ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 77, No. 2 (March, 2009), 453–487

BELIEF-FREE EQUILIBRIA IN GAMES WITH INCOMPLETE INFORMATION

JOHANNES HÖRNER
Yale University, New Haven, CT 06520-8281, U.S.A. and Centre for Economic Policy Research, London EC1V 0DG, U.K.

STEFANO LOVO
HEC School of Management, Paris and GREGHEC, 78351 Jouy-en-Josas, France

Econometrica, Vol. 77, No. 2 (March, 2009), 453–487

BELIEF-FREE EQUILIBRIA IN GAMES WITH INCOMPLETE INFORMATION

BY JOHANNES HÖRNER[1] AND STEFANO LOVO

We define belief-free equilibria in two-player games with incomplete information as sequential equilibria for which players' continuation strategies are best replies after every history, independently of their beliefs about the state of nature. We characterize a set of payoffs that includes all belief-free equilibrium payoffs. Conversely, any payoff in the interior of this set is a belief-free equilibrium payoff. The characterization is applied to the analysis of reputations.

KEYWORDS: Repeated game with incomplete information, Harsanyi doctrine, belief-free equilibria.

1. INTRODUCTION

THE PURPOSE OF THIS PAPER is to characterize the set of payoffs that can be achieved by equilibria that are robust to specification of beliefs. The games considered are two-player discounted repeated games with two-sided incomplete information and observable actions. The equilibria whose payoffs are studied are such that the players' strategies are optimal from any history on and independently of players' beliefs about their opponent's type. This concept is not new. It has been introduced in another context, namely in repeated games with imperfect private monitoring, in Piccione (2002) and Ely and Välimäki (2002), and further examined in Ely, Hörner, and Olszewski (2005). It is also related to the concept of ex post equilibrium that is used in mechanism design (see Crémer and McLean (1985)) as well as in large games (see Kalai (2004)). A recent study of ex post equilibria and related belief-free solution concepts in the context of static games of incomplete information was provided by Bergemann and Morris (2007).

To predict players' behavior in games with unknown parameters, a model typically includes specification of the players' subjective probability distributions over these unknowns, following Harsanyi (1967–1968). This is not necessary when belief-free equilibria are considered, as their characterization requires a relatively parsimonious description of the model. One needs to enumerate the set of possible states of the world and players' information partitions over these states, but it is no longer necessary to specify players' beliefs. Therefore, while solving for belief-free equilibria requires the game to be fully specified, it does not require that all players know all the parameters of the

[1]We thank Gabrielle Demange, Jeff Ely, Stephen Morris, Tristan Tomala, and, particularly, Larry Samuelson for useful comments. We are grateful to two anonymous referees and the editor for insightful comments and suggestions that substantially improved the paper. Stefano Lovo gratefully acknowledges financial support from the HEC Foundation.

 DOI: 10.3982/ECTA7134

model. In this sense, this idea is close to the original motivation of von Neumann and Morgenstern (1944) in defining "games of incomplete information" as games in which some parameters remain unknown, and it is consistent with misperceptions as defined by Luce and Raiffa (1957). Nevertheless, as in the case of games with perfect information, players are expected utility maximizers: players are allowed to randomize, and take expectations with respect to such mixtures when evaluating their payoff.[2] Our purpose is to characterize which equilibria do not require any probabilistic sophistication beyond that assumed in repeated games with perfect information.

Just as for ex post equilibria, belief-free equilibria enjoy the desirable property that the beliefs about the underlying uncertainty are irrelevant. This means that they remain equilibria when players are endowed with arbitrary beliefs. Such beliefs need not be derived by Bayes' rule from a common prior. Furthermore, the way in which players update their beliefs as the game unfolds is irrelevant. For instance, belief-free equilibria remain equilibria if we allow players to observe a signal of their stage-game payoff and to learn in this way about the other player's private information. Thus, belief-free equilibria are robust to all specifications of how players form and update their belief. In particular, belief-free equilibria are sequential equilibria (for any prior) satisfying any potentially desirable refinement. In a belief-free equilibrium, the players' strategies must be a subgame-perfect Nash equilibrium of the game of complete information that is determined by the joint of their private information. However, we do not view belief-free equilibrium as an equilibrium refinement per se. In fact, belief-free equilibria need not exist. The robustness that is demanded is extreme in the sense that it is not only required that the strategies be mutual best replies for a neighborhood of beliefs, but for all possible beliefs. Note that it is also stronger than the way it is modeled in the recent macroeconomics literature (Hansen and Sargent (2007)), since the property examined here treats all possible beliefs identically.

We provide a set of necessary conditions that belief-free equilibrium payoffs must satisfy, which defines a closed convex and possibly empty set. Conversely, we prove that every interior point of this set is a belief-free equilibrium payoff, provided that players are sufficiently patient. In the proof, we also show how to construct a belief-free equilibrium supporting any payoff in this set. This equilibrium has a recursive structure similar to standard constructions based on an equilibrium path and a punishment path for each player. While the set of belief-free equilibria is empty for some games, belief-free equilibria exist for large classes of games studied in economics such as, for example, most types of auctions, Cournot games, and Bertrand games. Constructing "belief-based" equilibria generally requires keeping track of beliefs and even of hierarchies of beliefs. This is usually untractable unless the information structure is quite

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

special or the game is repeated at most twice.[3] This problem does not arise with belief-free equilibria, thus offering a possible route for the analysis of dynamic economic interactions with relatively complex and realistic information structures at a time when there is a strong interest in modeling robustness in economics.

The set of belief-free equilibrium payoffs turns out to coincide with a set that plays a prominent role in the literature on Nash equilibria in games with one-sided incomplete information. Building on this literature, we describe the implications of the concept to the study of reputations. In particular, the Stackelberg payoff is equal to the lowest (belief-free) equilibrium payoff if the game is of conflicting interest, which is precisely the type of game typically used as example to show how surprisingly limited reputation effects are when players are equally patient. More generally, focusing attention on belief-free equilibria with equally patient players is shown to involve restrictions on the equilibrium payoff set similar to those of Nash equilibria when the informed player is infinitely more patient than the uninformed player.

As mentioned, the set of payoffs that characterizes belief-free equilibria has already appeared in the literature, at least in the case of one-sided incomplete information. In particular, Shalev (1994) considered the case of known-own payoffs (the uninformed player knows his own payoffs) and showed that the set of uniform (undiscounted) Nash equilibrium payoffs can be derived from this set. Closest to our analysis is Cripps and Thomas (2003), which considered the one-sided case with known-own payoffs as well, but with discounting. Most relevant here is their Theorem 2, which establishes that the payoffs in the strict interior of this set are Nash equilibria for all priors. In general, however, the set of Nash equilibrium payoffs is larger, as they demonstrated in their Theorem 3, which establishes a folk theorem. The work of Forges and Minelli (1997) is related as well. They showed how communication can significantly simplify the construction of strategies that achieve the Nash equilibrium payoffs. These simple strategies also appear in Koren (1992). The most general characterization of Nash equilibrium payoffs remains the one obtained by Hart (1985) for the case of one-sided incomplete information. A survey is provided by Forges (1992).

The assumptions of Bayesianism has already been relaxed in several papers. Baños (1968) and Megiddo (1980) showed that strategies exist that asymptotically allow a player to secure a payoff as high as in the game with complete information. Milnor (1954) reviewed several alternative criteria and discuss their relative merits. The topic has also been explored in computer science. Aghassi and Bertsimas (2006) used robust optimization techniques to provide an alternative concept in the case of bounded payoff uncertainty. Monderer and

[3]Examples can be found in financial economics, where sequential trade of a security is often modeled as a dynamic auction, and in industrial organization, where repeated Cournot games are used to model dynamic competition among firms.

Tennenholtz (1999) studied the asymptotic efficiency in the case in which players are non-Bayesian and monitoring is imperfect. All these papers either offer an alternative equilibrium concept or study what is asymptotically achievable without using any solution concept. Yet the strategy profiles that are characterized in these papers are not Bayesian equilibria (at least under discounting), which is a major difference with our paper.

As mentioned, the concept of belief-free equilibria has already been introduced in the context of games with complete but imperfect information. There, the restriction on the equilibrium pertains to the private history observed by the opponent. In both contexts, the application of the concept reduces the complexity of the problem (players need no longer keep track of the relevant beliefs) and yields a simple characterization. In games with imperfect private monitoring, this has further allowed the construction of equilibria in cases in which only trivial equilibria were known so far.

The next section introduces notations and definitions. Section 3 then provides the payoff characterization, identifying in turn necessary and sufficient conditions on payoffs that belief-free equilibrium payoffs satisfy. This section also gives a relatively short proof of sufficiency using explicit communication (the proof without such communication is given in Appendix) and provides counterexamples to existence, as well as sufficient conditions for existence. Section 4 applies the concept to the study of reputations.

2. NOTATION AND DEFINITIONS

We consider repeated games with two-sided incomplete information, as defined by Harsanyi (1967–1968) and Aumann and Maschler (1995). There is a $J \times K$ array of two-person games in normal form. The number of actions of player $i = 1, 2$ is the same across all $J \times K$ games. At the beginning of time and once for all, Nature chooses the game in the $J \times K$ array. Player 1 is told in which *row* $j = 1, \dots, J$ the true game lies, but he is not told which of the games in that row is actually being played. Player 2 is told in which *column* $k = 1, \dots, K$ the true game lies, but he is not told which of the games in that column is the true game. The row j (respectively, column k) is also referred to as player 1's (respectively, player 2's) *type*. Given some finite set B, $|B|$ denotes the cardinality of B and $\triangle B$ denotes the probability simplex over B. Also, given some set B, let $\operatorname{int} B$ denote its interior and $\operatorname{co} B$ denote its convex hull.

The stage game is a finite-action game. Let A_1 and A_2 be the finite sets of actions for players 1 and 2, respectively, where $|A_i| \geq 2$. Let $A = A_1 \times A_2$.

When the row is j and the column is k (for short, when the state is (j, k)), player i's reward (or payoff) function is denoted u_i^{jk} for $i = 1, 2$. We extend the domain of u_i^{jk} from pure action profiles $a \in A$ to mixed action profiles $\alpha \in \triangle A$

in the standard way. We let $u_1^k := \{u_1^{jk}\}_{j=1}^J$ and $u_2^j := \{u_2^{jk}\}_{k=1}^K$. The set of feasible payoffs in $\mathbb{R}^{J\times K} \times \mathbb{R}^{J\times K}$ is defined, as usual, as

$$\mathrm{co}\big\{\big((u_1^{jk}(a))_{(j,k)}, (u_2^{jk}(a))_{(j,k)}\big) : a \in A\big\}.$$

Let $M := \max|u_i^{jk}(a)|$, where the maximum is taken over players $i = 1, 2$, states (j, k), and action profiles $a \in A$. Given some payoff function u, let $\underline{u}$ or $\mathrm{val}\, u$ refer to the corresponding minmax payoff. We let $B_i^{jk}(\alpha_{-i})$ denote the set of player i's best replies in the stage game given state (j, k) and action α_{-i} of player $-i$. We omit the superscript k in case $|K| = 1$, that is, if the game is of one-sided incomplete information. If furthermore player 2's payoff does not depend on j, we write $B(\alpha_1)$ for his set of best replies.

As an example, consider the stage game given below. Since this game is dominance solvable and the dominant action depends on the state, ex post equilibria do not exist in the static game. Yet as we shall see, the repeated game admits a rich set of belief-free equilibria.

EXAMPLE 1—Prisoner's Dilemma With One-Sided Incomplete Information: Player 1 is informed of the true state (= the row), player 2 is not, and there is only one column ($J = 2$, $K = 1$). If the true game corresponds to $j = 1$, payoffs are given (in every period) by the prisoner's dilemma payoff matrix in which T is "cooperate" and B is "defect." If the true game corresponds to $j = 2$, payoffs are given by the prisoner's dilemma payoff matrix in which B is "cooperate" and T is "defect." The payoffs in the first state are given by

	T	B
T	$1, 1$	$-L, 1+G$
B	$1+G, -L$	$0, 0$

and in the second state by

	T	B
T	$0, 0$	$1+G, -L$
B	$-L, 1+G$	$1, 1$

We consider the repeated game between the two players. Players select an action in each period $t = 1, 2, \ldots$. Realized actions are observable, mixed actions and realized rewards are not.

