Bayesian games with contracts*

Françoise Forges†

September 2011 (with some corrections in November 2011)

Abstract

The set of all Bayesian-Nash equilibrium payoffs that the players can achieve by signing (conditional) commitments before playing a Bayesian game coincides with the set of all feasible, incentive compatible and interim individually rational payoffs of the Bayesian game. Furthermore, the various equilibrium payoffs, which are achieved by means of different commitment devices, are also the equilibrium payoffs of a universal, deterministic commitment game.

Keywords: Bayesian game, commitment, contract, incentive compatibility, interim individual rationality.

JEL classification: C70

---

*I wish to thank the Economics Department of Brown University, where I prepared the first version of this note in March 2011, for its hospitality. I am also grateful to Ehud Kalai for comments.

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

1 Introduction

Let us assume that, before making simultaneous strategic decisions, privately informed players can sign binding agreements. Can we characterize the set of all payoffs that they can reach in this way? Myerson (1991, chapter 6) gives an answer: the payoffs are the ones that are achieved by means of a random mechanism satisfying "informational incentive constraints" and "general participation constraints". However, as recently pointed out by Kalai et al. (2010), even in the case of two players with complete information, natural commitments may be conditional, so that some care is needed to avoid circularities. Consider for instance the price competition between two sellers; if every seller posts a price and commits to undercut his competitor's price by some amount, the outcome of the commitment strategies might not be well-defined.

Kalai et al. (2010) propose a model which overcomes the difficulties. They consider an arbitrary set D_i of commitment devices for player i ($i = 1, 2$); they do not impose any conceptual restriction on the D_i's. In order to account for conditional commitments, they assume that player i chooses, at the same time as a device d_i, a *response function* R_{d_i}, which describes how he would reply to the other player's commitment device d_{-i}. The key to avoid circular or endless reasoning is that R_{d_i} takes its value in the set A_i of actions of player i in the original underlying two person game. The pair of response functions (R_{d_1}, R_{d_2}) then amounts to a single "grand response function" $R : D_1 \times D_2 \to A_1 \times A_2$ which determines, without any ambiguity, the pair of actions to be played as a result of the players' commitments. The sets D_i's and a response function R transform the original game G into a contract game $G(\mathcal{D})$, where $\mathcal{D} \equiv (D_1, D_2, R)$. It is understood that participation in $G(\mathcal{D})$ is voluntary, in the sense that, in $G(\mathcal{D})$, every player may decide not to commit and just choose an action in the original game.

Kalai et al. (2010)'s main result can be roughly stated as follows: let G be a two-person strategic form game; the set of all Nash equilibrium payoffs that can be achieved in some commitment game $G(\mathcal{D})$ extending G coincides with the set of *feasible* and *individually rational* payoffs of G. As a by-product, Kalai et al. (2010) construct a *universal* commitment game extending G in which all these Nash equilibrium payoffs can be achieved at once. Finally, the results go through in the case of n players if correlated strategies are allowed.

In this note, we extend Kalai et al. (2010)'s result to n person games with incomplete information, namely Bayesian games. A relevant question is then the stage at which the players sign binding agreements: ex ante or interim. We follow the well-founded tradition according to which players make commitments after having learnt their types, namely, at the interim stage (see, e.g., Myerson (1991)). This assumption allows us to describe the commitment *possibilities* of the players in a *Bayesian* game G in the same way as in the case of complete information, i.e., *exactly* as in Kalai et al. (2010). In other words, the *set* of commitment devices (or, equivalently, the *set* of instructions to a mediator) that are available to a player is described independently of his private information. A natural justification for this model is that the authority implementing the contracts (namely, the mediator) does not know to which extent the players have private information. Of course, every player chooses his effective commitment device as a function of his type. Being modelled in the same way as in Kalai et al. (2010), our commitments can be conditional but do not give rise to any circularity.

