International Journal of Game Theory (1995) 24: 223–237

Game Theory

Cheap Play with No Regret

BHASKAR CHAKRAVORTI[1]
Economics Research Group, Information Sciences & Technologies Research Laboratory Bellcore, Morristown, NJ 07962-1910, USA

JOHN P. CONLEY[1]

Department of Economics, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA

Abstract: This paper studies a special class of differential information games with pre-play communication – games with "cheap play". We consider problems in which there are several rounds of payoff-irrelevant publicly observable choice (or discussion) of actions, followed by a final round in which actions are binding and payoff relevant. A natural focal subset of equilibria of such games in one that consists of equilibria involving *no regret.* Such games were first studied by Green and Laffont (1987), where a criterion called *posterior implementability* is introduced with the intention of identifying regret-free equilibria in games with cheap play. This is simply a restriction on the Bayesian equilibrium of the underlying one-shot game. If indeed such a restriction does characterize regret-freeness, then the analytics of such situations would be enormously simplified since one can ignore the messy extended-form of the cheap play game; merely examining the one-shot game is sufficient. We argue that regret-freeness of an equilibrium has a subtle distinction: regret-freeness in moves and regret-freeness in assessments. We show that the former causes the extended-form to be irrelevant; posterior implementability completely characterizes equilibria with regret-freeness in moves. The latter, on the other hand, does not yield a similar principle: the extended-form cannot be ignored.

1 Introduction

The focus of this short paper is on an important sub-class of games with pre-play communication – games with *cheap play.* Such games involve several rounds of payoff-irrelevant publicly observable choice (or discussion) of actions, followed by a final round in which actions are binding and payoff relevant. If there is differential information among the players, such "cheap play" in the earlier rounds conveys information.

In a standard model of pre-play communication, the equilibrium set is even larger than that expected in the absence of communication (for example, through "babbling"). As first observed by Green and Laffont (1987), the crucial aspect of cheap play that sets it apart from other forms of pre-play communication is that a natural focus is on equilibria exhibiting a stability property which we shall call *regret-freeness.* In such equilibria, players do not regret their choices made in the cheap play phase even after private information is updated. Thus, the notion of an "equilibrium" is extended to a dynamic, forward-looking context. The focal nature of such equilibria is evident in a variety of settings involving pre-contractual negotiations in "good faith" and is

[1] The authors would like thank the participants of the economic theory workshop at the University of Illinois, particularly Mark Feldman, for comments and stimulating discussions. Thc usual dsiclaimer applies.

0020–7276/95/3/223–237 $ 2.50

motivated in Green and Laffont (1987). Our objective is not to debate the merits of this refinement (including its existence properties), but to provide a characterization of regret-free sequential equilibria in games with cheap play. It is presumed that players initially narrow the set of equilibria to ones that are regret-free and then apply standard refinements.

Green and Laffont (1987), introduce a simple condition called *posterior implementability* to identify regret-free equilibria. The condition is a restriction on the Bayesian equilibria of the underlying one-shot game (with no communication). It requires that any list of actions chosen by the players in a Bayesian equilibrium for some realization of their private information must constitute an equilibrium in the game obtained by adjusting the players' beliefs according to Bayes' Law, conditioned on the actual actions chosen by everybody.

Clearly, in the absence of an implicit extended game (referred to as the *extended-form* in the sequel) involving several stages of cheap play, such a concept makes no sense. After all, if actions were chosen only once – and were considered binding commitments – then it is meaningless to ponder on whether or not a player regrets her choice after observing the choices of others. Hence, the posterior implementability restriction must imply that such an extended-form does exist but is irrelevant: studying the one-shot game is sufficient to identify regret-free equilibria.

A priori, it would appear that the extended-form would be hard to dispense with. Consider the following simple example in Figure 1.

	L	*R*
U	1, 2	0, 2
D	1, 3	0, 3

Fig. 1

There are two players: the row player chooses from $\{U, D\}$ and the column player chooses from $\{L, R\}$. Trivially, the strategy combination: row chooses U and column chooses a mixed strategy with equal weight on L and R is a posterior implementable (Nash) equilibrium in the one-shot game[2]. Even after observing the actual actions of

[2] The example is clearly an odd one, but the objective is to make a strong case for not dismissing the extended-form a priori. Also, observe that the example shows that posterior implementable equilibria do not preclude ones that involve dominated strategies. As we have mentioned earlier, our objective is to characterize regret-free equilibria and not to evaluate the merits of this notion as a refinement concept.

the other, no player would regret his strategy. However, the "story" underlying this notion is, presumably, an extended game in which there are several rounds of non-binding play followed by a terminal binding round. Suppose there are two rounds of such cheap play followed by a terminal round of firm commitments. A "regret-free" sequential equilibrium of the extended game must display some form of stationarity along the equilibrium path. For example, each player may be expected to use the (mixed) strategy given above in the first round and subsequently choose whatever action is realized in the first round, provided no defection from this rule is observed; alternatively, each player may be expected to use the (mixed) strategy above in every round, provided no defection from this rule is observed. However, the rationality of such stationary choices cannot be independent of the posited behaviour off-the-equilibrium path. The Green-Laffont concept provides no guide to off-equilibrium path behaviour.