Let $H^t = (A_1 \times A_2)^{t-1}$ be the set of all possible histories of actions h^t up to period t, with $H^1 = \emptyset$. A (behavioral) strategy for type j of player 1 (resp. type k of player 2) is a sequence of maps $s_1^j := (s_1^{j,1}, s_1^{j,2}, \ldots)$, $s_1^{j,t} : H^t \to \triangle A_1$ (resp. $s_2^k := (s_2^{k,1}, s_2^{k,2}, \ldots)$, $s_2^{k,t} : H^t \to \triangle A_2$). We define $s_1 := \{s_1^j\}_{j=1}^J$ and $s_2 := \{s_2^k\}_{k=1}^K$.

Consider the game of complete information given state (j, k). Given the common discount factor $\delta < 1$, player i's payoff in this game is the average discounted sum of expected rewards. A subgame-perfect Nash equilibrium of this game is defined as usual.

Our purpose is to characterize the payoffs that can be achieved, with low discounting, by a special class of Nash equilibria. In a *belief-free* equilibrium, each player's continuation strategy, after any history, is a best reply to his opponent's continuation strategy, independently of his beliefs about the state of the world and, therefore, independently of his opponent type. Such equilibria are trivially sequential equilibria that satisfy any belief-based refinement. At the same time, they do not require players to be Bayesian or to share a common prior. Because they are belief-free, they must, in particular, induce a subgame-perfect equilibrium in every complete information game that is consistent with the player's private information. Formally, a belief-free equilibrium is defined as follows.

DEFINITION 1: A strategy profile $s := (s_1, s_2)$ is a belief-free equilibrium if it is the case that, for all states (j, k), (s_1^j, s_2^k) is a subgame-perfect Nash equilibrium of the infinitely repeated game with stage-game payoffs given by (u_1^{jk}, u_2^{jk}).

As mentioned, belief-free equilibria have been previously introduced in and applied to games with imperfect private monitoring. With incomplete information but observable actions, there is no need for randomization on the equilibrium path. Indeed, in our construction, along the equilibrium path, players always have a strict preference to play some particular action. Of course, this action potentially depends on a player's private information (and on the history). In our construction, randomization is only necessary during punishment phases, as is standard in folk theorems that allow for mixed strategies to determine minmax payoffs, as we do.[4]

It follows from the definition of belief-free equilibria that even when different player's types use the same strategy, it would be weakly optimal for them to reveal their type (if there was a communication device). Indeed, by definition, the strategy profile that is played is an equilibrium of the underlying complete information game. This means that pooling belief-free equilibria are simply

[4]Yet a randomization device considerably simplifies the exposition. At the end of the Appendix, we indicate how to dispense with it.

"degenerately separating" belief-free equilibria. In particular, the payoffs of pooling belief-free equilibria are in the closure of the set of payoffs achieved by separating belief-free equilibria.

This finding implies that this concept is more restrictive than most refinements, since refinements do not usually prune all pooling equilibria that are not degenerate separating ones.[5]

3. CHARACTERIZATION

Any belief-free equilibrium determines a $J \times K$ array of payoffs (v_i^{jk}), for each player $i = 1, 2$. We first provide necessary conditions that such a pair of arrays must satisfy, before providing sufficient conditions that ensure they are achieved by some belief-free equilibrium.

3.1. *Necessary Conditions*

For definiteness, consider $i = 1$. Conditional on the column k he is being told, player 2 knows that player 1's equilibrium payoff is one among the coordinates of the vector $v_1^k = (v_1^{1k}, \dots, v_1^{Jk})$. Because the equilibrium is belief-free, player 1's payoff must be individually rational in the special case in which his beliefs are degenerate on the true column k. This means that, for a given column k, player 2's strategy s_2^k is such that player 1 cannot gain from deviating from s_1^j, for all $j = 1, \dots, J$. The existence of such a strategy s_2^k puts a restriction on how low player 1's payoff v_1^{jk} can be (in fact, a joint restriction on the coordinates of the vector v_1^k).

If $J = 1$, so that the game is of one-sided incomplete information, this restriction on player 1's payoff is standard: for each k, player 1 must receive at least as much as his minmax payoff (in mixed strategies) in the true game being played. In the general case however, the minmax level in one state depends on the payoffs in the other states, and there is a trade-off between these levels: punishing player 1 for one row may require conceding him a high payoff for some other row. Determining these minmax levels is not obvious. This is precisely the content of Blackwell's approachability theorem (Blackwell (1956)).

For a given $p \in \triangle\{1, \dots, J\}$ (resp. $q \in \triangle\{1, \dots, K\}$), let $b_1^k(p)$ (resp. $b_2^j(q)$) be the value for player 1 (resp. player 2) of the one-shot game with payoff matrix $p \cdot u_1^k$ (resp. $q \cdot u_2^j$). We say that a vector $v_1 \in \mathbb{R}^{J \times K}$ is *individually rational* for player 1 if it is the case that, for all $k = 1, \dots, K$,

$$p \cdot v_1^k \geq b_1^k(p) \quad \forall p \in \triangle\{1, \dots, J\},$$

[5]Indeed, many games admit "traditional" pooling equilibria in which individual rationality holds in expectation, but not conditional on every type of opponent.

where $v_1^k = (v_1^{1k}, \dots, v_1^{Jk})$. Similarly, $v_2 \in \mathbb{R}^{J\times K}$ is *individually rational* for player 2 if it is the case that, for all $j = 1, \dots, J$,

$$q \cdot v_2^j \geq b_2^j(q) \quad \forall q \in \triangle\{1, \dots, K\},$$

where $v_2^j = (v_2^{j1}, \dots, v_2^{jK})$. Blackwell's characterization ensures that if $v_1 \in \mathbb{R}^{J\times K}$ is individually rational for player 1, then for any column k, player 2 has a strategy $\hat{s}_2^k$ (referred to as a *punishment strategy* hereafter) such that player 1's average payoff cannot be larger than v_1^{jk} for all j independently of the strategy he uses. In a belief-free equilibrium, each player can guarantee that his payoff is individually rational, independently of the discount factor.[6]

NECESSARY CONDITION 1—Individual Rationality: If v_i is a belief-free equilibrium payoff, then it is individually rational.

In a belief-free equilibrium, play may depend on a player's private information. That is, player 1's equilibrium strategy s_1^j typically depends on the row j he is told, and player 2's strategy s_2^k depends on the row k. Since player 1's strategy s_1^j must be a best reply to s_2 independently of his beliefs, it must be a best reply to s_2^k, corresponding to beliefs that are degenerate on the true column k. In particular, s_1^j must be a better reply to s_2^k than $s_1^{j'}$, $j' \neq j$, when the row is j. While this might seem a weaker restriction than individual rationality, it is not implied by it, since it places restrictions on the equilibrium path. By deviating to $s_1^{j'}$ when the state is (j, k), player 1 induces the same distribution over action profiles as the one generating the payoff $v_1^{j'k}$ in state (j', k). This imposes additional restrictions on the equilibrium strategies.

To state this second necessary condition in terms of payoffs, observe that each pair (s_1^j, s_2^k) induces a probability distribution $\{\Pr\{a \mid (j,k)\} : a \in A\}_{(j,k)}$ over action profiles, where

$$\Pr\{a \mid (j,k)\} = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} \Pr\{a_t = a \mid (s_1^j, s_2^k)\}$$

and $\Pr\{a_t = a \mid (s_1^j, s_2^k)\}$ is the probability that action a is played in period t given the strategy profile (s_1^j, s_2^k).

[6]The punishments that can be imposed in the discounted game are lower than, but converge uniformly to, those that can be imposed in the undiscounted game. See Cripps and Thomas (2003) and references therein. We thank a referee for pointing out that individual rationality must hold for all discount factors.

NECESSARY CONDITION 2—Incentive Compatibility: If (v_1, v_2) is a pair of belief-free equilibrium payoff arrays, there must exist distributions $\{\Pr\{a \mid (j,k)\} : a \in A\}_{(j,k)}$ such that, for all (j,k),

$$v_1^{jk} = \sum_a \Pr\{a \mid (j,k)\} u_1^{jk}(a) \geq \sum_a \Pr\{a \mid (j',k)\} u_1^{jk}(a)$$

and

$$v_2^{jk} = \sum_a \Pr\{a \mid (j,k)\} u_2^{jk}(a) \geq \sum_a \Pr\{a \mid (j,k')\} u_2^{jk}(a).$$

If such distributions exist, we say that (v_1, v_2) is incentive compatible. While not every pair of payoff arrays is incentive compatible, there always exist some incentive compatible pairs, since the constraints are trivially satisfied for distributions $\Pr\{a \mid (j,k)\}$ that are independent of (j,k).

3.2. *Sufficient Conditions*

Let V^* denote the feasible set of pairs of payoff arrays satisfying Conditions 1 and 2. It is clear that V^* is convex. Our main result is the following.

THEOREM 1: *Fix some v in the interior of V^*. The pair of payoff arrays v is achieved in some belief-free equilibrium if players are sufficiently patient.*

This theorem establishes that the necessary conditions are "almost" sufficient. It is then natural to ask whether we can get an exact characterization. However, the strict inequalities corresponding to individual rationality cannot be generally weakened. One reason for this is that our optimality criterion involves discounting, while Blackwell's characterization of approachability is only valid for the undiscounted case. The strict inequalities corresponding to incentive compatibility may be weakened when V^* has nonempty interior. However, for the interesting case in which V^* has empty interior, this may not be possible.[7]

While belief-free equilibria need not exist, as shown in Section 3.4, they exist in a variety of games. For instance, the game in Example 1 admits a large set of belief-free equilibrium payoffs. Figures 1 and 2 display the resulting equilibrium payoffs.[8]

[7]Consider for instance the case of one-sided incomplete information: player 1 knows the row, but his payoff does not depend on the row, so the incentive compatibility constraints necessarily bind.

[8]Note that these are the projections of the equilibrium payoff pairs onto each player's payoff. It is not true that every pair of vectors selected from these projections is a pair of equilibrium payoff vectors. Incentive compatibility imposes some restrictions on the pairing. Details on

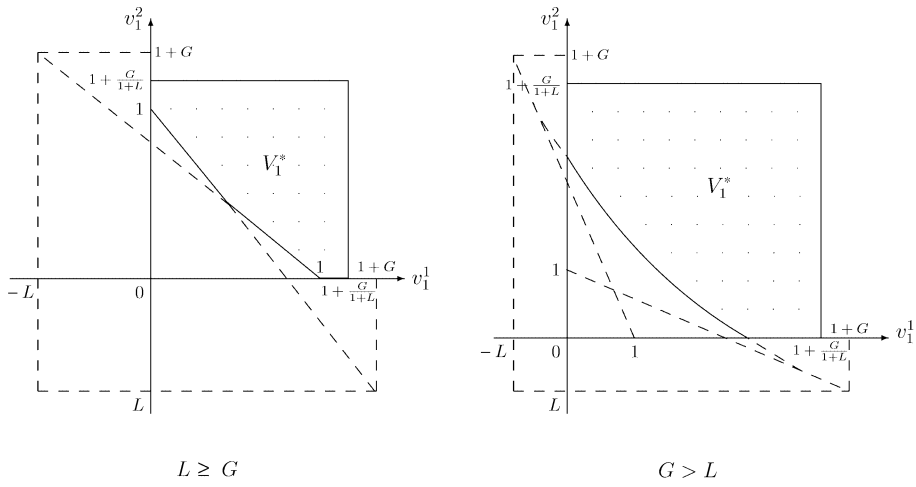

FIGURE 1.—Belief-free equilibrium payoffs for player 1 as $\delta \to 1$.

