Belief Free Incomplete Information Games

By

Dirk Bergemann and Stephen Morris

September 2007

COWLES FOUNDATION DISCUSSION PAPER NO. 1629

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

Belief Free Incomplete Information Games*

Dirk Bergemann[†] Stephen Morris[‡]

First Version:July 2007
This Version: September 2007

Abstract

We consider the following belief free solution concepts for games with incomplete information: (*i*) incomplete information rationalizability, (*ii*) incomplete information correlated equilibrium and (*iii*) ex post equilibrium. We present epistemic foundations for these solution concepts and establish relationships between them. The properties of these solution concepts are further developed in supermodular games and potential games.

KEYWORDS: Correlated Equilibrium, Rationalizability, Ex Post Equilibrium, Belief Free, Types, Payoff Types, Belief Types, Supermodular Games, Potential Games.

JEL CLASSIFICATION: C79, D82

*This research is supported by NSF Grant #SES-0518929.
[†]Department of Economics, Yale University, New Haven CT 06511, dirk.bergemann@yale.edu.
[‡]Department of Economics, Princeton University, Princeton NJ 08544, smorris@princeton.edu.

1 Introduction

In games with incomplete information, the private information of each agent is represented by his type. The type of each agent contains information about the preferences of the agents *and* information about the beliefs of the agents. The type of agent can therefore be decomposed into a payoff type and a belief type. The *payoff type* of an agent embodies information about the players payoff and the *belief type* embodies information about the players' belief and higher order beliefs. In games with incomplete information, the strategy of an agent may naturally depend on his entire type, namely his payoff type and his belief type. In large type spaces, and certainly in the universal type space, each agent may have different belief types associated with the same payoff type. The prediction of play in a game of incomplete information may therefore be sensitive to the payoff type as well as the belief type.

In this paper we consider three solution concepts for games of incomplete information which depend only on the payoff types but not on the belief types of the agent. The three solution concepts are (i) incomplete information rationalizability, (ii) incomplete information correlated equilibrium and (iii) ex post equilibrium. As these solution concepts do not depend on the beliefs and higher order beliefs of the agent, we refer to them as belief free solution concepts. Having defined the solution concepts, we give their epistemic foundations and establish relationships between the solution concepts. We then use these solution concepts in supermodular games and potential games with incomplete information. We should emphasize that these solution concepts have already been defined in the literature. Rather, the contribution of this paper is to present epistemic foundations for these solution concepts and establish their relationship to each other.

The notion of ex post equilibrium is the most demanding among the three solution concepts. The term "ex post equilibrium" is due to Cremer and McLean (1985).[1] It requires that in equilibrium, the strategic choice of each type of each player remains a best response ex post, that is after the payoff type of each agent has become public. The ex post equilibrium is frequently used as solution concept in mechanism design where the game is specifically designed so as to support truthtelling as an ex post equilibrium (e.g. Dasgupta and Maskin (2000), Jehiel and Moldovanu (2001) and Bergemann and Välimäki (2002)). In earlier work, Bergemann and Morris (2005) showed that the ex post equilibrium can

[1] Earlier, D'Aspremont and Gerard-Varet (1979) define the same notion in a private value environment as uniform equilibrium and Holmstrom and Myerson (1983) refer to uniform incentive compatibility in the direct mechanism of an interdependent value environment.

be understood as a solution concept which embeds robustness to beliefs and higher order beliefs in the following sense: a social choice function can be truthfully implemented in every type space in an interim equilibrium if and only if it can be truthfully implemented as an ex post equilibrium. The ex post equilibrium is thus a belief free solution concept as it requires that the strategies of the players remain an equilibrium for all possible beliefs and higher order beliefs. The objective of this paper is to describe belief free solution concepts in a general game theoretic environment. Here, and in contrast to the mechanism design environment, the notion of an ex post equilibrium is very demanding and in many games an ex post equilibrium does not exist. We shall therefore define and analyze weaker solution concepts, namely incomplete information correlated equilibrium and incomplete information rationalizability. In games with a finite number of actions and a finite number of payoff types, the existence of these solution concepts is proved by construction.

We say that an action is incomplete information rationalizable for a payoff type of an agent if it survives the process of iteratively elimination of dominated strategies. The notion of incomplete information rationalizability is belief free as the candidate action needs only to be a best response to *some* beliefs about the other agents actions and payoff types. This solution concept was studied under this name in Battigalli (1999); his work was incorporated in Battigalli and Siniscalchi (2003), where "Δ-rationalizability" is used to refer to a general dynamic version of rationalizability and Δ refers to common knowledge restrictions on beliefs. It is important to emphasize that the rationalizability of an action is defined with respect to the payoff type rather than the type of an agent as in the notions of interim rationalizability of Dekel, Fudenberg, and Morris (2006) and Ely and Peski (2006).