Suppose that player 2's strategy is to choose L in any zero-probability event. Suppose $[U, R]$ had been played in the first round of cheap play. If this action pair were played in the terminal round, the row player would get a zero payoff. Given the off-the-equilibrium path behaviour posited for the column player, the row player would always find it in her interest to defect to D in the second round of cheap play, thereby upsetting the posterior implementable equilibrium in an extended-form setting. Clearly, one cannot dismiss the extended-form, a priori.

Given that the definition of the set of equilibrium outcomes of a game with cheap play is potentially complicated, a result that shows that the extended-form is irrelevant has practical implications. For example, the design of optimal contracts in situations where such pre-play communication occurs becomes a straightforward programming problem and the messy extended game may be ignored.

Though Green and Laffont (1987) do not formally model the extended-form, they speculate that there may be some connection between perfect equilibria of the extended game induced by pre-play communication as la Farrell (1982) and posterior implementability. Our paper addresses this speculation. We develop a formal model of cheap play. We begin with the observation that a game-theoretic equilibrium, and a sequential equilibrium in particular, is a description of the information-contingent moves that a player makes and her assessments of the moves of other players; in an equilibrium, the assessments are consistent with the actual choiced made. We argue that the notion of a regret-free equilibrium involves a subtle distinction between no regret in the choice of moves and no regret in the choice of assessments.

We show that sequential equilibria of the extended game of cheap play with no regret in moves are completely characterized by the posterior implementability condition. Hence, under such a notion of regret-freeness, the "irrelevance of the extended-form" principle is true. On the other hand, no regret in assessments can never be captured by any restriction on the Bayesian equilibria of the one-shot underlying game; the extended-form cannot be ignored.

Section 2 contains the basic model. Section 3 presents equilibria with no regret in moves and Section 4 presents equilibria with no regret in assessments. The final section concludes.

2 The Model

A is a finite set of *outcomes*, $N = \{1, 2, ..., n\}$ is a set of *players*, and Θ is a finite set of *states of the world*. A state of the world $\theta \in \Theta$ is a profile $(\theta_i)_{i \in N}$. Each player i in N is characterized by a list $\langle \Theta_i, u_i, \tilde{\pi}_i, M_i \rangle$, which includes:

a set of possible *private observations* Θ_i,
a (von Neumann-Morgenstern) *utility function* $u_i : A \times \Theta \to \mathbb{R}$,
a *prior probability distribution on* Θ, $\tilde{\pi}_i : \Theta \to (0, 1]$, and
a finite set of *moves*, M_i.

In the sequel, for any set X_i, let $X \equiv \times_{i \in N} X_i$, $X_{-i} \equiv \times_{j \in N \setminus \{i\}} X_j$ and given $x_i \in X_i$, let $x \equiv (x_i)_{i \in N}$ and $x_{-i} \equiv (x_j)_{j \in N \setminus \{i\}}$. For any subset $C \subseteq N$, $x_C \equiv (x_i)_{i \in C}$ and $x_{-C} \equiv (x_j)_{j \in N \setminus C}$. Also, for any set Y, let $\Delta(Y)$ denote the set of randomizations over Y. Given a random variable $h : X \to \Delta(Y)$, we shall, with some abuse of notation, use $h(y \mid x)$ to denote the probability assigned to $y \in Y$ by the distribution $h(x)$.

An *outcome function* $g : M \to A$ specifies an outcome for every profile of moves. A *game (form)* Γ is a list $\langle A, N, M, g \rangle$. The *strategy space* for i in the game Γ is the set $S_i = \{s_i : \Theta_i \to \Delta(M_i)\}$. With slight abuse of notation, $s(\cdot \mid \theta)$ and $s_{-i}(\cdot \mid \theta_{-i})$ denote the joint probability distributions induced on M and M_{-i} by $s \in S$, and $s_{-i} \in S_{-i}$, respectively, given a realization $\theta \in \Theta$. Let the function $g * s : \Theta \to \Delta(A)$ be defined by:

$\forall a \in A, \forall \theta \in \Theta,$

$$g * s(a \mid \theta) = \begin{cases} \sum_{m' \in \{m : g(m) = a\}} s(m' \mid \theta) & \text{if } \exists m \text{ such that } g(m) = a \\ & \text{otherwise.} \end{cases}$$

This is the function that specifies the probability distribution induced on the set of outcomes given that s is played in the game in the state of the world θ.

Every player i updates her probability distribution on Θ upon observing an element of Θ_i using Bayes' Law. This is summarized by a *posterior probability distribution on* Θ_{-i}, $\pi_i : \Theta \to [0, 1]$, where $\pi_i(\theta)$ specifies the probability assigned by i to $\theta_{-i} \in \Theta_{-i}$, given the observation $\theta_i \in \Theta_i$.

The model thus far is assumed to be common knowledge in the sense of Aumann (1976).

A *Bayesian equilibrium* of Γ is a strategy profile $s \in S$ that specifies:
$\forall i \in N, \forall \theta_i \in \Theta_i, \forall s_i' \in S_i, s' \equiv (s_i', s_{-i}),$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \pi_i(\theta) g * s(a \mid \theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \pi_i(\theta) g * s'(a \mid \theta) u_i(a, \theta)$$

$\mathcal{E}_S(\Gamma) \subseteq S$ denotes the set of Bayesian equilibria of Γ, with $\mathcal{E}(\Gamma) = \{g * s : s \in \mathcal{E}_S(\Gamma)\}$ denoting the set of induced Bayesian equilibrium outcomes of Γ.