3.3. *Sketch of the Proof*

The proof of the theorem is constructive. A natural way to proceed is to follow Koren (1992) and others. First, players signal their type (through their choice of actions). Given the reported types, players then choose actions so as

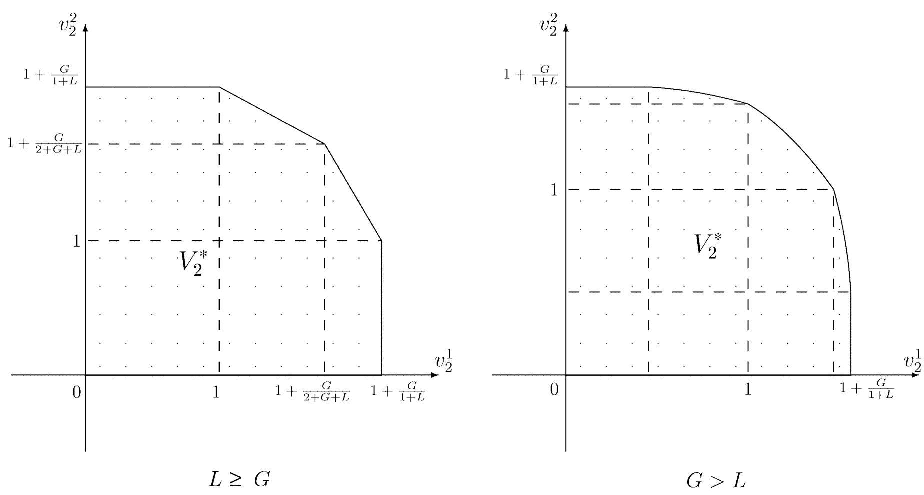

FIGURE 2.—Belief-free equilibrium payoffs for player 2 as $\delta \to 1$.

the derivation of individually rational and incentive compatible payoffs for Example 1 can be found in the working paper HEC CR 884/2008 available at http://www.hec.fr/hec/fr/professeurs_recherche/upload/cahiers/CR884SLOVO.pdf. We thank one referee for pointing out a mistake.

to generate the distribution over action profiles corresponding to these reports. If a player deviates in this second phase, he is minmaxed. Individual rationality guarantees that deviating after some report yields a lower payoff than equilibrium play does, independently of the state. Incentive compatibility ensures that truthful reporting is optimal.

However, such strategies typically fail to be sequentially rational. Minmaxing forever one's opponent need not be individual rational. While this issue can be addressed with the obvious modification, a more serious difficulty is that the resulting strategy profile still fails to be belief-free. In particular, if a player believes that the reported type is incorrect, following his prescribed continuation strategy is no longer individually rational.

The actual construction is therefore more involved, to ensure that beliefs are irrelevant after every possible history. To simplify exposition, we assume here that there is a public randomization device and that players can communicate at no cost in every period. These assumptions are dropped in the Appendix. So suppose that at the beginning of each period, a draw from the uniform distribution on the unit interval (independent of the state of nature and over time) is publicly observed, and suppose that at the beginning of the game (before the first draw is realized) and at the end of *every* period, players simultaneously make a report that is publicly observable. The set of possible reports is the set of rows and columns, respectively. Player 1 reports some $j' = 1, \dots, J$, while player 2 reports some $k' = 1, \dots, K$.

In every period, and using the most recent outcome of the randomization device as a correlation device, a correlated action profile is played that only depends on the last pair of reports made by the players. These correlated action profiles are such that each player obtains the desired payoff whenever $(j', k') = (j, k)$, that is, whenever reports are correct, and such that this payoff exceeds what can be obtained by misreporting, independently of the type truthfully reported by the opponent. Thus, players are willing to report their type truthfully, regardless of their beliefs. In case a player deviates from the prescribed action, he is then punished for finitely many periods. Making sure that play during such a punishment phase is also belief-free introduces some additional complications.

Because players report their types infinitely often, a player who believes that his opponent's report is incorrect still expects his opponent to revert to the true report in the next period. As a consequence, it is less costly for him to play for one period according to the report that he believes to be false than to deviate and to face a long punishment phase.

More formally, given some $v \in \operatorname{int} V^*$, we first describe the equilibrium strategies, and then check that these strategies (i) achieve v and (ii) are best replies that are belief-free.

Equilibrium Strategies

The play can be divided into phases, which are similar to states of an automaton. There are two kinds of phases. *Regular* phases last one period. *Punishment* phases can last from 1 to T periods, where T is to be specified. Regular phases are denoted $R^{jk}(\varepsilon_1, \varepsilon_2)$, where $\varepsilon_1, \varepsilon_2 \in \mathbb{R}$. Punishment phases are denoted P_1^k, P_2^j.

Actions

(i) *Regular phase*: In a regular phase $R^{jk}(\varepsilon_1, \varepsilon_2)$, actions are determined by the outcome of the public randomization device. Each action profile a is selected with probability $\Pr\{a \mid R^{jk}(\varepsilon_1, \varepsilon_2)\}$. Given

$$v_i^{jk}(R^{j'k'}(\varepsilon_1, \varepsilon_2)) := \sum_{a \in A} \Pr\{a \mid R^{j'k'}(\varepsilon_1, \varepsilon_2)\} u_i^{jk}(a)$$

and $v_i(R(\varepsilon_1, \varepsilon_2)) := \{v_i^{jk}(R^{jk}(\varepsilon_1, \varepsilon_2))\}_{(j,k)}$, these probabilities, along with some number $\bar{\varepsilon} > 0$, are chosen such that

$$v_i(R(\varepsilon_1, \varepsilon_2)) = v_i + \varepsilon_i \tag{1}$$

and

$$\begin{aligned} &v_1^{jk}(R^{jk}(\varepsilon_1, \varepsilon_2)) > v_1^{jk}(R^{j'k}(\varepsilon_1', \varepsilon_2')), \\ &v_2^{jk}(R^{jk}(\varepsilon_1, \varepsilon_2)) > v_2^{jk}(R^{jk'}(\varepsilon_1', \varepsilon_2')) \end{aligned} \tag{2}$$

for all $i = 1, 2$, $\varepsilon_i, \varepsilon_i' \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, $j' \neq j$, $k' \neq k$. This is possible for all sufficiently small $\bar{\varepsilon}$ by incentive compatibility, given that $v \in \operatorname{int} V^*$.

At the end of a regular phase, types are reported truthfully.

(ii) *Punishment phase*: The punishment phase lasts at most T periods. Without loss of generality, we describe here the actions and reports in phase P_1^k. Both the subscript (the identity of the punished player) and the superscript (the reported type by the punisher) remain constant throughout the phase. Decreasing $\bar{\varepsilon}$ if necessary, the (behavior) strategy $\widehat{s}_2^k$ of player 2 during the punishment phase P_1^k is such that, for some $\bar{\delta} < 1$ and all discount factors $\delta > \bar{\delta}$, the average discounted payoff of player 1 over the T periods, conditional on state (j, k), is no larger than $v_1^{jk} - 2\bar{\varepsilon}$. This is possible for all sufficiently large T by individual rationality, given that $v \in \operatorname{int} V^*$.

We further assume that T, $\bar{\delta}$, and $\bar{\varepsilon}$ satisfy, for all j, k, and $i = 1, 2$,

$$\begin{aligned} &-(1-\delta)M + \delta(v_i^{jk} - \bar{\varepsilon}) \\ &\quad > (1-\delta)M + \delta\big((1-\delta^T)(v_i^{jk} - 2\bar{\varepsilon}) + \delta^T(v_i^{jk} - \bar{\varepsilon})\big) \end{aligned} \tag{3}$$

and

$$(4) \qquad -(1-\delta^T)M + \delta^T v_i^{jk} > (1-\delta^T)M + \delta^T(v_i^{jk} - 2\bar{\varepsilon}/3).$$

To see that such T, $\bar{\delta}$, and $\bar{\varepsilon}$ exist, observe that for a fixed but small enough $\bar{\varepsilon} > 0$, (3) is satisfied for all T large enough and $\delta > \bar{\delta}$ for $\bar{\delta}$ close enough to 1. Increasing the value of $\bar{\delta}$ if necessary, (4) is then satisfied as well.

Returning to the specification of actions and reports, as long as the punishment phase P_1^k lasts (i.e., for at most T periods), player 2 plays according to $\widehat{s}_2^k$ (given k and the history starting in the initial period of P_1^k). Observe that $\widehat{s}_2^k$ need not be pure. Player 1 plays a best reply s_1^{jk} to $\widehat{s}_2^k$, conditional on the true column being k. Without loss of generality, we pick s_1^{jk} to be pure. Observe that s_1^{jk} may depend on j.

Players report truthfully types in all periods of the punishment phase.

(iii) *Initial phase*: As mentioned, players report types at the beginning of the game. These initial reports are made truthfully. The initial phase is the regular phase $R^{jk}(0, 0)$, where (j, k) are the initial reports.

Transitions

(i) *From a regular phase* $R^{jk}(\varepsilon_1, \varepsilon_2)$: If the action of player 1 (resp. player 2) differs from the prescribed action, while player 2 (resp. 1) plays the prescribed action, then the next phase is $P_1^{k'}$ (resp. $P_2^{j'}$), where k' (resp. j') is the report made at the end of the period by the corresponding player. (Observe that the message of the deviator plays no role here.) Otherwise: (a) if $(j', k') = (j, k)$ or both $j \neq j'$ and $k \neq k'$, the next phase is $R^{j'k'}(\varepsilon_1, \varepsilon_2)$, where (j', k') is the pair of messages in the period; (b) if $j \neq j'$ and $k = k'$ (resp. $j = j'$ and $k \neq k'$), the next phase is $R^{j'k'}(-\bar{\varepsilon}, \varepsilon_2)$ (resp. $R^{j'k'}(\varepsilon_1, -\bar{\varepsilon})$). In words, unilateral deviations from the prescribed action profile trigger a punishment phase, while inconsistencies in successive reports are punished via the payoff prescribed by the regular phase. Simultaneous deviations are ignored.

(ii) *From a punishment phase*: Without loss of generality, consider P_1^k, where k is player 2's report at the end of the last period before the punishment phase (so k is fixed throughout P_1^k). In what follows, all statements to histories and periods refer to the partial histories starting at the beginning of the punishment phase. Given $\widehat{s}_2^k$, define $H^k \subseteq H^T$ as the set of histories of length at most T for which there exists a (arbitrary) strategy s_1 of player 1 such that this history is on the equilibrium path for s_1 and $\widehat{s}_2^k$ as far as actions are concerned. That is, a history is not in H^k if and only if, in some period, the action of player 2 is inconsistent with $\widehat{s}_2^k$.

If $h^t \in H^k$ but $h^{t+1} \notin H^k$, the punishment phase stops at the end of period $t+1$ and the punishment phase $P_2^{j'}$ starts, where j' is player 1's report in period

$t+1$. Otherwise, the punishment phase continues up to the Tth period, and we henceforth let h denote such a history of length T. Let (j', k') denote the pair of reports in the last period of the punishment phase.

The next phase is then $R^{j'k'}(\varepsilon_1(h; P_1^k), \varepsilon_2(h; P_1^k))$ with $\varepsilon_1(h; P_1^k) \in [-\bar{\varepsilon}, 0]$ and $\varepsilon_1(h; P_1^k) = -\bar{\varepsilon}$ if $k' = k$, and

PROPERTY 1: $\varepsilon_1(h; P_1^k)$ is such that, if $k' \neq k$, playing the action specified in the punishment phase is optimal for player 1 along every history $h \in H^k$ under the state of the world (j', k') (recall that h specifies (j', k')).[9]

Inequality (4) guarantees that the variation of $\varepsilon_1(h; P_1^k)$ across histories h that is required is less than $2\bar{\varepsilon}/3$, so that this can be done with $\varepsilon_1(h; P_1^k)$ in $[-\bar{\varepsilon}, 0]$ for all histories h. As for $\varepsilon_2(h; P_1^k)$, it is in $[\bar{\varepsilon}/3, \bar{\varepsilon}]$ if $k' = k$ and in $[-\bar{\varepsilon}, -\bar{\varepsilon}/3]$ otherwise. Furthermore,

PROPERTY 2: $\varepsilon_2(\cdot; P_1^k)$ is such that, conditional on state (j', k') and after every history $h' \in H^k$ within the punishment phase, player 2 is indifferent over all sequences over action profiles (within the punishment phase) consistent with H^k, and prefers those to all others.

Given (4), this is possible whether $k' = k$ or not.

It is clear that the strategy profile yields the pair of payoff arrays $v = (v_1, v_2)$. It is equally clear that play is specified in a way that is independent of beliefs.

Verification That the Described Strategy Profile Is a Perfect Bayesian Equilibrium

Regular Phase $R^{jk}(\varepsilon_1, \varepsilon_2)$:

(i) Actions: Suppose that one player, say player 1, unilaterally deviates from the prescribed action profile. Then the punishment phase $P_1^{k'}$ starts, where k' is the announcement by player 2. Accordingly, the payoff from deviating is at most equal to the right-hand side of (3), while the payoff from following the prescribed strategy is at least the left-hand side of (3). The result follows.