We show that the set $\mathcal{NC}(G) \equiv \cup_{\mathcal{D}} N\left[G(\mathcal{D})\right]$ of all Bayesian-Nash equilibrium payoffs that can be achieved in some contract game $G(\mathcal{D})$ extending a given n person Bayesian game G coincides with the set $\mathcal{F}(G)$ of *feasible*, *incentive compatible* and *interim individually rational* payoffs of G, namely that $\mathcal{NC}(G) = \mathcal{F}(G)$. As mentionned above, this result is suggested in Myerson (1991, section 6.6), who gives precise definitions of the previous three properties[1]. However, Myerson (1991) does not allow for conditional commitments. Furthermore, he sticks to random mechanisms (while, as seen above, Kalai et al. (2010)'s response functions are deterministic) and does not propose any universal contract game. We propose a full extension of Kalai et al. (2010)'s result to games with incomplete information by constructing a universal deterministic commitment game $G(\mathcal{D}^*)$ extending G such that $N\left[G(\mathcal{D}^*)\right] = \mathcal{NC}(G) = \mathcal{F}(G)$.

Kalai et al. (2010)'s result is related to the well-known "folk theorem" which states that, if G is a game with complete information, the set of all equilibrium payoffs of the standard infinitely repeated game associated with G is the set of feasible and individually rational payoffs of G. Hence, under complete information, the infinite repetition of the game has the same effect as binding agreements in the one-shot game. This property no longer holds under incomplete information (see, e.g., Forges (1992)). More precisely, as

[1] In particular, interim individual rationality must be formulated in terms of vector payoffs (see also Hart (1985)).

shown in this note, contracts at the interim stage of a Bayesian game G enable the players to reach a large set of feasible and individually rational payoffs of the one-shot game G but even if G is a two person game, with a single informed player, the equilibrium payoffs of the infinitely repeated version of G cannot be described as the equilibrium payoffs of a one-shot game (see Hart (1985)).

2 Contracts in a Bayesian game

Let us fix a Bayesian game $G \equiv \left[N, \{T_i, A_i, u_i\}_{i\in N}, q\right]$, namely,

- a set of players N
- for every player i, $i \in N$, a (finite) set of types T_i, a (finite) set of actions A_i and a utility function $u_i : T \times A \to \mathbb{R}$, where $T = \prod_{i\in N} T_i$ and $A = \prod_{i\in N} A_i$
- a probability distribution q over T.

A *contract space* for G is defined by a set of instructions D_i for every player $i \in N$ and a response function $R : D \to \Delta(A)$, where $D = \prod_{i\in N} D_i$. We denote such a contract space as $\mathcal{D} = (D_i, i \in N; R)$. The interpretation is that every player gives his instructions to a mediator who is entitled to play G on his behalf. We do not impose any restriction (beyond standard measurability assumptions, see below) on the sets of instructions D_i. These *sets* are neutral, i.e., independent of the private information that the players may have, but of course, the players will typically give their instructions as a function of their type.

The formulation of the response function implicitly implies that, given the instructions of all the players, the mediator must be able to choose an $n-$tuple of actions in G, possibly by performing a lottery. This general definition is similar to the ones previously adopted in the literature (see, e.g., Myerson (1991, chapter 6), Ashlagi et al. (2009)). In particular, the formulation precludes the mediator to face inconsistent or circular instructions that would make him unable to select actions in A in a well-defined way[2].

[2]For instance,the instructions of the two sellers undercutting each other's prices (see the introduction) do not generate a response function.

