEOREM 4.1 WITH A COMMUNICATION DEVICE

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

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

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

for all i , all $\varepsilon_i \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, all $\varepsilon'_i \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, all (θ_i, θ_{-i}) and (θ'_i, θ_{-i}) such that $\kappa(\theta_i, \theta_{-i}) \neq \emptyset$ and $\kappa(\theta'_i, \theta_{-i}) \neq \emptyset$. Such a distribution exists for sufficiently small $\bar{\varepsilon} > 0$ given that $v \in \text{int } V^*$ 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 $\kappa(\theta) = \emptyset$. We distinguish two cases.

1. $\theta \in D$: by definition, there exist a set Ω_θ of players and types (i, θ'_i) such that $\kappa(\theta'_i, \theta_{-i}) \neq \emptyset$. Action profiles are selected according to a probability distribution $\mu_\theta(\varepsilon)$ in such a way that

$$u_i(k, \mu_\theta(\varepsilon)) < v_i^k + \varepsilon_i \quad (3)$$

for all $(i, \theta'_i) \in \Omega_\theta$, $k \in \kappa(\theta'_i, \theta_{-i})$ and all $\varepsilon_i \in [-\bar{\varepsilon}, \bar{\varepsilon}]$. Such a distribution exists for sufficiently small $\bar{\varepsilon} > 0$ given that $v \in \text{int } V^*$ 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.⁷ Player i plays $s_i^{\underline{a}_i}$ throughout the phase.

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

Initial phase

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

⁷To 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.

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

(b) $\kappa(\theta'_{-i}) = \emptyset$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_j = -\bar{\varepsilon}$ 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) $\kappa(\theta') \neq \emptyset$:

i. $\theta' = (\theta_{-i}, \theta'_i)$ for some $i \in N$ and $\theta'_i \neq \theta_i$: the next phase is $R^{\theta'}(-\bar{\varepsilon}, \varepsilon_{-i})$;

ii. otherwise, the next phase is $R^\theta(\varepsilon)$;

(b) $\kappa(\theta') = \emptyset$: the next phase is $E^{\theta'}(\varepsilon')$, where $\varepsilon'_i = -\bar{\varepsilon}$ 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) $\kappa(\theta'_{-i}) \neq \emptyset$: the next phase is $P^{\theta'_{-i}}$;

- (b) $\kappa(\theta'_{-i}) = \emptyset$: the next phase is $E^{\theta'}$ (ε'), where $\varepsilon'_j = -\bar{\varepsilon}$ 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) $\kappa(\theta') \neq \emptyset$: the next phase is $R^\theta(\varepsilon)$;
- (b) $\kappa(\theta') = \emptyset$: the next phase is $E^{\theta'}$ (ε'), where $\varepsilon'_i = -\bar{\varepsilon}$ 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) $\kappa(\theta') = \emptyset$: the next phase is $E^{\theta'}$ (ε'), where $\varepsilon'_i = -\bar{\varepsilon}$ if $(i, \theta''_i) \in \Omega_{\theta'}$ for some $\theta''_i \in \Theta_i$, and $\varepsilon'_i = \varepsilon_i$ otherwise;
- (b) $\kappa(\theta') \neq \emptyset$: 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 $(1 - \delta^T)(v_i^k - 2\bar{\varepsilon}) + \delta^T(v_i^k - \bar{\varepsilon})$, 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 $\varepsilon_j(\cdot; P^{\theta_{-i}})$

is in $[\bar{\varepsilon}/3, \bar{\varepsilon}]$ if $\theta'_j = \theta_j$, and it is in $[-\bar{\varepsilon}, -\bar{\varepsilon}/3]$ 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 $v \in \text{int } V^*$, we now pick $\bar{\varepsilon} > 0$ small to ensure that the probability distributions introduced above exist, and $\bar{\delta}$, 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

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

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

Given v and $\bar{\varepsilon} > 0$, these are all satisfied as $\delta^T \rightarrow 1$ and $T \rightarrow \infty$, so they are also satisfied for values of T and δ that are large enough. Inequality (6) guarantees that a variation of $2\bar{\varepsilon}/3$ 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 $v_i^k - \bar{\varepsilon}$, where $k \in \kappa(\theta')$. If he misreports, we further distinguish two cases:

- (a) $\kappa(\theta') = \emptyset$: assuming the other players report truthfully, the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\bar{\varepsilon}$. So player i 's payoff is at most $\max_{\theta'_i \neq \theta_i} (1 - \delta) u_i(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)) + \delta (v_i^k - \bar{\varepsilon})$, which is less than $v_i^k - \bar{\varepsilon}$, because of (3).
- (b) $\kappa(\theta') \neq \emptyset$: Player i gets at most $\max_{\theta'_i \neq \theta_i} (1 - \delta) u_i(k, \mu_{\theta'_i, \theta'_{-i}}(\varepsilon)) + \delta (v_i^k - \bar{\varepsilon})$, which is less than $(v_i^k - \bar{\varepsilon})$, 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 $-(1 - \delta^T)M + \delta^T(v_i^k + \bar{\varepsilon}/3)$. If he misreports, there are two cases:

- (a) $\kappa(\theta') = \emptyset$: the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\bar{\varepsilon}$, so his payoff is smaller than $(1 - \delta)M + \delta(v_i^k - \bar{\varepsilon}) < (1 - \delta^T)M + \delta^T(v_i^k - \bar{\varepsilon}/3)$, which is less than $-(1 - \delta^T)M + \delta^T(v_j^k + \bar{\varepsilon}/3)$ because of (6).
- (b) $\kappa(\theta'_i, \theta_{-i}) \neq \emptyset$: Player i gets at most $(1 - \delta^T)M + \delta^T(v_i^k - \bar{\varepsilon}/3)$ (assuming he reports truthfully at the end), which is less than $-(1 - \delta^T)M + \delta^T(v_i^k + \bar{\varepsilon}/3)$ 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 $\varepsilon_i(h; P^{\theta-i})$ such that $s_i^{\underline{a}_i}$ 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 $\varepsilon_j(h; P^{\theta-i})$ 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 $v_i^k - \bar{\varepsilon}$. If he misreports, we distinguish two cases:

1. $\kappa(\theta') = \emptyset$: the next phase is $E^{\theta'}(\varepsilon')$ with $\varepsilon'_i = -\bar{\varepsilon}$, so player i 's payoff is at most $\max_{\theta'_i \neq \theta_i} (1 - \delta) u_i(k, \mu_{\theta'_i, \theta'_-i}(\varepsilon)) + \delta (v_i^k - \bar{\varepsilon})$, which is less than $v_i^k - \bar{\varepsilon}$ because of (3).
2. $\kappa(\theta') \neq \emptyset$: player i gets at most $\max_{\theta'_i \neq \theta_i} (1 - \delta) u_i(k, \mu_{\theta'_i, \theta'_-i}(\varepsilon)) + \delta (v_i^k - \bar{\varepsilon})$, which is less than $v_i^k - \bar{\varepsilon}$ because of (2).

□

APPENDIX A.1: PROOF OF THEOREM 4.1 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 since they 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 (2009).

Because communication requires several periods, strategies must also specify how a player plays within a communication phase if his own previous action, or his opponent's previous action already precludes him from reporting correctly his private information. The construction must ensure that continuation strategies remain optimal for all states after such histories, 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.⁸) A communication phase lasts c periods, where

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

⁸It cannot be entirely ignored, since we must give i incentives that do not depend on his type.

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

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

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

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

denote a *report*, or *message* profile. For $k \in K$, let $u_i^C(k, \bar{\theta})$ denote player i 's average payoff from the communication phase if the state is k and the report is $\bar{\theta}$.⁹

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 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 $R(\bar{\theta}, \varepsilon)$, where $\kappa(\bar{\theta}) \neq \emptyset$, and $\varepsilon \in [-\bar{\varepsilon}, \bar{\varepsilon}]^N$, for some $\bar{\varepsilon} > 0$ to be specified.

A regular phase lasts at most n periods (to be specified), where $n > c$. We fix a (possibly

⁹This 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.

correlated) mixed action profile $\mu(\bar{\theta}, \varepsilon) \in \Delta A$ such that, $\forall k \in \kappa(\bar{\theta}), \forall i \in N, \forall \varepsilon, \varepsilon' \in [-\bar{\varepsilon}, \bar{\varepsilon}]^N$ and $\forall \theta'_i \in \Theta_i, \theta'_i \neq \bar{\theta}_i$, such that $\kappa(\theta'_i, \bar{\theta}_{-i}) \neq \emptyset$,

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

and

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

and

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

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

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

Transitions are described below.

Penitence Phase

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

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

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

Transitions are described below.

Punishment Phase

A *punishment phase*, indexed by i , is denoted $P_i(\bar{\theta}_{-i}, t)$, where $\bar{\theta}_{-i} \in \Theta_{-i}$ is such that $\kappa(\bar{\theta}_{-i}) \neq \emptyset$ 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 $s_{-i}^{\bar{\theta}_{-i}}$.

We pick $n, T, \bar{\delta} < 1$ and $\bar{\varepsilon}$ such that, $\forall \delta > \bar{\delta}, \forall k \in \kappa(\bar{\theta}_{-i})$, 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 $\bar{\theta}$, define

- whenever $\bar{\theta} \in \Theta, \forall \theta \in \Theta, \Delta_I(\theta, \bar{\theta}) := \{i \in N | \theta_i \neq \bar{\theta}_i\}$;
- whenever $\bar{\theta} \in \Theta, \bar{\theta} \in D, \Delta_D(\bar{\theta}) := \{i \in N | (i, \theta'_i) \in \Omega_{\bar{\theta}} \text{ for some } \theta'_i \in \Theta_i\}$;
- whenever $\bar{\theta} \notin \Theta, \Delta_U(\bar{\theta}) := \{i \in N | \bar{\theta}_i \notin \Theta_i\}$.

Given a unilateral deviation from a sequence $a(\bar{\theta}, \varepsilon)$, or from a mixed action $\mu(\bar{\theta}, \varepsilon)$, let Δ_A denote the index of the player who deviated.¹⁰ Finally, given a set $\Delta \subset N$, let $-\Delta := N \setminus \Delta$.

¹⁰Recall 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.

From a communication phase

The transition depends on the message $\bar{\theta}$ 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 $\bar{\theta} \in \Theta$, a regular or a penitence phase follows, while if $\bar{\theta} \notin \Theta$, either a punishment or a communication phase follows. If there is a unilateral deviation during Φ by player i , then if $\bar{\theta}_{-i} \in \Theta_{-i}$, 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

1. if $\bar{\theta}_{-i} \in \Theta_{-i}$, $\kappa(\bar{\theta}_{-i}) \neq \emptyset$: $P_i(\bar{\theta}_{-i}, T)$;
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 $\bar{\theta} \in \Theta$, $\kappa(\bar{\theta}) \neq \emptyset$: $R(\bar{\theta}, \varepsilon_{-\Delta_I(\theta, \bar{\theta})}, -\bar{\varepsilon}_{\Delta_I(\theta, \bar{\theta})})$;
- (b) if $\bar{\theta} \in \Theta$, $\bar{\theta} \in D$: $E(\bar{\theta}, \varepsilon_{-\Delta_D(\bar{\theta})}, -\bar{\varepsilon}_{\Delta_D(\bar{\theta})})$;
- (c) if $\Delta_U(\bar{\theta}) = \{i\}$, $\kappa(\bar{\theta}_{-i}) \neq \emptyset$: $P_i(\bar{\theta}_{-i}, n)$;
- (d) otherwise, C ;

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

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

punishment phase, is equal to, for all $k \in \kappa(\bar{\theta})$,

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

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

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

(c) otherwise, C .