(ii) Messages: (a) Assume first that player 1 has deviated from the recommended action profile, while player 2 has not. Because player 2 will correctly report the column k at the end of the punishment phase $P_1^{k'}$ that starts, he will get at most $(1-\delta^T)M + \delta^T(v_i^{jk} - \bar{\varepsilon}/3)$ by announcing $k' \neq k$, while he gets at least $-(1-\delta^T)M + \delta^T(v_i^{jk} + \bar{\varepsilon}/3)$ if he announces $k' = k$, so that player 2 has a strict incentive to report truthfully given (4). Given that player 1 has deviated, player 1's report is irrelevant, and so it is also optimal

[9]See Hörner and Olszewski (2006) for the details of an analogous specification.

for player 1 to report truthfully; (b) Otherwise, if player i (say player 2) reports the true state, he gets at least $v_i^{jk} - \bar{\varepsilon}$, while if he misreports, he gets at most $(1-\delta)\max_{k'} v_i^{jk}(R^{jk'}(\bar{\varepsilon}, \bar{\varepsilon})) + \delta(v_i^{jk} - \bar{\varepsilon})$. Therefore, (2) guarantees that neither player has an incentive to deviate. Note that whenever player i's reports contradict his previous report, his continuation payoff is at most $v_i^{jk} - \bar{\varepsilon}$, ensuring that no player benefits from misreporting his type.[10]

Punishment Phase: Without loss of generality, consider P_1^k.

(i) Messages: Observe first that all the messages in the punishment phase are irrelevant except in the last period of this punishment phase, whether this occurs after T periods or before. If such a history belongs to H^k, then truthful announcements are optimal because of (2), as in case (ii)(b) above. If such a history does not belong to H^k, then truthful announcements are also optimal as the situation is identical to the one described just above (case (ii)(a)).

(ii) Actions: The inequality (4) (for $i = 2$) along with Property 2 ensures that player 2 has no incentive to take an action outside of the support of the (possibly mixed) action specified by $\widehat{s}_2^k$ after every history $h \in H^k$ and that he is indifferent over all the actions within this support (whether his report k is correct or not). As for player 1, by definition his strategy is optimal in case k is the true column, and Property 1 guarantees that it remains optimal to play according to s_1^{jk} in state (j, k'), for all j, k'.

3.4. *Existence*

Strict individual rationality and incentive compatibility are stringent restrictions, implying that the set of belief-free individually rational payoffs is empty for some games. In the following, we discuss two examples in which the set V^* is empty and we provide two conditions ensuring nonemptiness. In Example 2 there is no feasible payoff that is individually rational for both players simultaneously. In Example 3 both the set of individually rational payoffs and the set of incentive compatible payoffs are nonempty, but their intersection is empty.

EXAMPLE 2—Nonexistence of Belief-Free Individually Rational Payoffs: Player 1 is informed of the true state (= the row); player 2 is not ($J = 2$, $K = 1$). The payoffs are either

	L	R
U	$10, -4$	$1, 1$
D	$1, 1$	$0, 0$

[10]Otherwise a player could profit from misreporting his type at the beginning of a punishment phase and in the next regular phase, and reverting to truthtelling afterward.

or

$$\begin{array}{c|c|c|} & L & R \\ \hline U & 0,0 & 1,1 \\ \hline D & 1,1 & 10,-4 \\ \hline \end{array}.$$

In each state, player 2 must be guaranteed at least 0 in a belief-free equilibrium: his equilibrium strategy must be optimal given any beliefs he may hold, including degenerate beliefs on the true state. His payoff must therefore be at least as large as his minmax payoff given the true state, which exceeds 0 in both states. This implies that the action profile yielding -4 to player 2 cannot be played more than a fifth of the time in equilibrium. Equivalently, this means that player 1's equilibrium payoff is at most $14/5$ in each state. However, if player 1 randomizes equally between U and D independently of the state, he is guaranteed to get at least 3 in one of the states, a contradiction. (This state typically depends on player 2's strategy. However, no strategy of player 2 can bring down player 1's payoff below 3 in both states simultaneously.)

In Example 2, player 2's payoff matrix depends on player 1's type. In a belief-free equilibrium, player 2 must get at least what he can guarantee when he knows player's 1 type. In this example, this is only possible, for all beliefs of player 2, if the equilibrium is separating, that is, if in equilibrium player 1 reveals his information. However, in this example a nonrevealing strategy yields a higher payoff to player 1 than any separating outcome that is individually rational for player 2, and so no belief-free equilibrium exists. This does not arise when the uninformed player does not need to know the state to secure his individually rational payoff. This gives rise to the following condition that guarantees that V^* is nonempty.

CONDITION 3: Consider a game of one-sided incomplete information in which player 1 is informed. If there exist $\alpha_2^* \in \triangle A_2$ and $\alpha_1^j \in B_1^j(\alpha_2^*)$ such that, for all $j = 1, \ldots, J$,

$$u_2^j(\alpha_1^j, \alpha_2^*) \geq \underline{u}_2^j,$$

then V^* is nonempty.

In fact, the above inequality implies that the payoffs obtained if player 2 plays α_2^* and player 1 plays his type-dependent best reply α_1^j are individually rational for both players as well as incentive compatible for player 1.

When the uninformed player always knows his own payoff, the strategy guaranteeing him his minmax payoff is independent of the informed player's type. Thus, Condition 3 always holds in games of one-sided incomplete information with known-own payoffs.[11] This is the main class of games examined in the literature on reputations. See Section 4.

The next example shows that, with two-sided incomplete information, known-own payoffs is not a sufficient condition for V^* to be nonempty.[12]

EXAMPLE 3—Nonexistence of Individually Rational and Incentive Compatible Payoffs: Each player is informed of his own payoffs. Player 1's payoff is

$$\begin{array}{c|c|c|}
 & L & R \\ \hline
T & 3 & 0 \\ \hline
B & 0 & 1 \\ \hline
\end{array}
\quad \text{or} \quad
\begin{array}{c|c|c|}
 & L & R \\ \hline
T & 1+\varepsilon & 1 \\ \hline
B & 0 & 0 \\ \hline
\end{array}$$

for $j = 1$ and $j = 2$, respectively. Player 2's payoff is

$$\begin{array}{c|c|c|}
 & L & R \\ \hline
T & 1 & 0 \\ \hline
B & 0 & 3 \\ \hline
\end{array}
\quad \text{or} \quad
\begin{array}{c|c|c|}
 & L & R \\ \hline
T & 0 & 1 \\ \hline
B & 0 & 1+\varepsilon \\ \hline
\end{array}$$

for $k = 1$ and $k = 2$, respectively, where $\varepsilon \in (0, 1/35)$.

Consider state $(j, k) = (2, 1)$. Player 1 can secure a payoff of at least 1. This requires that, in equilibrium, action T is used with frequency not smaller than $1 - \varepsilon/(1+\varepsilon)$. Player 2 can guarantee $3/4$, but as action profile $\{R, B\}$ cannot be played more than $\varepsilon/(1+\varepsilon)$ of the time, it follows that $\Pr\{T, L|(j, k) = (2, 1)\} > 3/4 - 3\varepsilon/(1 + \varepsilon)$. Applying a symmetric argument to player 2, we obtain that $\Pr\{B, R|(j, k) = (1, 2)\} > 3/4 - 3\varepsilon/(1 + \varepsilon)$. Consider now state $(j, k) = (1, 1)$. Each player may pretend that he is of type 2, so that his preferred outcome occurs at least $3/4 - 3\varepsilon/(1 + \varepsilon)$ of the time. Thus, the incentive compatibility constraints for player 1 and for player 2 in state $(1, 1)$ require that

$$3\Pr\{T, L|(j,k) = (1,1)\} + \bigl(1 - \Pr\{T, L|(j,k) = (1,1)\}\bigr) \geq 3\left(\frac{3}{4} - 3\frac{\varepsilon}{1+\varepsilon}\right)$$

and

$$\Pr\{T, L|(j,k) = (1,1)\} + 3\bigl(1 - \Pr\{T, L|(j,k) = (1,1)\}\bigr) \geq 3\left(\frac{3}{4} - 3\frac{\varepsilon}{1+\varepsilon}\right),$$

[11] For this class of games, Shalev (1994) showed that V^* is nonempty (Proposition 5, p. 253).

[12] Example 3 is inspired by Example 6.6 in Koren (1992) that establishes that Nash equilibria need not exist in undiscounted games with two-sided incomplete information.

respectively. However, there is no value of $\Pr\{T, L|(j,k) = (1,1)\}$ that satisfies both inequalities for $\varepsilon < 1/35$.

In Example 3, known-own payoffs guarantee that the set of individually rational payoffs is nonempty. Still, none of these payoff arrays is incentive compatible. This issue does not arise when there exists a distribution over action profiles that yields individually rational payoffs independently of the state. More formally, let $\alpha \in \triangle A$ be a distribution over action profiles and let $u_i^{jk}(\alpha)$ be player i's payoff in state (j,k) under the distribution α. Let $u_i(\alpha) := (u_i^{11}(\alpha), \dots, u_i^{JK}(\alpha))$.

CONDITION 4: If there exist $\alpha^* \in \triangle A$, $\widehat{\alpha}_1^j \in \triangle A_1$, $j = 1, 2, \dots, J$, and $\widehat{\alpha}_2^k \in \triangle A_2$, $k = 1, 2, \dots, K$, such that, for all (j,k) and i,

$$u_1^{jk}(\alpha^*) \geq u_1^{jk}(B_1^{jk}(\widehat{\alpha}_2^k), \widehat{\alpha}_2^k)$$

and

$$u_2^{jk}(\alpha^*) \geq u_2^{jk}(B_2^{jk}(\widehat{\alpha}_1^j), \widehat{\alpha}_1^j),$$

then V^* is nonempty.

The payoff array $(u_1(\alpha^*), u_2(\alpha^*))$ is obviously incentive compatible, since it is achieved by strategies that do not depend on players' types. The existence of "punishment" strategies $\widehat{\alpha}_1^j$, $\widehat{\alpha}_2^k$ that are independent of the other player's type guarantees that $(u_1(\alpha^*), u_2(\alpha^*))$ is individually rational.[13] Condition 3 relied on the existence of a strategy that secured a player his minmax payoff independently of the state. Condition 4 relies on the existence of a strategy that drives down an opponent's payoff below some target level independently of the state.

Condition 4 can be further simplified when a player punishment strategy is state-independent, that is, $\widehat{\alpha}_1^j = \widehat{\alpha}_1$ and $\widehat{\alpha}_2^k = \widehat{\alpha}_2$ for all states jk. This is the case in a variety of games commonly used in economics. For instance, most auction formats considered in the literature (including affiliated values, auctions with synergies, and multiunit auctions) satisfy it provided that the range of allowable bids includes the range of possible values of the units. In this case, any distribution α^* for which the winning price is below the lowest possible value and each bidder wins the auction with positive probability guarantees each player a positive payoff, while any punishment strategy $\widehat{\alpha}_i$ that sets a bid no smaller than the largest possible value drives player $-i$'s payoff to zero. Similar reasoning applies to Bertrand games and Cournot games provided that for some output range the market price is commonly known to exceed production cost. Thus,

[13]We thank a referee for pointing out that punishment strategies $\widehat{\alpha}$ can vary with the punishing player's type.

on the one hand, there always exists a way of sharing the market such that both players achieve a positive profit. On the other hand, each player can minmax his opponent by setting a low price in a Bertrand game or a high quantity in Cournot games, independently of the state.

So far, we have focused on conditions guaranteeing that V^* is nonempty. Yet Theorem 1 asserts the existence of belief-free equilibria only for payoffs in the interior of V^*. Focusing on the interior of V^* guarantees that it is possible to provide incentives for players to carry out punishments, as in standard proofs of folk theorems with perfect monitoring and complete information, and three or more players (see Fudenberg and Maskin (1986)). This may or may not be possible otherwise.

There are games for which the set V^* is nonempty, but its interior is empty. The problem may lie with individual rationality. For instance, V^* has empty interior in zero-sum games or in games in which a player has a strictly dominant action yielding a payoff independent of the opponent's action—the payoff corresponding to a Stackelberg type. We are not aware of any simple condition ensuring that strictly individually rational payoffs exist.