We say that a strategy profile forms an incomplete information correlated equilibrium if there exists some distribution over payoff types and actions such that every action taken by a payoff type of an agent is a best response given the distribution over payoff types and actions. The notion of incomplete information correlated equilibrium corresponds (up to some minor differences) to the universal Bayesian solution suggested in Forges (1993).

The epistemic foundations of incomplete information rationalizability and incomplete information correlated equilibrium present the natural generalizations of their complete information counterparts by Brandenburger and Dekel (1987) and Aumann (1987), respectively. In proposition 1 we show that a specific action of a payoff type is incomplete information rationalizable if and only if there exists type space and an interim equilibrium such that the message is an equilibrium action for a type with a given payoff type in the type space. Also in proposition 1, we show that an action is an element of an incomplete information correlated equilibrium for a payoff type of an agent if and only if there exists a type space with a common prior for which the specific message is a Bayes Nash equilibrium action for a type with that payoff type in the hierarchical type space. With respect to the ex post equilibrium, we show that a strategy profile forms an ex post equilibrium if and only if the strategies of the payoff types remain interim equilibrium strategies on all type spaces.

The three solution concepts are nested in the appropriate manner. We show that if for a given payoff type, an action is an element of an ex post equilibrium profile, then it is also an element of an incomplete information correlated equilibrium for the given payoff type. Likewise, if an action is an element of an incomplete information correlated equilibrium for a given payoff type, then it is also incomplete information rationalizable for the given payoff type.

In the case of supermodular games, the relationships between these three solution concepts can be further strengthened. In particular, we show that in generic supermodular games the set of rationalizable actions are single valued for all agents and all payoff types if and only if the set of correlated equilibrium actions are single valued. Moreover, if indeed they are single valued, then they form an ex post equilibrium. A second important class of games in this context is the class of Bayesian potential games. We show that if a game has a smooth concave potential for every payoff type profile and also has an ex post equilibrium, then the ex post equilibrium forms the unique incomplete information correlated equilibrium.

We restrict our attention to solution concepts for normal form (or static) games. In contrast, Kalai (2004) and Borgers and McQuade (2007) develop belief free solution concepts for extensive form games.

The remainder of the paper is organized as follows. Section 2 presents the belief free solution concepts in a finite environment. Section 3 gives the relevant definitions for compact action and payoff type spaces. We also consider a common interest game with quadratic payoffs to apply the solution concepts. Section 4 presents the epistemic foundations of the solution concepts. Section 5 establishes some relations between these solution concepts in general games. Section 6 obtains additional results in supermodular games. Section 7 considers Bayesian potential games and presents conditions for a unique incomplete information correlated equilibrium. Section 8 concludes.

2 Belief Free Solution Concepts

There are I players. Player i chooses an action $a_i \in A_i$ and has a *payoff type* $\theta_i \in \Theta_i$, where A_i and Θ_i are finite sets. To ensure a clear comparison with the literature, we also allow for uncertainty about "unknown payoff relevant variables," states that are not known by any agent; let Θ_0 be a finite set of unknown payoff relevant states, with typical element θ_0. Now $\Theta = \Theta_0 \times \Theta_1 \times ... \times \Theta_I$ is the relevant uncertainty space. We write $a_{-i} = (a_1, .., a_{i-1}, a_{i+1}, ..., a_I)$, $\theta_{-i} = (\theta_0, \theta_1, .., \theta_{i-1}, \theta_{i+1}, ..., \theta_I)$ and $\theta_{-\{0,i\}} = (\theta_1, .., \theta_{i-1}, \theta_{i+1}, ..., \theta_I)$. Now a belief free incomplete information game is given by the payoff functions $u = (u_i)_{i=1}^I$ where each $u_i : A \times \Theta \to \mathbb{R}$. Thus "payoff types" embody information about player's payoffs but we have no information about players' beliefs or higher order beliefs about other players' payoff types or unknown payoff relevant states.

We report natural generalizations of the complete information solution concepts of (correlated) rationalizability, correlated equilibrium and Nash equilibrium.

Definition 1 (Incomplete Information Rationalizability)
The incomplete information rationalizable actions $R = (R_i)_{i=1}^I$, each $R_i : \Theta_i \to 2^{A_i}/\varnothing$, are defined recursively as follows. Let $R_i^0(\theta_i) = A_i$,

$$R_i^{k+1}(\theta_i) = \left\{ a_i \in R_i^k(\theta_i) \;\middle|\; \begin{array}{l} \text{there exists } \mu_i \in \Delta(A_{-i} \times \Theta_{-i}) \text{ such that} \\ \text{(1) } \mu_i(a_{-i}, \theta_{-i}) > 0 \Rightarrow a_j \in R_j^k(\theta_j) \text{ for each } j \neq i \\ \text{(2) } a_i \in \arg\max\limits_{a_i'} \sum\limits_{a_{-i}, \theta_{-i}} u_i((a_i', a_{-i}), (\theta_i, \theta_{-i})) \mu_i(a_{-i}, \theta_{-i}) \end{array} \right\}$$