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

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

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

and if $\Phi = C_i$,

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

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

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

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

(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 $\theta = \bar{\theta}$, and $\varepsilon \in [-\bar{\varepsilon}, \bar{\varepsilon}]$ 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 θ_{-i} reported by the other players after the deviation and before the punishment phase to be correct; since his payoff at the beginning of the punishment phase is

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

then 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.¹¹

Consider next a punishment phase P_i . The definition of $\tilde{\varepsilon}_i$ guarantees that $s_i^{\frac{\alpha_i}{i}}$ is optimal for

¹¹Note 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.

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 .

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 $(m_1(\theta_1), \dots, m_N(\theta_N))$), 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, he 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 than $v_i - \bar{\varepsilon}$, 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 value of $\varepsilon_i \geq -\bar{\varepsilon}/4$ 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 $-\bar{\varepsilon}/4$; 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, $\varepsilon_i \geq -\bar{\varepsilon}/4$; if he reports differently, in the regular phase that will follow the second communication phase, $\varepsilon_i = -\bar{\varepsilon}$.

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 $\varepsilon_i = -\bar{\varepsilon} + \rho n_U$ in the regular phase that will follow the new communication C .

APPENDIX B: PROOF OF THEOREM 5.7

Necessity

We prove that if an information structure has a single majority component and is neither LWE nor has the all-or-nothing property, then there is a reward function (which satisfies KOP) such that V^* is empty. First observe that it is sufficient to focus on four information structures.

Proposition 7.1 *If an information structure has a single majority component and is neither LWE nor has the all-or-nothing property, then there there is a subset of three states such that the restriction of \mathcal{I} to this subset is of one of the following four types*

		k_1	k_2	k_3
1	k_1	*	*	
2	*	k_2	*	

A

		k_1	k_2	k_3
1	k_1	*	*	
2	*	k_2	*	
3	k_1	k_2	k_3	

B

		k_1	k_2	k_3
1	k_1	*	*	
2	*	k_2	*	
3	k_1	*	*	

C

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

D

where the entries are the types, or signals, of the players and it is understood that other players have no information on those states.

The proof is relegated to supplemental material (Appendix G).

Counter-examples

For each information structure, we present a counter-example, i.e. a reward function for which $V^*(u, \mathcal{I}) = \emptyset$.

A: a two-sided battle of the sexes We start by a counter-example due to Koren (1992), see also Hörner and Lovo (2009). There are three states k, k', k'' . The information of player 1 is $I_1(k) = \{k, k''\}$, $I_1(k') = \{k'\}$. The information of player 2 is $I_2(k) = \{k, k'\}$, $I_2(k'') = \{k''\}$. 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
	state k''	

	L	R
T	1, 1	1, 0
B	0, 0	0, 3
	state k'	

The proof that $V^* = \emptyset$ for this game is in Hörner and Lovo (2009). The main argument is the following. In state k' , player 1 has a dominant strategy, and individual rationality requires T to be played with frequency 1 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 frequency at least $3/4$ in state k . A symmetric argument for player 2 shows that (B, R) must be played with frequency at least $3/4$ in state k . These two requirements are mutually incompatible.

B: Adding a fully informed player Consider example 5.5. This corresponds to the previous game with the addition of a third player, player 3, who knows the state.

C: Adding a partially informed player Consider the game of case A again, and assume that there is a player 3 who has the same information as player 2. The payoff of player 3 does not depend on the state and is:

	L	R
T	3	$3 - \varepsilon$
B	3	0

u_3

In this game, V^* is empty. Assume for the sake of contradiction that there is a point v in V^* . Individual rationality of players 1 and 2 implies that in state k' , T is played with frequency 1, and (T, R) with frequency no more than $1/4$. Then, since player 1 is the only player to distinguish k and k' , incentive compatibility requires that the payoff v_1^k of player 1 in state k satisfies: $v_1^k \geq 3 \times \frac{3}{4}$. Since the sum of players 1 and 2's payoffs in state k is at most 4, this implies $v_2^k \leq \frac{7}{4}$. Individual rationality of players 1 and 2 also implies that in state k'' , R is played with frequency 1, and (T, R) with frequency no more than $1/4$. This implies that the payoff of player 3 in state k'' is such that: $v_3^{k''} \leq (3 - \varepsilon)/4$.

Consider now the following inconsistent reports: player 2 claims that the state is k'' and player 3 claims that the state is k . Joint rationality requires that there exists a distribution α of action profiles such that $v_2^k \geq u_2^k(\alpha)$ and $v_3^{k''} \geq u_3^{k''}(\alpha)$. This is impossible, because $v_2^k + v_3^{k''} \leq 7/4 + (3 - \varepsilon)/4 < 3 - \varepsilon$ for ε small and since for every action profile, $u_2^k + u_3^{k''} \geq 3 - \varepsilon$.

D: Adding two partially informed player Consider once again the game of case A, and assume that there is a third player, player 3, who has the same information as player 2, and a

fourth player, player 4, who has the same information as player 1. The payoff of player 3 is as in case C. The payoff of player 4 does not depend on the state and is:

	L	R
T	0	$3 - \varepsilon$
B	3	3

u_4

In this game, V^* is empty. Assume for the sake of contradiction that there is a point v in V^* . As in the previous example, individual rationality of players 1 and 2 in state k'' implies $v_3^{k''} \leq (3 - \varepsilon)/4$. Consider again the inconsistent reports in which player 2 claims that the state is k'' , while player 3 claims that the state is k . Since for every action profile $u_2^k + u_3 \geq 3 - \varepsilon$, joint rationality implies $v_2^k + v_3^{k''} \geq 3 - \varepsilon$ and thus $v_2^k \geq (3 - \varepsilon)3/4$.

By a symmetric argument, considering the inconsistent reports in which player 1 claims that the state is k' and player 4 claims that the state is k , we find $v_1^k \geq (3 - \varepsilon)3/4$. This implies that $v_1^k + v_2^k \geq (3 - \varepsilon)3/2 > 4$ for small ε , which is impossible.