On the other hand, the problem may lie with weak vs. strict incentive compatibility. Recall that weakly incentive compatible payoffs always exist, and suppose that some weakly incentive compatible payoff is strictly individually rational. We may then as well assume that the corresponding (distribution over) action profile(s) $\alpha \in (\triangle A)^{JK}$ is completely mixed.[14] Strict incentive compatibility is equivalent to $KJ(J-1)$ and $JK(K-1)$ linear inequalities, corresponding to player 1 and player 2, respectively. Given that incentive compatibility constraints only depend on differences in the distributions of outcomes corresponding to different reports, there are $(JK-1)$ distributions that can be chosen to find (strictly) incentive compatible payoffs. Thus, generically, this is possible if $(|A|-1)(JK-1)$ is at least as large as the number of constraints, $JK(J+K-2)$. Observe that $(J+K-1)(JK-1)-JK(J+K-2)=(J-1)(K-1)$. Therefore, a sufficient condition ensuring that, for a generic payoff matrix, V^* has nonempty interior whenever there exists some strictly individually rational, weakly incentive compatible payoff v, is

$$|A| \geq J+K.$$

If V^* is nonempty, but its interior is empty, belief-free equilibrium may or may not exist. For instance, in strictly dominant action games with a unique Stackelberg type—a class of games examined in the literature on reputations—a belief-free equilibrium always exist, although the interior of V^* is empty.

[14]To see this, observe that any state-independent action profile is weakly individually rational. Pick any such completely mixed action profile α' and consider the convex combination $\varepsilon\alpha' + (1-\varepsilon)\alpha$, $\varepsilon \in [0,1]$. Since the set of incentive compatible action profiles is convex, this linear combination is weakly incentive compatible, is completely mixed, and is strictly individually rational for small enough ε.

4. REPUTATIONS

We consider games with known-own payoffs and one-sided incomplete information. By Proposition 5 of Shalev (1994), the set V^* is nonempty in such games, and we restrict attention for now to games in which this set has nonempty interior, which guarantees that belief-free equilibria exist. Player 1 is the informed player, while player 2 is uninformed. We fix one payoff type of player 1—the *rational* type—and study how the lower bound on the limit of equilibrium payoffs as the discount factor tends to 1 varies with the addition of (finitely many) other payoff types. The supremum of this lower bound over these payoff types is called the *reputation payoff*.

Given some action $\alpha_1 \in \triangle A_1$ of the informed player, recall that $B(\alpha_1)$ is the set of best replies of player 2. The rational payoff type is denoted u_1. When considering two types only, we write u_1' for player 1's other payoff type.

The analysis of reputation is strikingly simple. Observe that if some other type u_1' is present, the rational type's payoff must be at least

$$\min_{\alpha \in \triangle A} u_1(\alpha) \quad \text{such that} \quad u_2(\alpha) \geq \underline{u}_2, \quad u_1'(\alpha) \geq \underline{u}_1'.$$

Indeed, player 2's strategy must be optimal if he assigns probability 1 to player 1's other (nonrational) type, so that the distribution over action profiles induced by player 1's other type against player 2 must be individually rational for both players. Yet player 1's rational type may mimic the other type. The supremum over u_1' of this expression gives then a lower bound on the reputation payoff. (Introducing more than two types can only increase this lower bound.) The dual problem is

$$\sup_{u_1', p \geq 0, q \geq 0} \quad p\underline{u}_2 + q\underline{u}_1' \quad \text{such that} \quad pu_2 + qu_1' \leq u_1.$$

Since the constraints can be taken to be binding, the reputation payoff is at least

$$\sup_{p \geq 0} \mathrm{val}(u_1 - p(u_2 - \underline{u}_2 1)),$$

where 1 is an $|A_1| \times |A_2|$ matrix with 1s as entries. This is the bound found for Nash equilibrium payoffs in the undiscounted case by Israeli (1999, Theorem 1) using Farkas' lemma. His proof shows that it is tight and achieved by $u_1' = -u_2$.[15] Since there usually is a trade-off between punishing player 1's rational type and his other type, punishing player 1's rational type might give his other type a payoff above his minmax. But if the other type's preferences are opposite to player 2's, player 2's payoff is below his minmax, a contradiction:

[15] As mentioned, zero-sum games have been ruled out by the assumption $\mathrm{int}\, V^* \neq \emptyset$. Nevertheless, there exist payoff types arbitrarily close to $u_1' = -u_2$ for which the assumption is satisfied, so that Israeli's analysis applies. See the online supplemental material (Hörner and Lovo (2009)).

By maximizing his payoff, player 2 minimizes the other type's payoff, implying that the rational type's payoff is high. Note that the reputation payoff is the lower bound on belief-free equilibrium payoffs, which may be higher than that of Nash or sequential equilibrium payoffs.

A standard concept in the analysis of reputations is the Stackelberg payoff, introduced by Fudenberg and Levine (1989).

DEFINITION 2: The Stackelberg payoff u_1^* is defined as

$$\sup_{\alpha_1 \in \triangle A_1} \min_{\alpha_2 \in B(\alpha_1)} u_1(\alpha_1, \alpha_2).$$

A sequence achieving the supremum is a Stackelberg sequence and its limit is a Stackelberg action.

We say that a reputation is *possible* in a given game if there exists some type u_1' such that, in all (belief-free) equilibria of the game, the rational type secures a payoff strictly above the minmax payoff. We also introduce a particular class of games (see, e.g., Schmidt (1993)).[16]

DEFINITION 3: A game has *conflicting interest* if some Stackelberg sequence minmaxes player 2.

In other words, in a game of conflicting interests, there exists a Stackelberg sequence $\{\alpha_1^n\}$ such that

$$\max_{\alpha_2 \in \triangle A_2} u_2(\alpha_1^n, \alpha_2) = \underline{u}_2.$$

THEOREM 2: *Fix a game of one-sided incomplete information with known-own payoffs in which player* 1 *is the informed player.*

(i) *The reputation payoff is equal to*

$$\sup_{\alpha_1 \in \triangle A_1} \min_{\alpha_2 : u_2(\alpha_1, \alpha_2) \geq \underline{u}_2} u_1(\alpha_1, \alpha_2).$$

(ii) *A reputation is possible if and only if, for some* $\alpha_1 \in \triangle A_1$,

$$\forall \alpha_2 \in \triangle A_2, \quad u_2(\alpha_1, \alpha_2) \geq \underline{u}_2 \quad \Rightarrow \quad u_1(\alpha_1, \alpha_2) > \underline{u}_1.$$

(iii) *The reputation and Stackelberg payoffs are equal if and only if, for any* $n \in \mathbb{N}$, *there exists* $\alpha_1^n \in \triangle A_1$ *such that*

$$\forall \alpha_2 \in \triangle A_2, \quad u_2(\alpha_1^n, \alpha_2) \geq \underline{u}_2 \quad \Rightarrow \quad u_1(\alpha_1^n, \alpha_2) \geq u_1^* - 1/n.$$

[16]Our definition is slightly stronger than the usual one, as minmaxing must occur along the sequence. If the supremum is a maximum, then one can take the constant sequence $\{\alpha_1\}$ and the definition coincides with Schmidt (1993). We thank a referee for an illuminating example.

This includes all games of conflicting interest.

The first conclusion is due to Israeli (1999). The second conclusion of the theorem follows immediately from the first. If a sequence $\{\alpha_1^n\}$ satisfying the condition of the third conclusion exists, then this sequence guarantees that the reputation payoff is at least as large as the Stackelberg payoff. Conversely, if the reputation payoff equals the Stackelberg payoff, then by definition of the reputation payoff, there must exist a sequence satisfying this condition. As for the last statement, observe that from the definition of a game of conflicting interest, given any term α_1^n of a Stackelberg sequence, the set of best replies to α_1^n and the set of individually rational actions for player 2 coincide. Therefore, plugging this Stackelberg sequence into the definition of the reputation payoff, it follows that the reputation payoff must be at least as large as, and therefore equal to, the Stackelberg payoff.[17]

Note that the Stackelberg payoff may or may not exceed the minmax payoff. That is, while the second conclusion characterizes when the reputation payoff exceeds the minmax payoff, the third makes no claim regarding the level of the reputation payoff.

A few more remarks are in order.

- Reputation may or may not be possible in games of common interest (Aumann and Sorin (1989)). This should not be surprising, since we allow for mixed strategies and, more importantly, incomplete information pertains to payoffs, not to the complexity of strategies.
- The theorem is reminiscent of results for Nash equilibrium payoffs with unequal discount factors. Schmidt (1993) showed that the Stackelberg payoff and the reputation payoff coincide in games of conflicting interests, when player 1 is sufficiently more patient than player 2 and both discount factors tend to 1. Cripps, Schmidt, and Thomas (1996) generalized this result by showing that the reputation payoff is as given in the first conclusion of the theorem and they provided an example that shows that the result is false with equal discounting. In the case of equal discounting, more severe restrictions are thus required to obtain reputational effects with Nash equilibria. Cripps, Dekel, and Pesendorfer (2005) showed that the Stackelberg payoff can be achieved when attention is restricted to a subclass of games with conflicting interest, namely games of strictly conflicting interest. It should come as no surprise that, unlike in the case of Nash equilibria, it is not necessary that player 1 be more patient than player 2 here. After all, the uninformed player must play a best reply to all possible beliefs. This alleviates the need for the informed player to build a reputation, which may be a costly enterprise, before enjoying it. Indeed, given that the general characterization of belief-free equilibrium payoffs is similar to the characterization of Nash

[17]We thank both referees for pointing this out, correcting an erroneous statement made in an earlier version.

equilibrium payoffs with no discounting (at least in the case of one-sided incomplete information), it is natural that our findings regarding reputations parallel those of Cripps and Thomas (1995) for the case of no discounting.

- Chan (2000) established that in strictly dominant action games (games in which player 1 has a strictly dominant action, and player 2's best reply yields the highest possible individually rational payoff to player 1), the rational type receives the Stackelberg payoff in any sequential equilibrium when the game is perturbed by adding a single commitment type who always play the Stackelberg action. See, for instance, the game in Figure 3. The reader may wonder how this result is consistent with our analysis, since a strictly dominant action game need not be a game of conflicting interest. Recall that we assumed so far that $\operatorname{int} V^* \neq \emptyset$, which rules out games in which the set of feasible and individually rational payoffs has empty interior. That is, we have excluded commitment types, as they correspond to payoff types with a dominant action whose payoff is independent of player 2's action.[18] If the game is perturbed by adding a single commitment type who always plays the Stackelberg action, a belief-free equilibrium exists if and only if there exists an action α_1 of player 1 such that (α_1, a_2) is a Nash equilibrium in the (stage) game of complete information between player 1's rational type and player 2, where a_2 is 2's best reply to the Stackelberg action. This condition is satisfied in strictly dominant action games and, indeed, it is then immediate that player 1's rational type secures his Stackelberg payoff in the belief-free equilibrium of any such game: Since player 2's strategy must be a best reply to all possible beliefs, including those which assign probability 1 to the commitment type, he must play a_2 in every period, and player 1's best reply is then to play his Stackelberg action. Observe, however, that reputation is fragile in such games: Consider replacing the single commitment type by any payoff type arbitrarily close to the commitment type, but for whom the dominant action does not yield a payoff independent of player 2's action. Then, according to the previous theorem, to determine the reputation payoff, we must minimize player 1's payoff from his Stackelberg action over player 2's individually rational actions rather than over his best replies only. In the example of Figure 3, this implies that for all nearby payoff types, the reputation payoff is $4/3$—still strictly above the minmax payoff of 0, so that a reputation is indeed possible, but below the Stackelberg payoff of 2. Since belief-free equilibria are sequential equilibria, this implies that reputation in strictly dominant action games is also fragile with respect to sequential equilibria.[19] In contrast, the reputational effects obtained by Cripps, Dekel,

[18]More precisely, player 1 has a dominant strategy in the repeated game for all discount factors if and only if he has a dominant strategy in the stage game yielding a payoff that is independent of player 2's action.

[19]Note, however, that this nongenericity is only with respect to the limit payoff set as the discount factor tends to 1. For a fixed discount factor, the Stackelberg action is a strictly dominant action in the supergame for payoff types sufficiently close to the Stackelberg type.

	L	R
T	$2,1$	$0,0$
B	$0,0$	$1,2$

FIGURE 3.—A strictly dominant action game.

and Pesendorfer (2005) in strictly conflicting interest games are robust, as the reputation payoff is continuous in the payoff parameters (as long as $\operatorname{int} V^* \neq \emptyset$).