Kalai et al. (2010), who focus on games with complete information, further require that the response function R be *deterministic*, namely takes the form of a function $R : D \to A$. This formulation strengthens the consistency requirement on R. Indeed, R can then be interpreted as the result of *individual* response functions. More precisely, if R is deterministic, we can write

$$\begin{aligned} R(d_1, ..., d_n) &= (R_1(d_1, ..., d_n), ..., R_n(d_1, ..., d_n)) \\ &= (R_{d_1}(d_2, ..., d_n), ..., R_{d_n}(d_1, ..., d_{n-1})) \end{aligned}$$

with $R_{d_i} : \prod_{j \neq i} D_j \to A_i$. Kalai et al. (2010) view D_i as a set of (conditional) commitment devices for player i and R_{d_i} as the individual response function which is associated with player i's device d_i. Player i's commitment to d_i amounts to committing to R_{d_i}, hence d_i can be interpreted as a conditional commitment. Until proposition 2, we stick to possibly random response functions, with values in $\Delta(A)$. In proposition 2, we show that such random response functions do not enable the players to achieve more outcomes than deterministic ones and we thus recover Kalai et al. (2010)'s interpretation[3].

As in previous models of games with contracts, every player will be allowed to reject any commitment and to act by himself in G. In order to capture this property, we define a contract space $\mathcal{D} = (D_i, i \in N; R)$ to be *voluntary* if for every player i, $D_i \supseteq A_i$ and for every $a_i \in A_i$ and $d_{-i} \in D_{-i}$, $R(a_i, d_{-i}) = \delta_{a_i} \otimes R_{-i}(d_{-i})$, where $\delta_{a_i} \in \Delta(A_i)$ chooses a_i with probability 1 and $R_{-i}(d_{-i}) \in \Delta(A_{-i})$. In other words, player i can just give the mediator the instruction to play (unconditionally) a_i on his behalf; in this case, the mediator plays a_i for player i, independently of the instructions of the other players, whatever these are[4]. Henceforth, we focus on voluntary contract spaces.

Let $\mathcal{D} = (D_i, i \in N; R)$ be a contract space. The (Bayesian) game $G(\mathcal{D})$ in which the players can sign a contract in $\mathcal{D}$ before playing G is described

[3]Except for the - temporary - possibility of randomization in R, our "contract spaces" are formally identical to Kalai et al. (2010)'s "spaces of commitment devices" or "device spaces". The terminology may nevertheless reflect more or less intended decentralization (see section 6.5 in Kalai et al. (2010) and our concluding remarks). We follow Myerson (1991)'s terminology.

[4]As Kalai et al. (2010), we conveniently model no participation in any contract as a very special form of commitment. Our definition of voluntary contract space is equivalent to Kalai et al.'s one and captures their two conditions of "unconditioned" and "private" play at the same time.

as follows:

- Types are chosen in T according to q and every player i privately observes his own type t_i

- Every player chooses instructions d_i in D_i and secretely transmits them to the mediator

- Given $d = (d_1, ..., d_n)$, the mediator chooses an $n-$tuple of actions $a \in A$ according to $R(d)$ (and makes the decision a in G on behalf of the players).

$G(\mathcal{D})$ will be referred to as a *contract game* extending G. $G(\mathcal{D})$ will be said *voluntary* if the underlying contract space $\mathcal{D}$ is voluntary. $G(\mathcal{D})$ will be said *deterministic* if the response function R in the underlying contract space $\mathcal{D}$ is deterministic, namely, $R : D \to A$.

Assuming an appropriate measurable structure on D_i for every $i \in N$, a *strategy* of player i in $G(\mathcal{D})$ is a mapping $\sigma_i : T_i \to \Delta(D_i)$, which thus allows player i to randomize. Bayesian-Nash equilibria in $G(\mathcal{D})$ are defined in the usual way. Let $N\left[G(\mathcal{D})\right]$ be the set of all Bayesian-Nash equilibrium payoffs of $G(\mathcal{D})$.

Our goal is to characterize the set $\mathcal{NC}(G)$ of all payoffs that can be achieved at a Bayesian-Nash equilibrium of some voluntary contract game extending G, namely $\mathcal{NC}(G) = \cup_{\mathcal{D}\, voluntary} N\left[G(\mathcal{D})\right]$. The characterization of $\mathcal{NC}(G)$ will be formulated in terms of feasible, incentive compatible and interim individually rational payoffs in G. We thus start by recalling the definition of these properties. We closely follow Myerson (1991, sections 6.4 and 6.6).