The following Proposition show that an information structure with a single majority component, that is neither LWE nor has the all-or-nothing property is necessary of one of type A, B, C or D.

LWE is sufficient

In this part we show that if the information structure is locally weakly embedded, then $V^*(\mathcal{I}, u) \neq \emptyset, \forall u \in S_{\mathcal{I}}$. Note that if there is a single majority component, LWE implies that there exists two players 1,2 and a partition of the set of states $K = K_1 \cup K_2$ such that:

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

- $I_1(k) = \{k\}$ for each $k \in K_1$.
- $I_2(k) = \{k\}$ for each $k \in K_2$.

We first prove the result assuming $K_1 = K$, i.e. player 1 is fully informed at each state.

Proposition 7.2 *Consider an information structure such that: player 1 knows the state and players $3, \dots, n$ have no information (i.e. $\forall k, I_1(k) = \{k\}, I_3(k) = \dots = I_N(k) = K$). Then $V^*(\mathcal{I}, u) \neq \emptyset, \forall u \in S_{\mathcal{I}}$.*

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

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

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

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

The set \mathcal{A}_θ is clearly compact.

Claim 7.3 \mathcal{A}_θ is non-empty.

Proof. We fix α_1 . For $i \geq 3$ consider the correspondence $F_i(\alpha_1, \cdot) : \times_{j \notin \{1, 2, i\}} \Delta A_j \rightarrow \Delta A_i$ 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:

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

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.

Let us denote $BR_{2,\theta}$ the best-reply correspondence of player 2 of type θ and $BR_{1,f}$ is the best-reply correspondence of player 1 when his payoff function is $f : A \rightarrow \mathbb{R}$. Consider the correspondence $\Phi_{f,\theta}$ from $\times_i \Delta A_i$ to itself defined as follows:

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

$\Phi_{f,\theta}$ has non-empty, convex and compact values and it is straightforward to check that it has a closed graph. It admits thus a fixed point $\bar{\alpha}$ by Kakutani's fixed point theorem. Clearly, $\bar{\alpha}$ is in \mathcal{A}_θ . Note that this profile has the additional property to be 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 property is used later on. This ends the proof of the Claim.

Note that this proves Lemma 5.8 by considering the special case where player 1 has a single action.

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^* . To phrase this definition, the informed player announces an action profile for all players but player 2, who takes a best-reply. Player 1 may choose the profile as she wishes, provided that it secures the minmax level of players 3, \dots , N , irrespective of the action of player 2.

Under the assumptions of Proposition 7.2, the constraints defining V^* are the following:

- *Individual rationality for player 1.*

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

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

- *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:

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

Let us check all these points.

Individual rationality for player 1. Fix θ and $q \in \Delta\theta$. It follows from our construction that:

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

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:

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

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, ..., n holds by construction of \mathcal{A}_θ . 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 pretends that the state is k' with $I_2(k') = \theta$, the induced action profile $\alpha^{k'}$ also belongs to \mathcal{A}_θ , and the resulting payoff for player 1 is at most $\max_{\alpha \in \mathcal{A}_\theta} u_1(k, \alpha)$.

Joint rationality. Suppose as above, that the true state is $k \in \theta$, but player 1 pretends that the true state is k' with $\theta' = I_2(k') \neq \theta$. Still, since the action of player 2 is dictated by her true type θ , so the induced action profile belongs to \mathcal{A}_θ . 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 $\alpha_{-2}^{k'}$ and the best-reply of player 2 of type θ' is $\alpha_{-2}^{k'}$ by construction. Player 2 has thus no incentive to misreport. This ends the proof of the proposition. \square

Now, let us take up the general LWE case where there is a partition of the set of states $K = K_1 \cup K_2$ such that: $I_l(k) = K$ for each k and each $l \neq 1, 2$, $I_1(k) = \{k\}$ for each $k \in K_1$, $I_2(k) = \{k\}$ for each $k \in K_2$.

Proposition 7.4 *Under KOP, then $V^*(\mathcal{I}, u)$ is non-empty for the above information structure.*

Proof. Let take $\{K_1, K_2\}$ be the partition of K such that K_1 is the largest subset of K satisfying $I_1(k) = \{k\}$ for each $k \in K_1$. Let $L := \{L_1, \dots, L_M\}$ be the partition over K_2 that is induced by player 1 information. Then we have,

- $\cup_m L_m = K_2$
- If $k, k' \in L_m$, then $I_1(k) = I_1(k')$
- If $k, k' \in L_m$, then $u_1(k, \alpha) = u_1(k', \alpha)$
- $\forall L_m \in L$, if $k \in L_m$, then there exists $k' \neq k, k' \in L_m$

Consider first the game Γ^{K_1} obtained by eliminating the states in K_2 . The only possible states for this game are those in K_1 and player 1 is the best informed player. This game has a non-empty V^* from Proposition 7.2. Consider the above construction for this restricted game and denote α_i^k , $i = 1, \dots, N$, $k \in K_1$ player i 's corresponding mixed action in state k . Note that α_i^k , $i \neq 2$ are such that when all players $i \neq 2$ play α_i^k , then player $j \geq 3$ payoff is individually rational independently of player 2's strategy.