- The result also sheds some light on the possible nonexistence of belief-free equilibria in games with two-sided incomplete information and known-own payoffs. Indeed, following the same logic, each player should be able to secure his reputation payoff in that case. However, nothing guarantees in general that it is feasible for both players to simultaneously achieve their reputation payoff.

5. CONCLUDING REMARKS

This paper has introduced a solution concept for two-player repeated games of incomplete information and has characterized the corresponding payoff set as the discount factor tends to 1. This characterization is simple. Payoffs must be individually rational (in the sense of Blackwell (1956)) and must correspond to probability distributions over action profiles that are incentive compatible, given the private information of each player. The relevance and effectiveness of this concept has been illustrated in the context of reputations.

There are several theoretical generalizations that demand attention. The information structure that we have considered in this paper is quite stylized, if standard. More generally, a player's information can be modeled as a partition over the states of nature. Second, attention has been restricted to two players. While the appropriate generalization of the incentive compatibility conditions is quite obvious, it is less clear how to define individual rationality in the case of three players or more, as Blackwell's characterization immediately applies to the case of two players only. Such a generalization could yield interesting insights for the study of reputations with more than two players.

For economic applications, it is also of interest to extend the characterization to the case of a changing state. For instance, Athey and Bagwell (2001) and Athey, Bagwell, and Sanchirico (2004) characterized the (perfect public) equilibrium payoffs of a repeated game between price-setting oligopolists whose marginal cost in each period is private information. These costs are assumed to be drawn independently across players and over time, according to some commonly known distribution. In this context, we may wish to know which of these payoffs remain equilibrium payoffs if all that firms know is that the costs are independent and identically distributed, but the underlying distributions are unknown. In this way, the concept of belief-free equilibrium may turn out to provide a useful robustness criterion in this literature.

APPENDIX: PROOF OF THEOREM 1

We first explain the construction without explicit communication, but with a randomization device. Communication is replaced by choices of actions, but since the set of actions may be smaller than the set of states, it may be necessary to use several periods to report types. We let $c-1$ denote the smallest such number given the number of states and actions, that is, c is the smallest integer such that $|A_1|^{c-1} \geq J$ and $|A_2|^{c-1} \geq K$ (recall that $|A_i| \geq 2$). Players will regularly report their type in rounds of c periods. For reasons that will become clear, in the last of these c periods, players have the opportunity, through the choice of a specific action, to signal that the report they have just made is incorrect.

Equilibrium Strategies

The play is again divided into phases. To guarantee that players' best replies are independent of their beliefs, even within a round of communication (especially if a player's own deviation during that round already prevents him from truthfully reporting his type), the construction must be considerably refined. For each player, we pick two specific actions from A_i, henceforth referred to as B and U. The pair of payoff arrays v is in the interior of V^* and is fixed throughout.

There are two kinds of phases. Regular phases last at most n periods and punishment phases last at most T periods, where n and T are to be specified. Regular phases are denoted $R^{jk}(\varepsilon_1, \varepsilon_2)$, where $j \in \{1, \dots, J\}$ and $k \in \{1, \dots, K\}$, or R^{xy}, where $x \in \{1, \dots, J, (L, n_1^U)\}$ and $y \in \{1, \dots, K, (L, n_2^U)\}$, with $n_i^U \in \{1, \dots, c\}$ and either $x = (L, n_1^U)$ or $y = (L, n_2^U)$, or both (L stands for "lie"). Punishment phases are denoted P_i, $i = 1, 2$. Let $\widehat{s}_2^k$ (resp. $\widehat{s}_1^j$) denote a (behavior) strategy of player 2 (resp. 1) such that player 1's (resp. player 2 's) payoff is less than $v_1^{jk} - 3\bar{\varepsilon}$ for all j and all strategies of player 1 (resp. $v_2^{jk} - 3\bar{\varepsilon}$ for all k and all strategies of player 2) for $\bar{\varepsilon} > 0$ small enough to be specified. Such strategies exist since $v \in \operatorname{int} V^*$. Further, let s_1^{jk} (resp. s_2^{jk}) denote some fixed pure best reply to $\widehat{s}_2^k$ (resp. $\widehat{s}_1^j$) given row j (resp. column k).

In several steps of the construction, a communication round of c periods takes place (within a phase). We fix a 1–1 mapping from states $\{1, \dots, J\}$ to J sequences $\{a_1^t\}_{t=1}^{c-1}$ of length $c-1$ ($a_1^t \in A_1$) and similarly fix a 1–1 mapping from states $\{1, \dots, K\}$ to K sequences $\{a_2^t\}_{t=1}^{c-1}$ of length $c-1$ ($a_2^t \in A_2$). If the play of player 1 during the first $c-1$ periods equals such a sequence and his action in period c equals B, we say that player 1 (or his play) reports the row j that maps into this sequence of actions. Similarly, if the play of player 2 during the first $c-1$ periods equals such a sequence and his action in period c equals B, we say that player 2 (or his play) reports the column k that maps into this sequence of actions. Otherwise, we say that player i (or his play) communicates (L, n_i^U), where U is the number of periods during these c periods in which player i chose action U. We shall provide incentives for player i to report the true row or column, rather than report (L, n_i^U) for any n_i^U, and to report (L, n_i^U) for any $n_i^U \geq 0$, rather than the incorrect row or column. Further, we provide incentives for player i to maximize this number n_i^U as soon as his sequence of actions $\{a_1^t\}_{t=1}^{\tau}$, $\tau \leq c-1$, is inconsistent with any of the sequences that the mapping maps into.

Actions

(i) *Regular phase*: A regular phase lasts at most $n > c$ periods, the last c of which is a communication round. During the first $n-c$ periods, for all regular phases indexed by j, k, and true column k', play proceeds as follows:

Phase	Player 1	Player 2
$R^{j(L,n_2^U)}$	$\widehat{s}_1^j$	$s_2^{jk'}$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	$a_1^{jk}(\varepsilon_1, \varepsilon_2)$	$a_2^{jk}(\varepsilon_1, \varepsilon_2)$
$R^{(L,n_1^U)(L,n_2^U)}$	$(U, \ldots, U)$	$(U, \ldots, U)$

The specification for $R^{(L,n_1^U)k}$ is the obvious analogue to the case $R^{j(L,n_2^U)}$. The action $a^{jk}(\varepsilon_1, \varepsilon_2)$ is to be specified. The strategies $\widehat{s}_1^j$ and $s_2^{jk'}$ are the same as in the punishment phase (note that the duration is not the same, however). The superscript jk' of the expression $s_2^{jk'}$ refers to the row j that indexes the regular phase $R^{j(L,n_2^U)}$ (which need not be the true row) and to the true column k'. This specification of actions is valid as long as (in the case of $R^{jk}(\varepsilon_1, \varepsilon_2)$ or $R^{(L,n_1^U)(L,n_2^U)}$) the history within the phase is consistent with these actions or if all deviations from the specified actions during this phase were simultaneous, and as long as (in the case of $R^{j(L,n_2^U)}$) the history within the phase is consistent with $\widehat{s}_1^j$ for some arbitrary s_2: As will be specified, a punishment phase is immediately entered otherwise. During the periods $n-c+1, \ldots, n-1$ of this phase, player 1 (resp. player 2) communicates the true row j (resp. true column k); if this is impossible given his play from period $n-c$ onward, he chooses U in every remaining period.

(ii) *Punishment phase*: Without loss of generality, consider P_1, where $T > 2c$ is to be specified. In the first c periods of this phase, player 1 plays U repeatedly while player 2 reports the true column (following the protocol described above). As in the regular phase, if this is impossible given player 2's play, he chooses U in every remaining period of this communication round. In the table below, we refer to the case in which the column reported is k as the case k, while $(L, (n_1^U, n_2^U))$ refers to any other case, where n_i^U is the number of times player i chose action U in periods $1, \dots, c$. Play in periods $c+1, \dots, T-c$ is then as follows:

Phase P_1	Player 1	Player 2
k	$s_1^{j'k}$	$\widehat{s}_2^k$
$(L, (n_1^U, n_2^U))$	U	U

This specification is valid up to period $T-c$ (i) in the case $(L, (n_1^U, n_2^U))$, as long as both players have played U in all periods since period $c+1$ or all deviations have been simultaneous, or (ii) in the case k, as long as the history since period $c+1$ is consistent with $\widehat{s}_2^k$ for some strategy s_1; otherwise, a new punishment phase is immediately entered (see below). Here, j' refers to the true row privately known to player 1.

In the last c periods of a punishment phase (assuming that the specification above remained valid up to period $T-c$), a communication round takes place, that is, players report the true row and column, and as soon as they fail to do so, play U repeatedly.

(iii) *Initial phase*: In the first c periods of the game, a communication round takes place, that is, players report the true row and column, and as soon as they fail to do so, play U repeatedly. In period $c+1$, the regular phase $R^{jk}(\varepsilon_1, \varepsilon_2)$ is entered if row j and column k are reported, where $\varepsilon_i \in [-\bar{\varepsilon}, \bar{\varepsilon}]$ is chosen so that the ex ante payoff in period 1 is exactly v^{jk} conditional on j and k being the true row and column. If player 1 reports j and player 2 reports (L, n_2^U) in the first c periods, the regular phase $R^{j(L,n_2^U)}$ is entered. Similarly, if player 1 reports (L, n_1^U), whereas player 2 reports k, the regular phase $R^{(L,n_2^U)k}$ is entered. Regular phase $R^{(L,n_1^U),(L,n_2^U)}$ is entered in the remaining case.

Transitions

From a Regular Phase

We have already mentioned what happens if there is a deviation during the first $n-c$ periods of such a phase: If a player makes a unilateral deviation during the first $n-c$ periods of a regular phase $R^{jk}(\varepsilon_1, \varepsilon_2)$ or $R^{(L,n_1^U)(L,n_2^U)}$, a punishment phase starts. If player 1 (player 2) unilaterally deviates, punishment phase P_1 (resp. P_2) is immediately entered. Similarly, if player 1 (resp.

player 2) deviates from $\widehat{s}_1^j$ (resp. $\widehat{s}_2^k$) during the first $n-c$ periods of a regular phase $R^{j(L,n_2^U)}$ (resp. $R^{(L,n_2^U)k}$), the punishment phase P_1 (resp. P_2) is immediately entered. From now on, we assume without repeating it that no such deviation occurs. In all tables that follow, $j' \neq j$ and $k' \neq k$.

(i) From $R^{jk}(\varepsilon_1, \varepsilon_2)$: The new phase depends on the last c periods of the phase. Define also $\rho := 2(1-\delta)\delta^{-\max(n,T)}M$. The quantity $\tilde{\varepsilon}_i^{jk}$ will be defined shortly. We have the following transitions:

Regular Phase	During Periods $n-c+1, \ldots, n$ of the Phase, Players 1 and 2 Report	Next Regular Phase
$R^{jk}(\varepsilon_1, \varepsilon_2)$	$(L, n_1^U), (L, n_2^U)$	$R^{(L,n_1^U)(L,n_2^U)}$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	$(L, n_1^U), k$	$R^{(L,n_1^U)k}$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	$(L, n_1^U), k'$	$R^{(L,n_1^U)k'}$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	j, k'	$R^{jk'}(\varepsilon_1, -\bar{\varepsilon})$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	j', k'	$R^{j'k'}(\varepsilon_1, \varepsilon_2)$
$R^{jk}(\varepsilon_1, \varepsilon_2)$	j, k	$R^{jk}(\varepsilon_1, \varepsilon_2)$

(We omit the obvious symmetric specification for reports $j, (L, n_2^U)$ and $j', (L, n_2^U)$.)