A *correlated strategy* (a mechanism, in Myerson (1991)'s terminology) for the set of all players N in G is a mapping $\mu : T \to \Delta(A)$, namely, $\mu = (\mu(.|t))_{t \in T}$, with $\mu(.|t) \in \Delta(A)$ for every $t \in T$. Let $\mathcal{C} = \Delta(A)^T$ be the set of all correlated strategies of N.

An $n-$tuple of vector payoffs (or simply, a payoff) $x = \left[(x_i(t_i))_{t_i \in T_i}\right]_{i \in N}$, is *feasible* in G if there exists a correlated strategy μ achieving x, namely

$$x_i(t_i) = \sum_{t_{-i}} q(t_{-i}|t_i) \sum_a \mu(a|t) u_i(t, a) \quad i \in N,\ t_i \in T_i \tag{1}$$

We then write $x_i(t_i) = U_i(\mu|t_i)$ and $x = U(\mu)$.

A feasible payoff $x = U(\mu)$ is *incentive compatible* (I.C.) if

$$U_i(\mu|t_i) \geq \sum_{t_{-i}} q(t_{-i}|t_i) \sum_a \mu(a|s_i, t_{-i}) u_i(t, a) \quad \text{for every } i \in N, t_i, s_i \in T_i \tag{2}$$

A payoff x is *interim individually rational* (INTIR) if, for every player i, there exists a correlated strategy τ_{-i} of players in $N \setminus \{i\}$ such that

$$x_i(t_i) \geq \max_{a_i \in A_i} \sum_{t_{-i}} q(t_{-i}|t_i) \sum_{a_{-i}} \tau_{-i}(a_{-i}|t_{-i}) u_i(t, a) \quad \text{for every } t_i \in T_i \tag{3}$$

We denote as $\mathcal{P}_{-i} = \Delta(A_{-i})^{T_{-i}}$ the set of correlated strategies of players in $N \setminus \{i\}$. Let $\mathcal{P} = \prod_{i \in N} \mathcal{P}_{-i}$. If τ_{-i} satisfies the previous inequalities, we refer to it as a *correlated punishment strategy* against player i. Observe that τ_{-i} treats all types of player i in the same way[5].

We denote as $\mathcal{F}(G)$ the set of feasible, incentive compatible and interim individually rational payoffs in G.

3 Characterization of feasible contracts

We start with the easier direction: the set $\mathcal{NC}(G)$ of all payoffs that can be achieved at a Bayesian-Nash equilibrium of some voluntary contract game extending G is included in the set of feasible, incentive compatible and interim individually rational payoffs in G.

Proposition 1 *Let $G(\mathcal{D})$ be a voluntary contract game extending G, and let x be a Bayesian-Nash equilibrium of $G(\mathcal{D})$. Then x is feasible, incentive compatible and interim individually rational, namely, $\mathcal{NC}(G) \subseteq \mathcal{F}(G)$.*

Proof: Let $\mathcal{D} = (D_i, i \in N; R)$ be the contract space and let $\sigma = (\sigma_i)_{i \in N}$ be the Bayesian-Nash equilibrium in $G(\mathcal{D})$ associated with x. We proceed in the same way as in standard proofs of the revelation principle. For every $t \in T$, let $\mu(.|t) \in \Delta(A)$ be the probability distribution over A induced by σ and R. Let $\mu = (\mu(.|t))_{t \in T} \in \mathcal{C}$ be the corresponding correlated strategy: $x = U(\mu)$. The payoff x is I.C. because σ is an equilibrium and thus player

[5] Under complete information, in the case of two players, the corresponding individual rationality level of player i is the standard minmax level, namely $\min_{\tau_{-i} \in \Delta(A_{-i})} \max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \tau_{-i}(a_{-i}) u_i(a_i, a_{-i})$.

i of type t_i cannot profit from using $\sigma_i(s_i)$, for $s_i \neq t_i$, as long as the other players follow σ.