Consider now $\underline{\Gamma}^{K_2}$ defined as the game where the only possible states are those in K_2 and the payoff function of player 2 in state $l \in K_2$ is $\hat{u}_2(l, \alpha) := -u_1(l, \alpha)$ while the payoff functions of all the other players are as in the original game. This is a game of known-own payoff where player 1 knows at least as much as player 2, i.e. player 1 is fully informed. Let α_i^l , $i = 1, \dots, N$, $l \in K_2$ be player i 's mixed action obtained by using the previous construction in $\underline{\Gamma}^{K_2}$ (player 1 announcing all players but player 2 stage game strategies and player 2 best replying). Note that if $l, l' \in L_m \subseteq K_2$, then $\alpha_i^l = \alpha_i^{l'}$. Also for this game α_i^l , $i \neq 2$ are such that when all players $i \neq 2$ play α_i^l , then player $j \geq 3$ payoff is individually rational independently of player 2 strategy. In addition player 1 incentive compatibility constraint implies that

$$\min_{\alpha_2} u_1(l, \alpha_{-2}^l, \alpha_2) \geq \min_{\alpha_2} u_1(l, \alpha_{-2}^{l'}, \alpha_2)$$

for any pair $l, l' \in K_2$. Note also that for any $l \in K_2$

$$\max_{k \in K_1} \min_{\alpha_2} u_1(l, \alpha_{-2}^k, \alpha_2) \leq \min_{\alpha_2} u_1(l, \alpha_{-2}^l, \alpha_2)$$

Otherwise, in game $\underline{\Gamma}^{K^2}$, player 1 would have chosen for state l a strategy α_{-2}^k for some $k \in K_1$ instead of α_{-2}^l .

Take a state $k \in K$ and a strategy profile α_{-2} for all players but 2 and choose

$$\beta_2^k(\alpha_{-2}) \in \arg \min_{\alpha_2} u_1(k, \alpha_{-2}, \alpha_2),$$

$$br_2^k(\alpha_{-2}) \in \arg \max_{\alpha_2} u_2(k, \alpha_{-2}, \alpha_2),$$

Consider the following construction:

- *Step 1: Types announcement.* Players 1 and 2 announce their types. Let θ_1 and θ_2 be their announcement.
- *Step 2.a: Regular Play.* If there is no contradiction in player 1 and 2 announcements, then
 - If the state $k \in K_1$ is announced, then players other than 2 play α_{-2}^k and player 2 takes a best reply.
 - If the state $l \in K_2$ is announced, then players other than 2 play α_{-2}^l and player 2 takes a best reply.

- *Step 2.b: Penitence Play.* If there is a contradicting announcement, then
 - If the announcement is $\theta_1 = k \in K_1$ and $\theta_2 = l \in K_2$, then players other than 2 play α_{-2}^k and player 2 plays $\beta_2^l(\alpha_{-2}^k)$.
 - If the announcement is $\theta_1 = L_m \in L$ and $\theta_2 = l' \notin L_m$, then players other than 2 play $\alpha_{-2}^{l'}$, with $l \in L_m$, and player 2 plays $\beta_2^{l'}(\alpha_{-2}^{l'})$.
 - If the announcement is $\theta_1 = L_m \in L$ and $\theta_2 \in K_1$, then players other than 2 play α_{-2}^l with $l \in L_m$, and player 2 takes a best-reply for the announced type θ_2 .

The interpretation is the following. As in the proof of Proposition 7.2, player 1 chooses an action profile α_2 that secures the minmax levels of the uninformed players (irrespective of the action of player 2), and player 2 takes a best-reply. On K_1 , the construction is essentially unchanged. On K_2 , player 1 makes this choice, expecting player 2 to be adversarial ($\hat{u}_2(l, \alpha) := -u_1(l, \alpha)$). If player 1 misreports, player 2 makes this expectation happen and minimizes the payoff of player 1. Otherwise, she takes a best reply according to her actual payoff function. This gives an incentive to player 1, who prefers player 2 to take her best reply rather than punishing him.

*Verification that the strategy define a payoff vector in V^**

Individual rationality. Consider the payoffs from the regular play. The payoff is clearly IR for player $i \geq 3$ from the construction. Player 2 payoff is IR since she is best replying to the other players strategies. In states $k \in K_1$, player 1's IR follows from the construction of the equilibrium of the game Γ^{K_1} . In a state $l \in K_2$, player 1's payoff is at least as large as the one in the game $\underline{\Gamma}^{K_2}$ which is individually rational. I.e., for $l \in K_2$ we have,

$$v_1^l = u_1(l, \alpha_{-2}^l, br_2^l(\alpha_{-2}^l)) \geq \min_{\alpha_2} u_1(l, \alpha_{-2}^l, \alpha_2) \geq \underline{u}_1^l.$$

Incentive compatibility. In states K_1 , incentive compatibility for player 1 follows the equilibrium of Γ^{K_1} . For the states l which belong to some $L_m \in L$, player 2's takes a best-reply to $(\alpha_{-2}^{L_m})$ which does not depend on $l \in L_m$, and incentive compatibility follows.

Joint Rationality. There are three possible types of contradicting announcement: First, $\theta_1 = k \in K_1$ and $\theta_2 = l \in K_2$; second $\theta_1 = L'_m \in L$ and $\theta_2 = l \notin L_m$; and third, $\theta_1 = L'_m \in L$ and $\theta_2 \in K_1$.

If the state is $k \in K_1$, player 2 has no incentive of announcing $\theta_2 = l \in K_2$ as by doing so she punishes player 1 instead of taking her best reply. That is,

$$u_2(k, \alpha_{-2}^k, \beta_2^l(\alpha_{-2}^k)) \leq u_2(k, \alpha_{-2}^k, br_2^k(\alpha_{-2}^k)) = v_2^k.$$

Similarly if player 1 announces $\theta_1 = k \in K$ when the state is $l \in L_m \subseteq K_2$, this triggers a punishment by player 2 and player 1's payoff is at most,

$$\max_{k \in K_1} u_1(l, \alpha_{-2}^k, \beta_2^l(\alpha_{-2}^k)) = \max_{k \in K_1} \min_{\alpha_2} u_1(l, \alpha_{-2}^k, \alpha_2) \leq$$

$$\leq \min_{\alpha_2} u_1(l, \alpha_{-2}^l, \alpha_2) \leq u_1(l, \alpha_{-2}^l, br_2^l(\alpha_{-2}^l)) = v_1^l$$