(ii) From $R^{(L,n_1^U)k}$ (and symmetrically from $R^{j(L,n_2^U)}$): We have the following transitions:

Regular Phase	During Periods $n-c+1, \ldots, n$ of the Phase, Players 1 and 2 Report	Next Regular Phase
$R^{(L,n_1^U)k}$	$(L, n_1'^U), (L, n_2'^U)$	$R^{(L,n_1'^U)(L,n_2'^U)}$
$R^{(L,n_1^U)k}$	$(L, n_1'^U), k$	$R^{(L,n_1'^U)k}$
$R^{(L,n_1^U)k}$	$(L, n_1'^U), k'$	$R^{(L,n_1'^U)k'}$
$R^{(L,n_1^U)k}$	$j, (L, n_2'^U)$	$R^{j(L,n_2'^U)}$
$R^{(L,n_1^U)k}$	j, k	$R^{jk}(\tilde{\varepsilon}_1^{jk} + \rho n_1^U, \varepsilon_2^{k;k}(h))$
$R^{(L,n_1^U)k}$	j, k'	$R^{jk'}(\varepsilon_1^{k;k'}(h) + \rho n_1^U, \varepsilon_2^{k;k'}(h))$

Here $\varepsilon_2^{k;k}(h) \in [3\bar{\varepsilon}/4, \bar{\varepsilon}]$, $\varepsilon_2^{k;k'}(h) \in [-\bar{\varepsilon}/2, -\bar{\varepsilon}/4]$, and $\varepsilon_1^{k;k'}(h) \in [-\bar{\varepsilon}, \bar{\varepsilon}]$ are computed as follows: $\varepsilon_2^{k;k}(\cdot)$ makes player 2 precisely indifferent over all histories h that are consistent with $\widehat{s}_2^k$, conditional on the true column being k; $\varepsilon_2^{k;k'}(\cdot)$ makes player 2 precisely indifferent over all histories h that are consistent with $\widehat{s}_2^k$, conditional on the true column being k', while $\varepsilon_1^{k;k'}(h)$ compensates player 1 for every period along h in which the action he took is the action specified by s_1^{jk}, so as to make sure that playing this action is optimal, conditional on the true state being (j, k') (reported in the last c periods).

(iii) Finally, from $R^{(L,n_1^U)(L,n_2^U)}$: We have the following transitions:

Regular Phase	During Periods $n-c+1,\dots,n$ of the Phase, Players 1 and 2 Report	Next Regular Phase
$R^{(L,n_1^U)(L,n_2^U)}$	$(L,n_1'^U),(L,n_2'^U)$	$R^{(L,n_1'^U)(L,n_2'^U)}$
$R^{(L,n_1^U)(L,n_2^U)}$	$(L,n_1'^U),k$	$R^{(L,n_1'^U)k}$
$R^{(L,n_1^U)(L,n_2^U)}$	j,k	$R^{jk}(\rho n_1^U,\rho n_2^U)$

(We omit the obvious symmetric specification for reports j, (L, n_2^U).)

From a Punishment Phase

Without loss of generality, consider P_1. We have already briefly mentioned what happens if there is a deviation during the periods $c+1,\dots,T-c$ of such a phase; in case $(L,(n_1^U,n_2^U))$, if player i unilaterally deviates from the play of U, the punishment phase P_i is immediately entered; in case k, if player 2 deviates from the support of the (possibly mixed) action specified by $\widehat{s}_2^k$, punishment phase P_2 is entered (no matter how player 1 has played). From now on, we assume without repeating it that no such deviation occurs up to period $T-c$. In case k, let h denote the history during the periods $c+1,\dots,T-c$.

(i) In case k, we observe the following transitions:

Punishment Phase P_1	During Periods $T-c+1,\dots,T$ of the Phase, Players 1 and 2 Report	Next Regular Phase
Case k	$(L,n_1'^U),(L,n_2'^U)$	$R^{(L,n_1'^U)(L,n_2'^U)}$
Case k	$(L,n_1'^U),k$	$R^{(L,n_1'^U)k}$
Case k	$(L,n_1'^U),k'$	$R^{(L,n_1'^U)k'}$
Case k	j,k	$R^{jk}(\rho n_1^U-\bar{\varepsilon},\varepsilon_2^{k;k}(h))$
Case k	j,k'	$R^{jk'}(\varepsilon_1^{k;k'}(h),\varepsilon_2^{k;k'}(h))$
Case k	$j,(L,n_2'^U)$	$R^{j(L,n_2'^U)}$

(ii) In case $(L,(n_1^U,n_2^U))$, the transitions are described by the next table.

It is clear from this specification that the strategy profile described here is belief-free, since actions are always determined by the history and possibly by a player's own type (in case he is minmaxed), but not on his beliefs about his opponent's type.

Punishment Phase P_1	During Periods $T-c+1, \ldots, T$ of the Phase, Players 1 and 2 Report	Next Regular Phase
Case $(L, (n_1^U, n_2^U))$	$(L, n_1'^U), (L, n_2'^U)$	$R^{(L,n_1'^U)(L,n_2'^U)}$
Case $(L, (n_1^U, n_2^U))$	$(L, n_1'^U), k$	$R^{(L,n_1'^U)k}$
Case $(L, (n_1^U, n_2^U))$	$j, (L, n_2'^U)$	$R^{j(L,n_2'^U)}$
Case $(L, (n_1^U, n_2^U))$	j, k	$R^{jk}(\rho n_1^U, \rho n_2^U)$

Specification of $\bar{\varepsilon}$, $a_1^{jk}(\varepsilon_1, \varepsilon_2)$, δ, T, n, and $\tilde{\varepsilon}_i^{jk}$:

Since v is in the interior of V^*, it is possible to find $\bar{\varepsilon} > 0$, as well as, for all $(\varepsilon_1, \varepsilon_2), (\varepsilon_1', \varepsilon_2') \in [-2\bar{\varepsilon}, 2\bar{\varepsilon}]$, probability distributions over A, $\Pr\{\cdot \mid R^{jk}(\varepsilon_1, \varepsilon_2)\}$, such that for all j, k, j', k', and $i = 1, 2$, defining

$$v_i^{jk}(R^{j',k'}(\varepsilon_1, \varepsilon_2)) := \sum_{a \in A} \Pr\{a \mid R^{j'k'}(\varepsilon_1, \varepsilon_2)\} u_i^{jk}(a),$$

it is the case that, for $j' \neq j$ and $k' \neq k$,

$$\text{(A1)} \qquad \begin{aligned} &v_1^{jk}(R^{j,k}(\varepsilon_1, \varepsilon_2)) > v_1^{jk}(R^{j'k}(\varepsilon_1', \varepsilon_2')) \quad \text{and} \\ &v_2^{jk}(R^{j,k}(\varepsilon_1, \varepsilon_2)) > v_2^{jk}(R^{jk'}(\varepsilon_1', \varepsilon_2')). \end{aligned}$$

Furthermore, if $\{a_1^t\}_{t=1}^c$ and $\{a_2^t\}_{t=1}^c$ are the sequences corresponding to reports j and k, for all δ close enough to 1 and n large enough, we can pick those distributions so that player i's average discounted payoff under state (j, k) from the sequence $\{a_1^t, a_2^t\}_{t=1}^c$ followed by $n - c$ repetitions of the action profile determined by $\Pr\{a \mid R^{jk}(\varepsilon_1, \varepsilon_2)\}$ is exactly equal to $v_i^{jk} + \varepsilon_i$. Observe that in the equilibrium described above, all values of ε_i are in $[-\bar{\varepsilon}, \bar{\varepsilon}]$. Furthermore, since v is in the interior of V^*, we may assume that player 1's (resp. player 2's) average discounted payoff under state (j, k) given that player 2 uses $\widehat{s}_2^k(\varepsilon)$ (resp. $\widehat{s}_1^j(\varepsilon)$) for $n - 2c$ periods, followed by any arbitrary play during c periods, is at most $v_1^{jk} + \varepsilon$ (resp. $v_2^{jk} + \varepsilon$) for $\varepsilon > -3\bar{\varepsilon}$.

Consider the inequalities

$$\text{(A2)} \qquad v_1^{jk} + \varepsilon_1 > (1 - \delta^c)M + \delta^c(1 - \delta^n)(v - 2\bar{\varepsilon}) + \delta^{n+c}(v_1^{jk} + \tilde{\varepsilon}_1^{jk} + c\rho),$$

$$\text{(A3)} \qquad v_1^{jk} + \varepsilon_1 < -(1 - \delta^{n+c})M + \delta^{n+c}(v_1^{jk} + \tilde{\varepsilon}_1^{jk}),$$

$$\text{(A4)} \qquad v_1^{jk} - \bar{\varepsilon} > (1 - \delta^c)M + \delta^c(1 - \delta^{n-c})(v_1^{jk} - 2\bar{\varepsilon}) + \delta^n(v_1^{jk} - \bar{\varepsilon}).$$

Given $\bar{\varepsilon}$, fixing δ^n, inequality (A4) is satisfied as $\delta \to 1$, provided that the value of δ^n is large enough. Similarly, given $\bar{\varepsilon}$, fixing δ^n, inequality (A2) is satisfied as $\delta \to 1$ for $\tilde{\varepsilon}_1^{jk} = -\bar{\varepsilon}$ and (A3) is satisfied for $\tilde{\varepsilon}_1^{jk} = 3\bar{\varepsilon}/4$, provided that the value of δ^n is large enough and $\varepsilon_1 < \bar{\varepsilon}/2$ (recall that $\rho = 2(1 - \delta)\delta^{-\max(n,T)}M \to 0$ for fixed $\delta^{-\max(n,T)}$). Observe that the left-hand side of (A4) is the lowest possible payoff for player 1, evaluated in the first period of a communication round, concluding either a punishment phase or a regular phase, if he reports his true row j and player 2 reports his true column k, while the right-hand side is the most he can expect by reporting another row $j' \neq j$ when player 2 reports his true column k. Similarly, the left-hand sides of (A2) and (A3) are player 1's payoff, evaluated in the first period of a communication round concluding either a punishment phase or a regular phase, if he reports his true row j and player 2 reports his true column k (and the upcoming regular phase is $R^{jk}(\varepsilon_1, \varepsilon_2)$), while the right-hand side of (A2) (resp. (A3)) is the highest (resp. lowest) payoff he can expect if he reports (L, n_1^U) for some n_1^U. Therefore, if $\varepsilon_1 < \bar{\varepsilon}/2$, by the intermediate value theorem, we can find different values of $\tilde{\varepsilon}_1^{jk} \in (-\bar{\varepsilon}, 3\bar{\varepsilon}/4)$ so that the payoff from reporting the true row exceeds the payoff from reporting (L, n_1^U) for all n_1^U, which in turn exceeds the payoff from reporting another row $j' \neq j$, provided player 2 reports the true column. If $\varepsilon_1 \geq \bar{\varepsilon}/2$, we can set $\tilde{\varepsilon}_1^{jk} = 0$: In that case as well, the same ordering obtains provided that the value of δ^n is large enough as $\delta \to 1$. The values $\tilde{\varepsilon}_2^{jk}$ are defined similarly.

Consider now the two inequalities

$$\text{(A5)} \qquad -(1-\delta^n)M + \delta^n(v_1^{jk} + 3\bar{\varepsilon}/4) > (1-\delta^n)M + \delta^n(v_1^{jk} + \rho c),$$

$$\text{(A6)} \qquad -(1-\delta^n)M + \delta^n v_1^{jk} > (1-\delta^n)M + \delta^n(v_1^{jk} - \bar{\varepsilon}).$$

Conditional on player 2 reporting (L, n_2^U) for some n_2^U, the left-hand side of (A5) is the lowest possible payoff for player 1, evaluated in the first period of a communication round concluding either a punishment phase or a regular phase, if he reports his true row j, while the right-hand side is the highest payoff he can get if he reports (L, n_1^U) for some n_1^U. Similarly, the left-hand side of (A6) is the lowest possible payoff for player 1, evaluated in the first period of a communication round concluding either a punishment phase or a regular phase, if he reports (L, n_1^U) for some n_1^U, while the right-hand side is the highest payoff he can get if he reports another row $j' \neq j$. Observe that both inequalities hold, given $\bar{\varepsilon}$, letting $\delta \to 1$, provided δ^n is large enough.

Finally, observe that the choice of ρ trivially ensures that, conditional on having started reporting (L, n_1^U) for some n_1^U, player 1 has strict incentives to play U in all remaining periods of the communication round, no matter where this round takes place.

Similar considerations hold for player 2. To summarize, we have shown that we can ensure that both players prefer to report their true type, in any communication round, than to report (L, n_i^U) for all n_i^U; that, conditional on reporting (L, n_i^U) for some n_i^U, player i has strict incentives to choose U in any remaining period of the communication round; and that they prefer to report (L, n_i^U) for any n_i^U than to report an incorrect row or column; all this, provided that δ^n (and δ^T) is fixed but large enough, by taking $\delta \to 1$, given $\bar{\varepsilon}$.