Next, we shall model the following situation. Prior to committing themselves to moves in the game, Γ, the players engage in pre-play communication. The communication is in the form of "cheap play" – each player makes or proposes a move for herself with the knowledge that it can be withdrawn at no cost. We allow for several rounds of such non-binding and payoff-irrelevant (simultaneous-move) play.

Let T denote the number of rounds of simultaneous-move choices available to the players such that the choices made from M in the first $(T-1)$ rounds are non-binding and payoff-irrelevant. The moves chosen from M in the T-th round are binding commitments and determine the final outcome. The introduction of these rounds yields a *cheap play extension of* Γ, and is denoted Γ^T.

Let M^t denote the t-fold product of M for all $t \in \{1, ..., T\}$. For a given sequence of play, $m^T \in M^T, m(t)$ is used to denote the profile of moves chosen in the t-th round. An information set for player i at any $t > 1$ in the extensive-form game induced by Γ^T (referred to in the sequel, simply, as Γ^T) is characterized by a pair $(\theta_i, m^{t-1}) \in \Theta_i \times M^{t-1}$ and an information set for i at $t = 1$ is, obviously, $\theta_i \in \Theta_i$. Let $\mathcal{H}$ denote the set of all information sets in Γ^T, with $\mathcal{H}_i$ identifying the subset of information sets that belong to player i. A strategy in Γ^T for player i is a function $\sigma_i : \mathcal{H}_i \to \Delta(M_i)$ and is given by a sequence $(s_i^t)_{t=1}^T$, where $s_i^1 \in S_i$, and for all $t > 1, s_i^1 : \Theta_i \times M^{t-1} \to \Delta(M_i)$. Let Σ_i denote the strategy space for i in Γ^T, A *system of beliefs* is an n-tuple of functions $\{\beta_i : \Theta \times \bigcup_{t=1}^T M_{-i}^t \to [0, 1]\}_{j \in N}$ such that $\sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, m_{-i}^t) = 1$ for every $i \in N, \theta_i \in \Theta_i$ and $m_{-i}^t \in \bigcup_{t'=1}^T M_{-i}^{t'}$. The function β_i specifies the probability distribution that player i assigns to Θ_{-i}, given her information θ_i and the previous moves of others. Let $\mathbb{B}$ denote the space of such systems of beliefs.

Given $\sigma = (s^t)_{t=1}^T$, let the function $g * \sigma : \Theta \to \Delta(A)$ be defined by:
$\forall a \in A, \forall \theta \in \Theta,$

$$g * \sigma(a \mid \theta) = \begin{cases} \sum_{m' \in \{m: g(m)=a\}} \tilde{\sigma}(m'(T) \mid \theta) & \text{if } \exists m \text{ such that } g(m) = a \\ & \text{otherwise,} \end{cases}$$

where $\tilde{\sigma}(m'(T) \mid \theta)$ denotes the probability that m' will be played in the T-th round under the strategy profile σ in state θ.

Also, for any $\bar{m}_i \in M_i$, for any $t' \in \{1, ..., T\}$, if $m_i(t) = \bar{m}_i$ for all $t \in \{1, ..., t'\}$, then we write m_i^t as $[\bar{m}_i]$; $[\bar{m}]$ and $[\bar{m}_{-i}]$ are defined in the obvious manner.

A *sequential equilibrium* of Γ^T is a pair $(\sigma, \beta) \in \Sigma \times \mathbb{B}$ that satisfies:
$\forall i \in N, \forall h_i = (\theta_i, m^{t-1}) \in \mathcal{H}_i, \forall \sigma_i' \in \Sigma_i, \sigma' \equiv (\sigma_i', \sigma_{-i}),$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^t) g * \sigma(a \mid \theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^t) g * \sigma'(a \mid \theta) u_i(a, \theta).$$

where β is compatible with the use of Bayes' Law whenever it is applicable.

$\mathcal{E}_{\Sigma \times \mathbb{B}}(\Gamma^T) \subseteq \Sigma \times \mathbb{B}$ denotes the set of sequential equilibria of Γ^T, with $\mathcal{E}(\Gamma^T) \equiv \{g * \sigma : \exists \beta \in \mathbb{B} \text{ such that } (\sigma, \beta) \in \mathcal{E}_{\Sigma \times \mathbb{B}}(\Gamma^T)\}$ denoting the set of induced sequential equilibrium outcomes of Γ^T.

In this paper, we will concentrate on the subset of sequential equilibria of Γ^T that are "regret-free". Such equilibria have the property that despite the changes in the beliefs of the players over the rounds of cheap play, the choices made in the early rounds remain optimal even when the time to make a binding commitment arrives.