Consider now the announcement $\theta_1 = L_m$ and $\theta_2 = l' \in K_2$ with $l' \notin L_m$. If the state is $l \in L_m$, then player 2 has no incentive to announce l' as this would trigger a punishment yielding

$$u_2(l, \alpha_{-2}^l, \beta_2^{l'}(\alpha_{-2}^l)) \leq u_2(l, \alpha_{-2}^l, br_2^{l'}(\alpha_{-2}^l)) = v_2^{l'}.$$

If the state is $l' \notin L_m$, then player 1 has no incentive to announce L_m as this would trigger a punishment yielding at most

$$\begin{aligned} \max_{l' \in K_2} u_1 \left(l', \alpha_{-2}^l, \beta_2^{l'}(\alpha_{-2}^l) \right) &= \max_{l' \in K_2} \min_{\alpha_2} u_1 \left(l', \alpha_{-2}^l, \alpha_2 \right) \leq \\ &\leq \min_{\alpha_2} u_1 \left(l', \alpha_{-2}^{l'}, \alpha_2 \right) \leq u_1 \left(l', \alpha_{-2}^{l'}, br_2^{l'}(\alpha_{-2}^{l'}) \right) = v_1^{l'} \end{aligned}$$

where the first inequality follows from the IC constraint of player 1 in game $\underline{\Gamma}^{K_2}$.

Finally, consider the announcement $\theta_1 \subseteq K_2$ and $\theta_2 \subseteq K_1$. This occurs for instance if the state is $k \in K_1$ and player 1 announces $L_m \subseteq K_2$. Let $l \in L_m$, then player 1's payoff in the induced penitence play is

$$u_1 \left(k, \alpha_{-2}^l, br_2^{\theta_2}(\alpha_{-2}^l) \right) \leq u_1 \left(k, \alpha_{-2}^k, br_2^{\theta_2}(\alpha_{-2}^k) \right) = v_1^k,$$

otherwise, in game Γ_{K_1} player 1 would have chosen strategies α_{-2}^l as the equilibrium strategies for state k . Similarly, if the state $l \in L_m$ and player 2 announces $\theta_1 = K_1$, then player 2 payoff in the penitence phase is $u_2 \left(l, \alpha_{-2}^l, br_2^{\theta_2}(\alpha_{-2}^l) \right) \leq u_2 \left(l, \alpha_{-2}^l, br_2^l(\alpha_{-2}^l) \right) = v_2^l$, since player 2 prefers to take the best-reply according to her actual payoff function. \square

Supplemental Material to “Belief-free Equilibria in Games with Incomplete Information: Characterization and Existence”

APPENDIX C: PROOF OF THEOREM 5.3

Sufficiency is outlined in the Section 5. For the necessity part, assume that there are two states k, ℓ such at most players 1 and 2 distinguish these two states. Consider the following example, due to Renault (2001). There are three players 1, 2, 3, and we consider only the states k, ℓ . Other

players have no influence on rewards, and rewards in other states do not depend on actions.

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 ℓ 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		

First, assume that only player 1 knows the state and assume that $V^*(\mathcal{I}, u)$ is non-empty. The IR condition for player 3 implies that he plays E in state k and W in state ℓ . Since the preference ordering of player 1 is the opposite of the one of player 3, this violates the IC condition.

Assume now that players 1 and 2 know the state. 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 ℓ , player 3 plays W and the payoff vector in state ℓ is $(0, 0, 1)$.

Now, suppose that player 1 announces k and player 2 announces ℓ : either the true state is k and player 2 is misreporting or the true state is ℓ and player 1 is misreporting. The JR condition implies that there exists a distribution of action profiles α such that $u_1(\ell, \alpha) \leq 0$ and $u_2(k, \alpha) \leq 0$. This is impossible since for each action profile a , $u_1(\ell, a) + u_2(k, a) = 1$. □

APPENDIX D: PROOF OF THEOREM 5.11

Sufficiency: For each state k , fix a vector v^k that is individually rational in the complete information game corresponding to state k , i.e., $v^k \geq \underline{u}^k$. We show that $v := \{v^k\}$ is in V^* . This profile is chosen to be individually rational. *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 choice but to be inconsistent with the other players, or go along with the identification of the state. The distribution corresponding to the bad outcome can be used to deter a player from deviating.

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

Example 7.5 (*This example is adapted from Hörner and Lovo, 2009*). There are two states k, k' , and two players. Player 1 is informed of the state, player 2 is not. The payoff matrix in states k and k' are the following:

	L	M	R		L	M	R
T	10, -4	1, 1	10, -4		0, 0	1, 1	10, -4
B	1, 1	0, 0	-1, -4		1, 1	10, -4	-1, -4
	<i>state k</i>				<i>state k'</i>		

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. □

APPENDIX E: PROOF OF THEOREM 5.12

Necessity can be shown by considering a two-player two-sided game where both players are essential. In this context a counter-example is found in Koren (1992) and in Hörner and Lovo (2009). This example is also in Appendix B (example A). To prove sufficiency, consider a game with known-own payoffs and a bad outcome, and an information structure with at most one essential player per state. Partition the set of states as

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

where for each $k \in K_0$, there is no essential player at k , and for each $s = 1, \dots, S$, there exists a unique player i_s who is essential at states in K_s . That is,

- 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 one 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 $k' \notin K_s = I_{-i}(k)$, 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 $s = 1, \dots, S$, 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 locally weakly embedded. Thus, from Theorem 5.7, V_s^* is non-empty. Let us choose a payoff array in this set, for each s . We construct the overall equilibrium as follows. Let players announce their information:

- If the announcements identify a state $k \in K_0$, v^k is implemented.
- If after 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 played.