Turning now to actions, we must consider

$$\text{(A7)}\qquad (1-\delta^{n+1})M+\delta^{n+1}(1-\delta^{T-n})(v_i^{jk}-2\bar{\varepsilon})+\delta^{T+1}(v_i^{jk}-\bar{\varepsilon}) < -(1-\delta^{n})M+\delta^{n}(v_i^{jk}-\bar{\varepsilon}),$$

$$\text{(A8)}\qquad (1-\delta^{n+1})M+\delta^{n+1}(1-\delta^{T-n})(v_i^{jk}-2\bar{\varepsilon})+\delta^{T+1}(v_i^{jk}-\bar{\varepsilon}) < -(1-\delta^{T})M+\delta^{T}(v_i^{jk}-\bar{\varepsilon}/2),$$

$$\text{(A9)}\qquad \bigl(1-\delta^{\max(T,n)}\bigr)M+\delta^{\max(T,n)}(v_i^{jk}-\bar{\varepsilon}/2) < -\bigl(1-\delta^{\max(T,n)}\bigr)M+\delta^{\max(T,n)}(v_i^{jk}-\bar{\varepsilon}/4).$$

Observe that all three inequalities hold for both $i=1,2$, given $\bar{\varepsilon}$, for δ^T and n fixed as $\delta\to 1$. This ensures that, given $\bar{\varepsilon}$, we can choose n, T, and δ to satisfy all the inequalities above. As for the interpretation, (A7) ensures that player i does not want to deviate during any regular phase, (A8) ensures that player i does not want to deviate during the punishment phase P_{-i}, and (A9) ensures that we can pick $\varepsilon_2^{k;k}(\cdot)$ and $\varepsilon_2^{k;k'}(\cdot)$ within a range of values not exceeding $\bar{\varepsilon}/4$ in case k and after phases $R^{(L,n_1^u),k}$ and $R^{j,(L,n_2^u)}$. Indeed, the left-hand sides of (A7) and (A8) are the highest payoffs player i can hope for by deviating at any time (outside communication rounds), while the right-hand side of (A7) (resp. (A8)) is the lowest payoff he can expect by sticking to the equilibrium strategies in a regular phase (resp. in a punishment phase). Note that $(1-\delta^T)M$ is the highest payoff player 1 (resp. player 2) can get when using strategy $\widehat{s}_1^j$ (resp. $\widehat{s}_2^k$) during the punishment phase P_{-i} over all actions consistent with his equilibrium strategy, while $-(1-\delta^T)M$ is the lowest such payoff. Inequality (A9) guarantees therefore that there exist functions $\varepsilon_1^{k;k'}$ and $\varepsilon_1^{k;k}$ whose ranges do not exceed $\bar{\varepsilon}/4$ such that player 1 is playing a best reply, given $\varepsilon_1^{k;k'}(\cdot)$, whether or not the true column is k.

To conclude, it remains to show that public randomization can be dispensed with. Observe that this device is used in exactly one place. For all pairs (j,k), and all $(\varepsilon_1,\varepsilon_2)\in[-2\bar{\varepsilon},2\bar{\varepsilon}]$, if $\{a_1^t,a_2^t\}_{t=1}^c$ is the sequence of action profiles corresponding to the reports (j,k), and for all δ close enough to 1, the public randomization device guarantees that we can find a correlated action profile such that the average discounted payoff from the sequence $\{a_1^t,a_2^t\}_{t=1}^c$ followed by $n-c$ repetitions of this correlated action profile yields a payoff $v_i^{jk}+\varepsilon_i$ to player i, in state (j,k). Observe now that all incentives in the regular phase are strict, so that they would also be satisfied, for all δ close enough to 1, as long as the continuation payoff $\hat{v}_i^t$ in period t of the regular phase is within $2\bar{\varepsilon}+\hat{\varepsilon}$ (rather than within $2\bar{\varepsilon}$) of v_i^{jk} for some $\hat{\varepsilon}>0$ sufficiently small and all $t=1,\dots,n$. Observe now that, following Fudenberg and Maskin (1991) (which itself builds on Sorin (1986)), we can find n large enough, so that for all δ close enough to 1, there exists a sequence of sequences $\{\{a_1^t(\nu),a_2^t(\nu)\}_{t=1}^n\}_{\nu=1}^\infty$, with $\{a_1^t(\nu), a_2^t(\nu)\}_{t=1}^c = \{a_1^t, a_2^t\}_{t=1}^c$ for all ν, such that (i) the average discounted payoff from the infinite play

$$\{a_1^1(1), a_2^1(1), \dots, a_1^n(1), a_2^n(1), a_1^1(2), a_2^1(2), \dots\},$$

obtained by concatenation of the elements of this sequence, is equal to $v_i^{jk} + \varepsilon_i$, and that (ii) the continuation payoff from any period t onward in this infinite play is within $\hat{\varepsilon}$ of $v_i^{jk} + \varepsilon_i$.[20] It is then clear how to modify the specification above: Increase n and choose δ close to 1, if necessary, to guarantee the existence of such sequences; if players are in the vth consecutive regular phase $R^{jk}(\varepsilon_1, \varepsilon_2)$, with reports (j, k) that agreed in all those phases, play in that vth phase is given by $\{a_1^t(\nu), a_2^t(\nu)\}_{t=1}^n$. (Note that, in general, the continuation payoff of i at the beginning of the vth phase is not exactly $v_i^{jk} + \varepsilon_i$, so $(\varepsilon_1, \varepsilon_2)$ only refers to the continuation payoff achieved in the first such regular phase, or more precisely, from the communication phase that immediately precedes this first regular phase onward.) If a deviation occurs or consecutive reports disagree, a new such sequence of consecutive plays $\{\{a_1^t(\nu), a_2^t(\nu)\}_{t=1}^n\}_{v=1}^\infty$ starts in the next regular phase (or more precisely, from the communication phase that immediately precedes this first regular phase), given the new values of $(\varepsilon_1, \varepsilon_2)$.

REFERENCES

AGHASSI, M., AND D. BERTSIMAS (2006): "Robust Game Theory," *Mathematical Programming*, Ser. B, 107, 231–273. [455]

ATHEY, S., AND K. BAGWELL (2001): "Optimal Collusion With Private Information," *RAND Journal of Economics*, 32, 428–465. [476]

ATHEY, S., K. BAGWELL, AND C. SANCHIRICO (2004): "Collusion and Price Rigidity," *Review of Economic Studies*, 71, 317–349. [476]

AUMANN, R. J., AND M. B. MASCHLER (1995): *Repeated Games With Incomplete Information*. Cambridge, MA: The MIT Press. [456]

AUMAN, R., AND S. SORIN (1989): "Cooperation and Bounded Recall," *Games and Economic Behavior*, 1, 5–39. [474]

BAÑOS, A. (1968): "On Pseudo-Games," *Annals of Mathematical Statistics*, 39, 1932–1945. [455]

BERGEMANN, D., AND S. MORRIS (2007): "Belief Free Incomplete Information Games," Discussion Paper 1629, Cowles Foundation, Yale University. [453]

BLACKWELL, D. (1956): "An Analog of the Minmax Theorem for Vector Payoffs," *Pacific Journal of Mathematics*, 6, 1–8. [459,476]

CHAN, J. (2000): "On the Non-Existence of Reputation Effects in Two-Person Infinitely-Repeated Games," Working Paper, Johns Hopkins University. [475]

CRÉMER, J., AND R. MCLEAN (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 53, 345–361. [453]

CRIPPS, M., AND J. THOMAS (1995): "Reputation and Commitment in Two-Person Repeated Games Without Discounting," *Econometrica*, 63, 1401–1419. [475]

[20]The construction of Sorin (1986) and Fudenberg and Maskin (1991) guarantees that n and δ can be chosen independently of $v^{jk} + \varepsilon$.

——— (2003): "Some Asymptotic Results in Discounted Repeated Games of One-Sided Incomplete Information," *Mathematics of Operations Research*, 28, 433–462. [455,460]

CRIPPS, M., E. DEKEL, AND W. PESENDORFER (2005): "Reputation With Equal Discounting in Repeated Games With Strictly Conflicting Interests," *Journal of Economic Theory*, 121, 259–272. [474-476]

CRIPPS, M., K. SCHMIDT, AND J. THOMAS (1996): "Reputation in Perturbed Repeated Games," *Journal of Economic Theory*, 69, 387–410. [474]

ELY, J., AND J. VÄLIMÄKI (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102, 84–105. [453]

ELY, J., J. HÖRNER, AND W. OLSZEWSKI (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73, 377–415. [453]

FORGES, F. (1992): "Non-Zero-Sum Repeated Games of Incomplete Information," in *Handbook of Game Theory*, Vol. 1, ed. by R. J. Aumann and S. Hart. Amsterdam: North Holland. [455]

FORGES, F., AND E. MINELLI (1997): "A Property of Nash Equilibria in Repeated Games With Incomplete Information," *Games and Economic Behavior*, 18, 159–175. [455]

FUDENBERG, D., AND D. LEVINE (1989): "Reputation and Equilibrium Selection in Games With a Single Patient Player," *Econometrica*, 57, 759–778. [473]

FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem in Repeated Games With Discounting or With Incomplete Information," *Econometrica*, 54, 533–554. [471]

——— (1991): "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, 53, 428–438. [484,485]

HANSEN, L. P., AND T. J. SARGENT (2007): *Robustness*. Princeton, NJ: Princeton University Press. [454]

HARSANYI, J. C. (1967–1968): "Games With Incomplete Information Played by Bayesian Players," *Management Science*, 14, 159–182, 320–334, 486–502. [453,456]

HART, S. (1985): "Nonzero-Sum Two-Person Repeated Games With Incomplete Information," *Mathematics of Operations Research*, 10, 117–153. [455]

HÖRNER, J., AND S. LOVO (2009): "Supplement to 'Belief-Free Equilibria in Games With Incomplete Information'," *Econometrica Supplemental Material*, 77, http://www.econometricsociety.org/ecta/Supmat/7134_extensions.pdf. [472]

HÖRNER, J., AND W. OLSZEWSKI (2006): "The Folk Theorem With Almost-Perfect Monitoring," *Econometrica*, 74, 1499–1544. [466]

ISRAELI, E. (1999): "Sowing Doubt Optimally in Two-Person Repeated Games," *Games and Economic Behavior*, 28, 203–216. [472,474]

KALAI, E. (2004): "Large Robust Games," *Econometrica*, 72, 1631–1665. [453]

KOREN, G. (1992): "Two-Person Repeated Games Where Players Know Their Own Payoffs," Working Paper, New York University. [455,462,469]

LUCE, D., AND H. RAIFFA (1957): *Games and Decisions*. New York: Wiley. [454]

MEGIDDO, N. (1980): "On Repeated Games With Incomplete Information Played by Non-Bayesian Players," *International Journal of Game Theory*, 9, 157–167. [455]

MILNOR, J. (1954): "Games Against Nature," in *Decision Processes*, ed. by R. M. Thrall, C. H. Coombs, and R. L. Davis. New York: Wiley, 49–59. [455]

MONDERER, D., AND M. TENNENHOLTZ (1999): "Dynamic Non-Bayesian Decision Making in Multi-Agent Systems," *Annals of Mathematics and Artificial Intelligence*, 25, 91–106. [454-456]

PICCIONE, M. (2002): "The Repeated Prisoner's Dilemma With Imperfect Private Monitoring," *Journal of Economic Theory*, 102, 70–83. [453]

SCHMIDT, K. (1993): "Reputation and Equilibrium Characterization in Repeated Games With Conflicting Interests," *Econometrica*, 61, 325–351. [473,474]

SHALEV, J. (1994): "Nonzero-Sum Two-Person Repeated Games With Incomplete Information and Known-Own Payoffs," *Games and Economic Behavior*, 7, 246–259. [455,469,472]

SORIN, S. (1986): "On Repeated Games With Complete Information," *Mathematics of Operations Research*, 11, 147–160. [484,485]

VON NEUMANN, J., AND O. MORGENSTERN (1944): *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press. [454]

Dept. of Economics, Yale University, Box 208281, New Haven, CT 06520-8281, U.S.A. and Centre for Economic Policy Research, London EC1V 0DG, U.K.; johannes.horner@yale.edu

and

Dept. of Finance and Economics, HEC School of Management, Paris and GREGHEC, 1 rue de la Liberation, 78351 Jouy-en-Josas, France; lovo@hec.fr.

Manuscript received April, 2007; final revision received August, 2008.