Thus, some form of stationarity in strategies over time seems to capture the notion of no regret. There are primarily two approaches to modelling equilibria with no regret depending on the perspective one adopts regarding what an "equilibrium strategy" means. In equilibrium, a strategy for player i performs a dual role. It is a contigency plan for i specifying which move i should make in every state of i's information. On the other hand, it is a probability assessment by all $j \neq i$ about the moves that i is likely to make. Hence, to say that an equilibrium is regret-free could mean one of two things: (i) player i does not regret her moves made in the cheap play rounds, despite the information conveyed by cheap play; and (ii) each $j \neq i$ does not regret the assessment made about i's move, despite the information conveyed by cheap play.

In the subsequent sections, we shall address each form of no regret in turn.

3 Equilibria with No Regret in Moves

We begin with "no regret in moves" (NRM) and the associated notion of stationarity is denoted "NRM-stationarity".

A sequential equilibrium (σ, β) is *NRM-stationary* if $\sigma = (s^t)_{t=1}^T$ satisfies:
$\forall t \in \{2, ..., T\}, \forall i \in N, \forall (\theta_i, m^{t-1} \in \mathcal{H}_i,$
$[\exists \bar{m}_i \in M_i$ such that $m_i^{t-1} = [\bar{m}_i]$, and $s_i^1(\bar{m}_i \mid \theta_i) > 0] \Rightarrow [s_i^t(\bar{m}_i \mid \theta_i, m^{t-1}) = 1]$.

$\mathcal{E}^{NRM}_{\Sigma \times \mathbb{B}}(\Gamma^T) \subseteq \Sigma \times \mathbb{B}$ denotes the set of NRM-stationary sequential equilibria of Γ^T, with $\mathcal{E}^{NRM}(\Gamma) \equiv \{g * \sigma : \exists \beta \in \mathbb{B}$ such that $(\sigma, \beta) \in \mathcal{E}^{NRM}_{\Sigma \times \mathbb{B}}(\Gamma^T)\}$ denoting the set of NRM-stationary sequential equilibrium outcomes of Γ^T.

NRM-stationarity requires that the players stick to the moves chosen in the first round, provided that the observed moves in the first round have positive probability in the postulated equilibrium and the moves observed in all subsequent rounds are the same as the ones observed in the first round.

Next, we define a simple condition on the set of Bayesian equilibria of the underlying one-shot game Γ called *posterior implementability.* It was introduced by Green and Laffont (1987).

For all $i \in N$ and $s \in S$, define the function $\pi_i^s : \Theta \times M \to [0, 1]$ by
$\forall \theta \in \Theta, \forall m \in M,$

$$\pi_i^s(\theta, m) = \frac{s_{-i}(m_{-i} \mid \theta_{-i})\pi_i(\theta)}{\sum\limits_{\theta'_{-i} \in \Theta_{-i}} s_{-i}(m_{-i} \mid \theta'_{-i})\pi_i(\theta'_{-i}, \theta_i)}.$$

A Bayesian equilibrium s is *posterior implementable* if
$\forall \theta \in \Theta, \forall m \in M$ such that $s(m \mid \theta) > 0, \forall i \in N, \forall m'_i \in M_i,$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \pi_i^s(\theta, m) u_i(g(m), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \pi_i^s(\theta, m) u_i(g(m'_i, m_{-i}), \theta)$$

$\mathcal{E}_S^{PI}(\Gamma) \subseteq S$ denotes the set of posterior implementable Bayesian equilibria of Γ, with $\mathcal{E}^{PI}(\Gamma) = \{g * s : s \in \mathcal{E}_S^{PI}(\Gamma)\}$ denoting the set of posterior implementable Bayesian equilibrium outcomes of Γ.

The following theorem shows that NRM-stationarity is completely characterized by posterior implementability; hence, the extended-form of the game with cheap play rounds is irrelevant.

Theorem 1: $\mathcal{E}^{NRM}(\Gamma^T) = \mathcal{E}^{PI}(\Gamma)$ for all $T > 1$.

Proof: The theorem is a consequence of Lemmas 1 and 2. We begin by defining an additional piece of notation.

For any $i \in N$, let $C_i(m^{T-1}), s, \theta_i) \equiv \{j \in N$: either (i) $m_j^{T-1} \neq [m_j]$ for any $m_j \in M_j$ or (ii) given $(\theta_i, m^{T-1}) \in \mathcal{H}_i$, and $m_j \in M_j$, $m_j^{T-1} = [m_j]$ implies that for all $\theta_j \in \Theta_j$ such that $\sum_{\theta_{-\{ij\}} \in \Theta_{-\{ij\}}} \pi_i(\theta) > 0, s_j(m_j \mid \theta_j) = 0\}$.

Lemma 1: $\mathcal{E}^{PI}(\Gamma) \subseteq \mathcal{E}^{NRM}(\Gamma^T)$.

Proof of Lemma 1: Choose $s \in \mathcal{E}^{PI}(\Gamma)$. Also, choose $\sigma = (s^t)_{t=T}^T \in \Sigma$ such that

(a) $s^1 = s$,

(b) for all $i \in N$, for all $(\theta_i, m^{T-1}) \in \mathcal{H}_i$, if $m_i^{T-1} \neq [m_j]$ for all $m_i \in M_i$ such that $s_i(m_i \mid \theta_i) > 0$, then $s_i^T(\cdot \mid \theta_i, m^{T-1}) = s_i(\cdot \mid \theta_i)$.