The induced payoff array 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 other choice than letting the state be revealed or being inconsistent with the other players. 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 is now aware that the true 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 says that this announcement is inconsistent with some other player's report. Player i_s has thus no other choice than letting K_s be revealed or inducing the bad outcome. This provides a weak incentive to tell the truth. \square

APPENDIX F: PROOF OF THEOREM 6.1

Define

$$u'_1 := \sup_{\{p_i \geq 0: i=2, \dots, N\}} \text{val}(u_1 - \sum_{i=2}^N p_i(u_i - \underline{u}_i 1)),$$

and

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

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 $(p_2, \dots, p_N) \geq 0$ and $\alpha_1 \in \Delta A_1$ such that

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

Conversely, for every $\varepsilon > 0$, there exists $\alpha_1 \in \Delta A_1$ such that $\min_{\alpha_{-1} \in Y(\alpha_1)} u_1(\alpha_1, \alpha_{-1}) \geq u_1'' - \varepsilon$.

Therefore, fixing $\alpha_1 \in \Delta A_1$, for every $\alpha_{-1} \in \mathbb{R}_+^{|A-1|}$,

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

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

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

Therefore,

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

We now show that the bound is attained by $u_1^k = -u_k, \forall k = 2, \dots, N$. Given some equilibrium, let $\mu^i \in \Delta A$ 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 $p \in \Delta\{1, \dots, N\}$,

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

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

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

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

$$v_1^k := \max_{\mu \in \Delta A} \{u_1^k(\mu) : u_1(\mu) \leq u'_1, u_i(\mu) \geq \underline{u}_i, \forall i \geq 2\},$$

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 ensures 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 $p \in \Delta\{1, \dots, K\}$,

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

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

$$u_i(\alpha) \geq \underline{u}_i \cdot \alpha, u'_1 \cdot \alpha \geq u_1(\alpha) \Rightarrow v_1^k \cdot \alpha \geq u_1^k(\alpha).$$

By Farkas' Lemma, for every $k = 1, \dots, K$, there exists $\gamma^k \geq 0, \lambda_i^k \geq 0$ such that $v_1^k \cdot 1 - u_1^k \leq \gamma^k(u'_1 \cdot 1 - u_1) + \sum_{i=2}^N \lambda_i^k(u_i - \underline{u}_i \cdot 1)$. Therefore, for all $p \in \Delta\{1, \dots, K\}$,

$$\text{val}(\sum_{k=1}^K p_k u_1^k) \leq \sum_{k=1}^K p_k v_1^k - \sum_{k=1}^K p_k \gamma^k u'_1 + \text{val}(\sum_{k=1}^K p_k (\gamma^k u_1 - \sum_{i=2}^N \lambda_i^k (u_i - \underline{u}_i \cdot 1))),$$

and so individual rationality for player 1 is satisfied if

$$\sum_{k=1}^K p_k \gamma^k u'_1 \geq \text{val}(\sum_{k=1}^K p_k (\gamma^k u_1 - \sum_{i=2}^N \lambda_i^k (u_i - \underline{u}_i \cdot 1))).$$

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

$$u'_1 \geq \text{val}(u_1 - \sum_{i=2}^N \nu_i (u_i - \underline{u}_i \cdot 1)),$$

which is satisfied by definition of u'_1 . □

APPENDIX G: PROOF OF PROPOSITION 7.1

Take an information structure \mathcal{I} with a single majority component and say that player i is *trivial* if $I_i(k) = K$ for all k ; player i is *non-trivial* otherwise.

Lemma 7.6 *If there are at most two non-trivial players, then either \mathcal{I} is LWE or there is a subset of three states, such that the restriction of \mathcal{I} to this subset is of type A.*

Proof. Let 1, 2 be the two non-trivial players. If it holds for each k that $I_1(k) \subseteq I_2(k)$ or $I_2(k) \subseteq I_1(k)$, then it is LWE. Otherwise there exists a state c such that the two sets $I_1(c), I_2(c)$ are not comparable. That is, there exists c' and c'' such that $c' \in I_1(c) \setminus I_2(c)$ and $c'' \in I_2(c) \setminus I_1(c)$. The subset $\{c, c', c''\}$ is as required. \square

Proposition 7.7 *If there are at least three non-trivial players, then either \mathcal{I} has the all-or-nothing property, or there is a subset of three states such that the restriction of \mathcal{I} to this subset is of type A, B, C or D.*

Proof. The proof is by induction on the number of states. First, assume that there are only three states. We denote by E the 3-state, 3-player, all-or-nothing information structure:

	k_1	k_2	k_3
1	k_1	*	*
2	*	k_2	*
3	*	*	k_3

E

Lemma 7.8 *A 3-state information structure which has only one majority component and which is not LWE is A, B, C, D or E.*

Proof. We prove this by enumeration.

First, because the information is not LWE, there must exist 2 players, say player 1, 2, and three states, denoted k_1, k_2, k_3 , such that $k_1 \notin I_1(k_3), k_2 \in I_1(k_3), k_2 \notin I_2(k_3), k_1 \in I_2(k_3)$. That

is, there must exist two players with non-comparable information at some state. We discuss the information of the other players.

1. If all other players have no information, this is A. Otherwise:
 2. If some player (player 3) is fully informed:
 - (a) If all other players have no information, this is B.
 - (b) If player 4 has some information, there is more than one majority component. For instance, if player 4 has the same information as player 1, $\{k_3\}$ is a majority component. The reasoning is the same if player 4 has the same information as player 2. If the information of player 4 is $I_4(k_3) \neq I_4(k_1) = I_4(k_2)$, we have the same conclusion: three players (1, 3, 4) can distinguish k_1 and k_3 , and three players (2, 3, 4) can distinguish k_2 and k_3 , so $\{k_3\}$ is a majority component.
3. If no player is fully informed, but some player (player 3) is partially informed:
 - (a) If all other players have no information, this is C (up to a relabelling of players) or E.
 - (b) If player 4 also has partial information, all other players being uninformed, then it is either D or there is more than one majority component. By symmetry we may assume that players 3 and 4 have the same information. If it is the same as that of player 1 (resp. player 2) then $\{k_3\}$ is a majority component. Otherwise, it is equivalent to E, with a fourth player having the same information as 1, 2 or 3. In this case, one sees easily that if the fourth player has the same information as (e.g.) player 1, $\{k_1\}$ is a majority component.