To show that x is INTIR, let, for every player i, $\tau^*_{-i} \in \mathcal{P}_{-i} = \Delta(A_{-i})^{T_{-i}}$ be the conditional distribution induced by R and σ_{-i} when player i "does not participate", i.e., gives the instruction to play some given action a_i. Since $\mathcal{D}$ is voluntary, τ^*_{-i} is well-defined and independent of a_i (and, of course, of player i's type). If x is not INTIR, there exist a player i, a type t_i and an action a^*_i such that (3) is not satisfied. Hence player i of type t_i profitably deviates by giving the instruction a^*_i and facing τ^*_{-i}. ■

Example: Let us illustrate the previous result on Kalai et al. (2010)'s example 5, which shows that in Bayesian games, commitments may be used as means of communication. Kalai et al. (2010)'s commitment devices in that example form an interesting "indirect mechanism". The proof of the previous revelation principle associates a correlated strategy, namely a "direct mechanism" with the indirect one.

Kalai et al. (2010) describe the basic Bayesian game G as follows: "a treasure is buried in one of three locations L_1, L_2 and L_3 with equal probability. It takes two to dig for it. Player 1 lives in L_1 and if the treasure is buried there he knows it. Likewise, player 2, who lives in L_2, knows if the treasure is buried there. The players cannot communicate, but they can move simultaneously to any location. If they happen to meet at the location of the treasure they dig for it". They show that, in any equilibrium of G, the probability that a player reaches the location of the treasure cannot be higher than 2/3.

They propose a commitment space in which each player i has two devices (or possible instructions), d^s_i (for stubborn) and d^f_i (for flexible). The device d^s_i chooses the location L_i, independently of the device of the other player. The device d^f_i chooses the location L_{-i} against the device d^s_{-i} of the other player, but chooses L_3 against the device d^f_{-i} of the other player. The response function R is thus

$$\begin{array}{ccc} & d^s_2 & d^f_2 \\ d^s_1 & (L_1, L_2) & (L_1, L_1) \\ d^f_1 & (L_2, L_2) & (L_3, L_3) \end{array}$$

In the contract game in which players choose devices, Kalai et al. (2010) consider the following strategies: every player i chooses d^s_i when the treasure is at his location and d_i^f otherwise. This strategy profile is an equilibrium, which induces the following (pure) correlated strategy μ:

$$\begin{array}{ccc} & L_2 & \textit{not } L_2 \\ L_1 & (L_1, L_2) & (L_1, L_1) \\ \textit{not } L_1 & (L_2, L_2) & (L_3, L_3) \end{array}$$

This example illustrates well how a deterministic response function R can be generated by decentralized, individual response functions of every player to the other player's device. Decentralization somehow disappears in the correlated strategy μ resulting from the revelation principle. This phenomenon is of course well-known in implementation theory and more generally in mechanism design. In any case, proposition 1 characterizes the largest set of payoffs $\mathcal{F}(G)$ that can be achieved if players can voluntarily commit themselves.

We now investigate the converse of proposition 1. The following first step is readily in Myerson (1991, sections 6.1 and 6.6).

Lemma 1 *Let x be a feasible, incentive compatible and interim individually rational payoff in G. There exists a voluntary contract game $G(\mathcal{D})$ extending G such that x is a Bayesian-Nash equilibrium payoff of $G(\mathcal{D})$, namely, $\mathcal{F}(G) \subseteq \mathcal{NC}(G)$.*