(c) for all $i \in N$, for all $(\theta_i, m^{T-1}) \in \mathcal{H}_i$, if there exists $m_i \in M_i$ such that $m_i^{T-1} = [m_i]$ and $s_i(m_i \mid \theta_i) > 0$, then $s_i^T(m_i \mid \theta_i, m^{T-1}) = 1$.

Let $\hat{\beta}_i(\theta_j, m_{-i}^{T-1})$ and $\hat{\pi}_i(\theta_j, \theta_i)$ denote the marginal probabilities on Θ_j consistent with $\beta_i(\theta, m_{-i}^{T-1})$ and $\pi_i(\theta)$, respectively.

Define $\beta \in \mathbb{B}$ such that it satisfies:
for all $i \in N$, for all $(\theta_i, m^{T-1}) \in \mathcal{H}_i$, if $j \notin C_i(m^{T-1}, s, \theta_i)$, then, given $m_j^{T-1} = [m_j]$

$$\hat{\beta}_i(\theta_j, m_{-i}^{T-1}) = \frac{s_j^1(m_j \mid \theta_j)\hat{\pi}_i(\theta_j, \theta_i)}{\sum_{\theta_j \in \Theta_j} s_j^1(m_j \mid \theta_j')\hat{\pi}_i(\theta_j', \theta_i)}$$

if $j \in C_i(m^{T-1}, s, \theta_i)$,

$$\hat{\beta}_i(\theta_j, m^{t-1}) = \hat{\pi}_i(\theta_j, \theta_i).$$

Next, we need to show that the pair (σ, β) is indeed an equilibrium.

Choose $i \in N$ and $\theta_i \in \Theta_i$. By definition of $\mathcal{E}^{PI}(\Gamma)$, for all $m^{T-1} \in M^{T-1}$ and $m \in M$ such that $m^{T-1} = [m]$ and $c_i(m^{T-1}, s, \theta_i) = \emptyset$, by construction of β, the following inequality is met for all $\tilde{m}_i \in M_i$,

$$\sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m_{-i}])u_i(g(m), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m_{-i}])u_i(g(m_{-i}, \tilde{m}_i), \theta).$$

On the other hand, if $C_i(m^{t-1}, s, \theta_i) \neq \emptyset$, for some $m^{T-1} \in M^{T-1}$, since $s \in \mathcal{E}_S^{PI}(\Gamma)$, for all $s'_i \in S_i$, $s' \equiv (s'_i, s_{-i})$, and by construction of β, we have:

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^{T-1}) g * \bar{s}(a \mid \theta) u_i(a, \theta) \geq$$
$$\geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^{T-1}) g * s'(a \mid \theta) u_i(a, \theta),$$

where $\bar{s}_j = s_j$ if $j \in C_i(m^{t-1}, s, \theta_i)$ and $\bar{s}_j(m_j \mid \theta_j) = 1$ for all $\theta_j \in \Theta_j$ if $j \notin C_i(m^{t-1}, s, \theta_i)$ with $m_j^{T-1} = [m_j]$. In words: in each round each player is supposed to choose the initial move realized by applying s in the first round. If some player i deviates from this rule, some other player j may or may not be able to detect the deviation. If the deviation is undetected, then j's beliefs about i's play in the terminal round are that i will choose the first round move. If the deviation is detected, j ignores all information conveyed in the cheap play rounds, and expects i to play according to s in the terminal round. By posterior implementability, any move chosen by any i as part of the equilibrium s is a best response to any $(n-1)$-tuple of moves assigned positive probability by the same equilibrium, regardless of the information conveyed by observation of the $(n-1)$-tuple. Hence, by construction of the beliefs, we have sequential rationality.

Thus, we have:

for all $i \in N$, for all $h_i = (\theta_i, m^{t-1}) \in \mathcal{H}_i$, for all $\sigma'_i \in \Sigma_i$, $\sigma' \equiv (\sigma'_i, \sigma_{-i})$,

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^t) g * \sigma(a \mid \theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m_{-i}^t) g * \sigma'(a \mid \theta) u_i(a, \theta).$$

By definition, $(\sigma, \beta) \in \mathcal{E}_{\Sigma \times \mathbb{B}}(\Gamma^T)$ and satisfies NRM-stationarity. By construction $g * s = g * \sigma$. Thus, $\mathcal{E}^{PI}(\Gamma) \subseteq \mathcal{E}^{NRM}(\Gamma^T)$. □

Lemma 2: $\mathcal{E}^{NRM}(\Gamma^T) \subseteq \mathcal{E}^{PI}(\Gamma)$.

Proof of Lemma 2: Choose $(\sigma, \beta) \in \mathcal{E}_{\Sigma \times \mathbb{B}}^{NRM}(\Gamma^T)$, with $\sigma = (s^t)_{t=1}^T$. By sequential rationality, and NRM-stationarity, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m \in M$, such that $s^1(m \mid \theta) > 0$, for all $\tilde{m}_i \in M_i$, [1] must hold:

$$\sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m_{-i}]) u_i(g(m), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m_{-i}]) u_i(g(m_{-i}, \tilde{m}_i), \theta). \qquad [1]$$

But by NRM-stationarity,

$$\beta_i(\theta, [m_{-i}]) = \frac{s_{-i}^1(m_{-i} \mid \theta_{-i}) \pi_i(\theta)}{\sum_{\theta'_{-i} \in \Theta_{-i}} s_{-i}^1(m_{-i} \mid \theta'_{-i}) \pi_i(\theta'_{-i}, \theta_i)}$$
$$= [s_{-i}^1(m_{-i} \mid \theta_{-i}) \pi_i(\theta)] K(m_{-i}),$$

where $K(m_{-i})$ is a constant term.