(c) Finally, if players 4 and 5 have partial information, there is more than one majority component. There are three types of partial information and five players. Either three of them have the same information and they can then distinguish states. Or the information structure is the symmetric one, with two duplicated players, which leads back to the previous case. \square

Let us do now the induction step. Take $|K| > 3$ and assume that the statement of Proposition 7.7 holds for $|K| - 1$. We consider an information structure with $|K|$ states which has only one majority component, at least three non-trivial players and which is not all-or-nothing.

Consider the relation on states defined as aRb iff $\nu(a, b) \leq 2$, and consider also the graph of this relation. \mathcal{I} has only one majority component means that this graph is connected. Note that if we delete a state and all its adjacent edges, we obtain the graph of the relation on the restricted set of states. Take now two states a and b such that there is a path in the graph from a to b with maximal length among the paths in this graph. The graph obtained by suppressing a (resp. b) is still connected. Indeed, any other point c is connected to b (resp. a) by a path that does not go through a (resp. b), since otherwise, this would contradict the maximality of the path from a to b . It follows that $\mathcal{I}_{K \setminus \{a\}}$ (resp. $\mathcal{I}_{K \setminus \{b\}}$) has only one majority component.

If $\mathcal{I}_{K \setminus \{a\}}$ or $\mathcal{I}_{K \setminus \{b\}}$ has at least three non-trivial players and is not symmetric, we are done by induction. Assume otherwise.

Case A. Both $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$ have at least three non-trivial players and are all-or-nothing. First, the non-trivial players are the same for $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$. Indeed, let i be non-trivial for $\mathcal{I}_{K \setminus \{a\}}$. There exists $k \neq a$ such that $I_i(k) \cap K \setminus \{a\} = \{k\}$, so that $I_i(k) \subseteq \{k, a\}$. Then i cannot be trivial in $\mathcal{I}_{K \setminus \{b\}}$: for a trivial player $I_{i, K \setminus \{b\}}(k)$ contains at least three states. Let now $1, \dots, m$ be these non-trivial players.

Let K_1, \dots, K_m be the partition induced by $\mathcal{I}_{K \setminus \{a\}}$ on $K \setminus \{a\}$. Since $\mathcal{I}_{K \setminus \{b\}}$ is all-or-nothing, there is a unique player, say player 1, such that $I_{1, K \setminus \{b\}}(a) = \{a\}$. So that $I_1(a) \subseteq \{a, b\}$.

- If $b \in K_1$, consider two other non-trivial players j, l and $c' \in K_l$. By the all-or-nothing property of $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$, one has $a \in I_j(c')$ and $b \in I_j(c')$. So that j does not distinguish a and b . Now, either $I_1(a) \neq I_1(b)$ and \mathcal{I} is all-or-nothing, or $I_1(a) = I_1(b)$ and no player distinguishes a from b . In both cases, this is a contradiction.
- If $b \notin K_1$, say $b \in K_2$. If $I_1(a) = I_1(b)$, take c in K_3 . By the all-or-nothing property of $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$, $c \in I_1(b)$ contradicting $I_1(a) \subseteq \{a, b\}$. Thus $I_1(a) \neq I_1(b)$, that is $I_1(a) = \{a\}$ and by the all-or-nothing property of $\mathcal{I}_{K \setminus \{a\}}$, $I_1(b) = K \setminus (K_1 \cup \{a\})$. By the all-or-nothing property of $\mathcal{I}_{K \setminus \{a\}}$, no player, except player 2, distinguishes b from other states in K_2 and by the all-or-nothing property of $\mathcal{I}_{K \setminus \{b\}}$, no player, except player 1, distinguishes a from other states in K_1 . Thus I has the all-or-nothing property, a contradiction.

Case B. Both $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$ have at most two non-trivial players. If $\mathcal{I}_{K \setminus \{a\}}$ or $\mathcal{I}_{K \setminus \{b\}}$ is not LWE, we are done by lemma 7.6. Assume to the contrary that both are LWE. Then $\mathcal{I}_{K \setminus \{a, b\}}$ is LWE as well which implies that the two non-trivial players are the same in $\mathcal{I}_{K \setminus \{a\}}$ and $\mathcal{I}_{K \setminus \{b\}}$, say players 1 and 2. This implies that suppressing a or b changes some player, say player 3, from non-trivial to trivial, which is not possible.

Case C. $\mathcal{I}_{K \setminus \{b\}}$ has at least three non-trivial players and has the all-or-nothing property and $\mathcal{I}_{K \setminus \{a\}}$ has at most two non-trivial players. If $\mathcal{I}_{K \setminus \{a\}}$ is not LWE, we are done by lemma 7.6. Assume the contrary and consider $\mathcal{I}_{K \setminus \{a, b\}}$. This is both all-or-nothing and LWE. This shows that the non-trivial players from $\mathcal{I}_{K \setminus \{a\}}$ are non-trivial in $\mathcal{I}_{K \setminus \{b\}}$ as well, and that $\mathcal{I}_{K \setminus \{b\}}$ has exactly three non-trivial players called henceforth 1, 2, 3. Suppressing a transforms, say player 1, from non-trivial to trivial. So it must be the case that $I_1(a) = \{a\}$ and $I_1(k) = K \setminus \{a\}$ for $k \neq a$.

Let us choose now $c \neq a$ such that $I_2(c) = I_2(a)$ (which exists, because player 1 is the only informed player at a) and assume that $I_3(c) \subset I_2(c)$. Take $d \in I_2(c) \setminus I_3(c)$. The information structure on $\{a, c, d\}$ is of type C or D, depending on whether player 3 can distinguish a from c or not. If it is not the case that one can choose such a d (even by exchanging the roles of 2 and 3), it means that players 2 and 3 have the same information structure. One just has to choose a , $c \neq a$ such that $I_2(c) = I_3(c) = I_2(a)$ and $d \neq a$ outside of $I_2(c)$, to end up with a type C. This concludes the proof. \square