Proof: Let $\mu \in \mathcal{C}$ be a correlated strategy such that $x = U(\mu)$. Since x is INTIR, there exists $(\tau_{-i})_{i\in N} \in \mathcal{P} = \prod_{i\in N} \mathcal{P}_{-i}$ such that (3) is satisfied for every $i \in N$. With the understanding that $T_i \cap A_i = \emptyset$, let $D_i = T_i \cup A_i$, $i \in N$, and

$$\begin{aligned} R(d_1, ..., d_n) &= \mu(.|t_1, ..., t_n) \text{ if } d_i = t_i \in T_i \text{ for every } i \in N \\ &= \delta_{a_i} \otimes \tau_{-i}(.|t_{-i}) \text{ if } d_i = a_i \in A_i \text{ for some } i \in N \text{ and } d_j = t_j \in T_j \text{ for } j \neq i \\ &= \text{some arbitrary } a^* \in A \text{ otherwise} \end{aligned}$$

Let $\mathcal{D} = (D_i, i \in N; R)$ be the associated contract space. $G(\mathcal{D})$ is a voluntary contract game, with our convention that every instruction a_i in A_i is interpreted as "non-participation" of player i. The strategy profile σ defined by $\sigma_i(t_i) = t_i$ for every $i \in N$, $t_i \in T_i$ (in which every player "participates", whatever his type) is an equilibrium of $G(\mathcal{D})$, with payoff x. ■

In lemma 1, a different contract space, namely a different response function R, is constructed for every payoff x in $\mathcal{F}(G)$, which amounts to assuming that the mediator implementing the contract for the players knows x. Our next step is to construct a single contract game, which can be used for all $x \in \mathcal{F}(G)$.

Lemma 2 *There exists a voluntary contract game $G(\mathcal{D}^*)$ extending G such that every feasible, incentive compatible and interim individually rational payoff in G is a Bayesian-Nash equilibrium payoff of $G(\mathcal{D}^*)$, namely, $\mathcal{F}(G) \subseteq N[G(\mathcal{D}^*)]$.*

Proof: Let us set $C_i = \mathcal{C} \times T_i \times \prod_{j \neq i} \mathcal{P}_{-j}$ for every $i \in N$ and $D_i^* = C_i \cup A_i$ (with the understanding that $C_i \cap A_i = \emptyset$). The interpretation is that every player i decides either "to participate" by giving an instruction in C_i, which contains a correlated strategy and possible punishments, or "not to participate" by choosing an (unconditional) action a_i. More precisely, we describe the response function R^* by distinguishing several cases:

(1) $d_i \in C_i$ for every $i \in N$; let us set $d_i = (\mu_i, t_i, (\tau^i_{-k})_{k \neq i})$.

- **(1.a)** If $\mu_1 = ... = \mu_n = \mu$, then $R^*(d_1, ..., d_n) = \mu(.|t_1, ..., t_n)$.
- **(1.b)** Otherwise, if $n = 2$, $R^*(d_1, d_2) = \tau_1(.|t_1) \otimes \tau_2(.|t_2)$; if $n \geq 3$ and for some j and every $i \neq j$, $\mu_i = \mu$ and $\tau^i_{-j} = \tau_{-j}$, then $R^*(d_1, ..., d_n) = \delta_{a_j} \otimes \tau_{-j}(.|t_{-j})$, where a_j is chosen arbitrarily in A_j.

(2) For some j, $d_j = a_j \in A_j$ and for every $i \neq j$, $d_i = (\mu_i, t_i, (\tau^i_{-k})_{k \neq i}) \in C_i$, with $\tau^i_{-j} = \tau_{-j}$. Then $R^*(d_1, ..., d_n) = \delta_{a_j} \otimes \tau_{-j}(.|t_{-j})$, as in (1.b), but a_j is player j's choice.

(3) In all other cases, $R^*(d_1, ..., d_n) =$ some arbitrary $a^* \in A$.