Hence, by [1] the posterior implementability criterion is met. Next, we need to check that $\mathcal{E}^{NRM}(\Gamma^T) \subseteq \mathcal{E}(\Gamma)$.

By substituting in [1] and multiplying through by $[K(m_{-i})]^{-1}$, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m \in M$ such that $s^1(m \mid \theta) > 0$, for all $\tilde{m}_i \in M_i$, [2] must hold:

$$\begin{aligned} &\sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m), \theta) \geq \\ &\geq \sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m_{-i}, \tilde{m}_i), \theta). \end{aligned} \qquad [2]$$

Since [2] is true for each $m_{-i} \in M_{-i}$, such that $s^1_{-i}(m_{-i} \mid \theta_{-i}) > 0$, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m_i \in M_i$ such that $s^1_i(m_i \mid \theta_i) > 0$, and for all $\tilde{m}_i \in M_i$, [3] must hold:

$$\begin{aligned} &\sum_{m_{-i} \in M_{-i}} \sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m), \theta) \geq \\ &\geq \sum_{m_{-i} \in M_{-i}} \sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m_{-i}, \tilde{m}_i), \theta). \end{aligned} \qquad [3]$$

Thus, $s^1 \in \mathcal{E}_S(\Gamma)$. By *NRM*-stationarity, $g * \sigma = g * s^1$. Hence $\mathcal{E}^{II}(\Gamma^T) \subseteq \mathcal{E}^{PI}(\Gamma)$. □

4 Equilibria with No Regret in Assessments

In this section, we consider "no regret in assessments" (NRA) and the associated notion of stationarity is denoted NRA-stationarity.

Given $(\sigma, \beta) \in \mathcal{E}_{\Sigma \times \mathbb{B}}(\Gamma^t)$ and $t \in \{1, ..., T\}$, let

$$\tilde{M}^t(\sigma, \beta) \equiv \{m^t \in M^t : m^t \text{ is realized with positive probability in the equilibrium } (\sigma, \beta)\}.$$

A sequential equilibrium (σ, β) is *NRA-stationary* if $\sigma = (s^t)_{t=1}^T$ satisfies:

$\forall t \in \{2, ..., T\}, \forall i \in N, \forall \theta_i \in \Theta_i, \forall m^{t-1} \in \tilde{M}^{t-1}(\sigma, \beta), s^t_i(\theta_i, m^{t-1}) = s^1_i(\theta_i)$.

$\mathcal{E}^{NRA}_{\Sigma \times \mathbb{B}}(\Gamma^T) \subseteq \Sigma \times \mathbb{B}$ denotes the set of NRA-stationary sequential equilibria of Γ^T, with $\mathcal{E}^{NRA}(\Gamma) \equiv \{g * \sigma : \exists \beta \in \mathbb{B} \text{ such that } (\sigma, \beta) \in \mathcal{E}^{NRA}_{\Sigma \times \mathbb{B}}(\Gamma^T)\}$ denoting the set of NRA-stationary sequential equilibrium outcomes of Γ^T.

We shall, first, make the following observation:

Observation 1: There is no logical relationship between $\mathcal{E}^{NRA}(\Gamma^T)$ and $\mathcal{E}^{PI}(\Gamma)$ for any $T > 1$.

Proof: Consider the following examples. The first example shows that $\mathcal{E}^{NRA}(\Gamma^T)$ does not contain $\mathcal{E}^{PI}(\Gamma)$, and the second example shows that the reverse containment may not be expected either. Any non-genericity in the payoffs is purely for the purposes of keeping the examples as simple as possible and is not critical for the arguments.

Example 1: Consider the following game, Γ.

	L	R
U	5, 1	2, 1
D	4, 1	0, 1

(θ_1, θ_2)

	L	R
U	0, 1	4, 1
D	2, 1	5, 1

(θ_1, θ_2')

Fig. 2

The players are labelled 1 and 2 and they choose from the sets $\{U, D\}$ and $\{L, R\}$ respectively. Player 2 has private information; her set of possible private observations is $\{\theta_2, \theta_2'\}$. Player 1 is uninformed and has a single possible observation. Player 1's posterior distribution on Θ_2 is $\pi_1(\theta_2) = \pi_1(\theta_2') = \frac{1}{2}$.

Trivially, player 2 is indifferent among her strategies. Consider the following Bayesian equilibrium s:

$$s_1(U \mid \theta_1) = \frac{1}{2}.$$

$$s_2(L \mid \theta_2) = \frac{2}{3}; \quad s_2(L \mid \theta_2') = \frac{1}{3}.$$

Both L and R are assigned positive probability by the equilibrium. Conditional upon the observation of L, the probability that player 1 assigns to θ_2 is $\frac{2}{3}$ and in the event that R is observed, the probability assigned to θ_2 is $\frac{1}{3}$. Upon observing either one of player 2's moves, the payoff to player 1 from choosing either U or D is $\frac{10}{3}$. Hence, the Bayesian equilibrium given above is also posterior implementable.

The equilibrium s induces the following distribution on the outcome space in each state.

	L	R
U	$\frac{1}{3}$	$\frac{1}{6}$
D	$\frac{1}{3}$	$\frac{1}{6}$

(θ_1, θ_2)

	L	R
U	$\frac{1}{6}$	$\frac{1}{3}$
D	$\frac{1}{6}$	$\frac{1}{3}$

(θ_1, θ_2')

Fig. 3

Suppose that there exists a *NRA*-stationary strategy in the game Γ^T, denoted $\sigma = (\hat{s}^t)_{t=1}^T$, that also yields the same distribution on outcomes. By *NRA*-stationarity, $\hat{s}^1$ must satisfy:

$$\begin{bmatrix} \hat{s}_1^1(U \mid \theta_1) \\ \hat{s}_1^1(D \mid \theta_1) \end{bmatrix} \begin{bmatrix} \hat{s}_2^1(L \mid \theta_2) & \hat{s}_2^1(R \mid \theta_2) \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{6} \end{bmatrix}$$

$$\begin{bmatrix} \hat{s}_1^1(U \mid \theta_1) \\ \hat{s}_1^1(D \mid \theta_1) \end{bmatrix} \begin{bmatrix} \hat{s}_2^1(L \mid \theta_2') & \hat{s}_2^1(R \mid \theta_2') \end{bmatrix} = \begin{bmatrix} \frac{1}{6} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{3} \end{bmatrix}$$

The unique solution to the equations above is $\hat{s}^1 = s$. However, σ is not an equilibrium strategy. Suppose that player 2 chooses the play L in the first round. Player 1's payoff from playing U in the second round is:

$$\frac{2}{3}\left(\frac{2}{3}(5) + \frac{1}{3}(2)\right) + \frac{1}{3}\left(\frac{1}{3}(0) + \frac{2}{3}(4)\right) = \frac{32}{9},$$

which is strictly greater than the payoff from playing D in the second round, given by:

$$\frac{2}{3}\left(\frac{2}{3}(4) + \frac{1}{3}(0)\right) + \frac{1}{3}\left(\frac{1}{3}(2) + \frac{2}{3}(5)\right) = \frac{28}{9}.$$

The strategy σ_1 is not a best response to σ_2 for the case $T = 2$. It can be checked that the same conclusion would be obtained for arbitrary $T(> 1)$. Thus, $\mathcal{E}^{PI}(\Gamma)$ is not a subset of $\mathcal{E}^{NRA}(\Gamma^T)$.

Example 2: Consider the following game, Γ.

	L	R
U	0, 1	2, 1
D	2, 1	0, 1

(θ_1, θ_2)

	L	R
U	2, 1	0, 1
D	0, 1	2, 1

(θ_1', θ_2)

Fig. 4

The players are labelled 1 and 2 and they choose from the sets $\{U, D\}$ and $\{L, R\}$ respectively. Player 1 has private information; his set of possible private observations is $\{\theta_1, \theta_1'\}$. Player 2 is uninformed. Player 2's posterior distribution on Θ_1 is $\pi_2(\theta_1) = \pi_2(\theta_1') = \frac{1}{2}$.

Trivially, player 2 is indifferent among her strategies. Consider the following Bayesian equilibrium s:

$$s_1(U \mid \theta_1) = \frac{2}{3}; \quad s_1(U \mid \theta_1') = \frac{1}{3}.$$

$$s_2(L \mid \theta_2) = \frac{1}{2}.$$

It is easily checked that $\sigma = (s^t = s)_{t=1}^T$ is a sequential equilibrium strategy profile (which satisfies *NRA*-stationarity, by definition) as well. For any $T > 1, T-1$ rounds of cheap play do not invalidate the best-response property of either player's strategy.

The equilibrium σ induces the following distribution on the outcome space in each state. (Fig. 5)

Let $\hat{s} \in S$ be a strategy profile in the game Γ. If $\hat{s}$ must yield the distribution over outcomes given above, it must satisfy:

$$\begin{bmatrix} \hat{s}_1(U \mid \theta_1) \\ \hat{s}_1(D \mid \theta_1) \end{bmatrix} [\hat{s}_2(L \mid \theta_2)\ \hat{s}_2(R \mid \theta_2)] = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} \end{bmatrix}$$

$$\begin{bmatrix} \hat{s}_1(U \mid \theta_1') \\ \hat{s}_1(D \mid \theta_1') \end{bmatrix} [\hat{s}_2(L \mid \theta_2)\ \hat{s}_2(R \mid \theta_2)] = \begin{bmatrix} \frac{1}{6} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{3} \end{bmatrix}$$

	L	R
U	$\frac{1}{3}$	$\frac{1}{3}$
D	$\frac{1}{6}$	$\frac{1}{6}$

(θ_1, θ_2)

	L	R
U	$\frac{1}{6}$	$\frac{1}{6}$
D	$\frac{1}{3}$	$\frac{1}{3}$

(θ_1', θ_2)

Fig. 5

The unique solution to the equations above is $\hat{s} = s$. However, s is not posterior implementable. Suppose that player 1's observation is θ_1 and he observes player 2's play of R. Player 1's payoff from playing U clearly dominates the payoff from playing D, given this observation.