Let $\mathcal{D}^* = (D_i^*, i \in N; R^*)$; $G(\mathcal{D}^*)$ is a voluntary contract game. Let $x \in \mathcal{F}(G)$; as in the proof of lemma 1, let $\mu \in \mathcal{C}$ be such that $x = U(\mu)$ and let $(\tau_{-i})_{i \in N} \in \mathcal{P} = \prod_{i \in N} \mathcal{P}_{-i}$ be such that (3) is satisfied. Consider the following strategy profile σ in $G(\mathcal{D}^*)$: $\sigma_i(t_i) = (\mu, t_i, (\tau_{-k})_{k \neq i})$ for every $i \in N$, $t_i \in T_i$. σ forms an equilibrium in $G(\mathcal{D}^*)$ and its payoff is x. ■

In the previous proof, the contract game $G(\mathcal{D}^*)$ does not depend on the payoff $x \in \mathcal{F}(G)$ but still formally depends on the game G, through the sets of types and actions. To make the contract game fully universal, i.e., independent of the underlying Bayesian game, it suffices to ask the players to report the parameters of G as part of their strategies in the contract game, and to implement the contract only if they agree on these parameters[6]. Kalai et al. (2010)'s device space for games with complete information is "universal" in the sense of lemma 2, namely, independent of the payoff $x \in \mathcal{F}(G)$.

As emphasized earlier, Kalai et al. (2010)'s device space has a further desirable property: it involves a *deterministic* response function, which is easily generated by individual response functions. If we think of the contract as being implemented by a mediator, deterministic response functions are useful in making it possible to check that the mediator indeed followed the players' instructions.

The next proposition will be established by modifying the proof of lemma 2. As in Kalai et al. (2010), the *jointly controlled lotteries* introduced by Aumann et al. (1968) (see also Aumann and Maschler (1985)) are a basic tool for the construction. In the case of incomplete information, the players perform jointly controlled lotteries in parallel, for every possible $n-$tuple of types t.[7]

Proposition 2 *There exists a voluntary,* ***deterministic*** *contract game $G(\mathcal{D}^*)$ extending G such that every feasible, incentive compatible and interim individually rational payoff in G is a Bayesian-Nash equilibrium payoff of $G(\mathcal{D}^*)$, namely, $\mathcal{F}(G) \subseteq N\left[G(\mathcal{D}^*)\right]$.*

Proof: Going back to the proof of lemma 2, recall that $C_i = \mathcal{C} \times T_i \times \prod_{j \neq i} \mathcal{P}_{-j}$. Let us first replace $\mathcal{C} = \Delta(A)^T$ with $\mathcal{C}' = \mathcal{C} \times [0,1]^T$. Let us assume that a partition of $[0,1]$ into subintervals $I_\alpha(a)$ of length $\alpha(a)$ is associated to every $\alpha \in \Delta(A)$. When player i gives the instruction $\mu_i \in \mathcal{C}$, he also selects a number $r_i(t) \in [0,1]$, for every $t \in T$. Let us assume that we are in case (1.a), namely that $\mu_1 = ... = \mu_n = \mu$ and that the reported types are $t = (t_1, ..., t_n)$. Instead of performing the lottery $\mu(.|t)$, the mediator chooses $a \in A$ if $\sum_{i \in N} r_i(t) \bmod 1 \in I_{\mu(.|t)}(a)$. At equilibrium, every player i chooses his numbers $r_i(t)$ uniformly in $[0,1]$, independently for every t.

[6] Note the formalism of the proof of lemma 2 does not require that the mediator interprets the sets T_i in any particular way. The mediator gets an n-tuple of inputs in $\prod_{i \in N} D_i^*$ from the players, which enables him to select an n-tuple of actions in A to be played on behalf of the players.

[7] In the next proposition, the sets of commitment devices D_i are infinite. Kalai et al. (2010) show that in general, this assumption is necessary to get all payoffs in $\mathcal{F}(G)$ as equilibrium payoffs of a universal contract game. They also give an approximation result.

This guarantees that a is chosen with probability $\mu(a|t)$, even in case of a unilateral deviation in the choice of the numbers.