Thus, $\mathcal{E}^{NRA}(\Gamma^T)$ is not a subset of $\mathcal{E}^{PI}(\Gamma)$.

We conclude with the following proposition:

Theorem 2: There is no refinement of $\mathcal{E}(\Gamma)$, say $\mathcal{E}^*(\Gamma)$, such that $\mathcal{E}^*(\Gamma) = \mathcal{E}^{NRA}(\Gamma^T)$ for all $T > 1$.

Proof: The argument is based on the simple principle that the set of outcomes realized in $\mathcal{E}^{NRA}(\Gamma^T)$ is dependent on length of the cheap play game, T. Hence, the normal-form simultaneous move game Γ which ignores this information cannot have a single solution concept whose outcomes are identical to $\mathcal{E}^{NRA}(\Gamma^T)$ for arbitrary T. The following example shows that the value of T affects the set $\mathcal{E}^{NRA}(\Gamma^T)$.

Example 3: Consider the following game, Γ. (Fig. 6)

The players are labelled 1 and 2 and they choose from the sets $\{U, D\}$ and $\{L, R\}$ respectively. Player 1 has private information; his set of possible private observations is $\{\theta_1, \theta_1'\}$. Player 2 is uninformed. Player 2's posterior distribution on Θ_1 is $\pi_2(\theta_1) = \frac{1}{3}$.

Consider the following strategies:

$$s_1^1(U \mid \theta_1) = \frac{2}{3}; \quad s_1^1(U \mid \theta_1') = \frac{1}{3}.$$

$$s_2^1(L \mid \theta_2) = \frac{1}{2}.$$

	L	R
U	1, 1	1, 0
D	1, 1	1, 0

(θ_1, θ_2)

	L	R
U	1, 0	1, 1
D	1, 0	1, 1

(θ_1', θ_2)

Fig. 6

By construction, after one round of cheap play, regardless of whether U or D is played, player 2 assigns equal probability to the two states θ_1 and θ_1'. It may be checked that for $T = 2, \sigma = (s^t)_{t=1,2} \in \mathcal{E}_{\Sigma}^{NRA}(\Gamma^T)$. However, since, in round 1, L i a dominated strategy in expectation for Player 2, $s^1 \notin \mathcal{E}(\Gamma)$. □

5 Concluding Remarks

To summarize, we have focused on a special class of games with pre-play communication via cheap play. A natural focal subset of equilibria of such games are ones that involve regret-freeness on the part of the players. Such equilibria have been analyzed by Green and Laffont (1987) using the concept of posterior implementability. An implication of this condition is that the extended-form of the game induced by the rounds of communication is irrelevant and it is enough to simply focus on a subset of the equilibria of the one-shot game. We have shown that posterior implementability completely characterizes one notion of regret-freeness: "no regret in moves." We argue that there is an alternative notion of regret-freeness – "no regret in assessments" which is not captured by posterior implementability or any condition that ignores the extended-form of the induced game.

We have highlighted two approaches to the question of regret-freeness. There are, of course, other forms of stationary play, for example:

- Players choose to play the move that they had chosen in the immediately preceding round, regardless of the information conveyed by earlier choices.
- In every round, the players choose the move they made in the first round, regardless to the information they receive in the interim.

- In any round players choose one of the moves chosen in the past.
- In every round of play, the players choose the move they made in the first round, regardless of the information they receive in the interim, provided in the past they have chosen the first round move at every round.

Note that in all these cases, the move that is being repeated in subsequent rounds may or may not occur with positive probability in equilibrium.

We could probably add many other specifications of behaviour off the equilibrium path. However, such specifications are not particularly interesting. All such conditions require, among other things, that for any realization of play, say m, in some cheap play round, the move m must be repeated in the final round. By definition of equilibrium, this requires that for each i, each $m_i \in M_i$ must be a best response (in expected utility) to every $m_{-i} \in M_{-i}$, i.e. each $m \in M$ must yield the same expected payoff to any player i. Any perturbation of the payoffs will upset such a delicate structure. Generically, this implies that such restrictions of stationarity off-the-equilibrium path will almost always yield an empty equilibrium set.

In Chakravorti (1992), a special case of the analysis given earlier is applied. In the paper, the equivalence between the stationary pure strategy Bayesian equilibria of the extended game and pure strategy posterior implementable equilibria of the static game is considered. Note that all of the definitions of stationarity given in this paper are equivalent when attention is restricted to pure strategies and to regret-freeness only in parts of the game that are on the equilibrium path.

References

Aumann R (1976) Agreeing to disagree. Annals of Stat 4 : 1236–1239

Chakravorti B (1990) Efficiency and mechanisms with no regret. Int Econ Review 33 : 45–60

Farrell J (1982) Communication in games. Mimeo, MIT

Green J, Laffont J-J (1987) Posterior implementability in a two-person decision problem. Econometrica 55 : 69–94

Received October 1992