One can proceed similarly for the punishment strategies, by replacing $\mathcal{P}_{-j}$ by $\mathcal{P}'_{-j} = \mathcal{P}_{-j} \times [0,1]^{T_{-j}}$, for every $j \in N$. ■

Putting propositions 1 and 2 together, we get

Corollary *There exists a voluntary, deterministic contract game $G(\mathcal{D}^*)$ extending G such that $N[G(\mathcal{D}^*)] = \mathcal{NC}(G) = \mathcal{F}(G)$.*

4 Concluding remarks

While deterministic voluntary contract games are designed to capture decentralized individual commitment devices, the equilibria constructed in the proof of Kalai et al. (2010)'s theorem 1 and in our proposition 2 require a great deal of coordination between the players. Selection among the many equilibria exacerbates the problem (see Kalai et al. (2010), section 6.5). At first sight, the role of the coordinating authority might appear even more important in the case of incomplete information. However, as we pointed out earlier, in the proof of proposition 2, the mediator does not need to know what the sets T_i's are. What really seems to matter in the model of commitment is that some authority is entitled to make decisions on behalf of the players, as a function of their instructions. So the issue of decentralized commitment seems to be independent of the underlying information structure.

Kalai et al. (2010) point out that "when dealing with commitments in Bayesian games, there are several modeling alternatives. For example, the individual commitments can be done before or after the private information is revealed". In this note, we focus on *interim* commitments, with the usual (w.o.l.g.) understanding that interim participation in a contract means that all types of a player should participate. Myerson (e.g., 1991) provides substantial arguments in favor of this point of view. As shown in this note, interim commitments are tractable because the players' sets of instructions can be described in a neutral way, as in the case of complete information, while the players give their effective instructions as a function of their information. In particular, a single stage is considered at which every player gives all his instructions at once. Ex ante commitment typically requires two stages: a first one to commit on a contingent plan (i.e., instructions as a function of the type) and a second one to report the information necessary to implement the contingent plan (see, e.g., Forges et al. (2002)).

In this note, we implicitly focus on commitment by the set of *all players*, submitted to *individual* participation constraints[8]. As soon as there are more than two players, information being complete or not, *group participation* constraints may matter as well. These are definitely easier to capture at the ex ante stage (see, e.g., Forges et al. (2002) and Biran and Forges (2011)).

References

[1] Ashlagi I., D. Monderer and M. Tennenholtz (2009), "Mediators in position auctions", *Games and Economic Behavior* 67, 2-21.

[2] Aumann, R. J. and M. Maschler (1995), *Repeated Games of Incomplete Information*, Cambridge: M.I.T. Press.

[3] Aumann, R. J., M. Maschler and R. Stearns (1968), "Repeated games with incomplete information: an approach to the nonzero sum case", Reports to the U.S. Arms Control and Disarmament Agency, ST-143, Chapter IV, 117-216.

[4] Biran, O. and F. Forges (2011), "Core-stable rings in auctions with independent private values", *Games and Economic Behavior* 73, 52-64.

[5] Forges, F. (1992), "Repeated games of incomplete information : non-zero-sum", in : R. Aumann et S. Hart , Handbook of Game Theory with Economic Applications, Elsevier Science Publishers, North-Holland, chapter 6, 155-177.

[6] Forges, F., J.-F. Mertens and R. Vohra (2002), "The ex ante incentive compatible core in the absence of wealth effects", *Econometrica* 70, 1865-1892.

[7] Hart S. (1985), "Nonzero-sum two-person repeated games with incomplete information", *Mathematics of Operations Research* 10, 117–153.

[8] We adopt the same framework as Kalai et al. (2010) in that respect. They establish their result in details for two-person games with complete information. They explain in section 6.2 that by considering correlated punishments, their technique can be used in the case of n players.

[8] Kalai, A., E. Kalai, E. Lehrer and D. Samet (2010), "A commitment Folk theorem", *Games and Economic Behavior* 69, 127-137.

[9] Myerson, R. (1991), *Game theory: analysis of conflict*, Harvard University Press.