for each $k = 1, 2, ...,$ and

$$R_i(\theta_i) = \bigcap_{k \geq 0} R_i^k(\theta_i).$$

This solution concept was studied under this name in Battigalli (1999); this work was incorporated in Battigalli and Siniscalchi (2003), where "Δ-rationalizability" is used to refer to a general dynamic version of rationalizability and Δ refers to common knowledge restrictions on beliefs. Their definition reduces to the one above in a static setting when Δ is the empty set. Battigalli and Siniscalchi (2003) assumed that all payoff relevant variables are known by some agent, so, in our language, the set Θ_0 is a singleton. This solution concept has played an important role in our work on robust full implementation (see Bergemann and Morris (2001), Bergemann and Morris (2007a) and Bergemann and Morris (2007b)). Note that $R_i(\theta_i)$ is non-empty for each i and $\theta_i \in \Theta_i$ by construction.

Definition 2 (Incomplete Information Correlated Equilibrium)
A probability distribution $\mu \in \Delta(A \times \Theta)$ is an incomplete information correlated equilibrium (ICE) of u if for each i, θ_i, a_i and a_i',

$$\begin{aligned} & \sum_{a_{-i},\theta_{-i}} u_i\left(\left(a_i, a_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) \mu\left(\left(a_i, a_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) \\ \geq & \sum_{a_{-i},\theta_{-i}} u_i\left(\left(a_i', a_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) \mu\left(\left(a_i, a_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right). \end{aligned}$$

Many incomplete information versions of correlated equilibrium have been defined: Forges (1993) proposed five "legitimate" ones. The above definition corresponds to the most general in Forges (1993) - i.e., the universal Bayesian approach of section 6. She does not explicitly incorporate "payoff types", i.e., payoff relevant variables that are known by one agent. Thus, in our language, it is as if each Θ_i were a singleton. She notes in proposition 4 that the set of payoffs that might arise under this solution concept is equal to the set of payoffs that might arise under more stringent solutions concepts if players are allowed to observe some sufficiently rich private signals. Forges (1993) deals with two player games, but Forges (2006) discusses the straightforward extension to many players; here she refers to the solution concept as the "Bayesian solution" of the game. The type correlated equilibria of Cotter (1994) are essentially equivalent to this definition, with the proviso that he fixes the prior distribution on Θ, so the equilibrium describes a distribution on A conditional on each realized θ.[2]

Every game has an ICE: it is enough to fix any $\theta^* \in \Theta$, let $\alpha = (\alpha_i)_{i=1}^I \in \underset{i=1}{\overset{I}{\times}} \Delta(A_i)$ be any Nash equilibrium of the complete information game $(u_i(\cdot, \theta^*))_{i=1}^I$, and let

$$\mu(a,\theta) = \begin{cases} \prod_{i=1}^{I} \alpha_i(a_i), \text{ if } \theta = \theta^*, \\ 0, \text{ otherwise.} \end{cases}$$

Moreover, for any $\psi \in \Delta(\Theta)$, there exists an ICE whose marginal on Θ is ψ: for every $\theta \in \Theta$, let

$$\alpha^\theta(\theta) = \left(\alpha_i^\theta\right)_{i=1}^I \in \underset{i=1}{\overset{I}{\times}} \Delta(A_i)$$

[2] In an earlier paper, Cotter (1991), analyzes the notion of a correlated equilibrium with type dependent strategies. In the correlated equilibrium with type dependent strategies the randomization device is restricted to be independent of the type of each player. In the current definition we allow the correlation device to depend on the type profile realization.

be any Nash equilibrium of the complete information game $\left(u_i\left(\cdot,\theta^*\right)\right)_{i=1}^I$, and let

$$\mu\left(a,\theta\right)=\psi\left(\theta\right)\prod_{i=1}^{I}\alpha_i^{\theta}\left(a_i\right).$$

We denote the set of actions taken by payoff type θ_i of agent i in some incomplete information correlated equilibrium by $C_i\left(\theta_i\right)$, so formally we have

$$C_i\left(\theta_i\right)=\left\{a_i\in A_i\left|\exists\text{ ICE }\mu\text{ and }\left(a_{-i},\theta_{-i}\right)\in A_{-i}\times\Theta_{-i}\text{ such that }\mu\left(\left(a_i,a_{-i}\right),\left(\theta_i,\theta_{-i}\right)\right)>0\right.\right\}. \tag{1}$$

We define an ex post equilibrium. A *payoff type strategy* for player i is a function $s_i:\Theta_i\to A_i$.

Definition 3 (Ex Post Equilibrium)
A payoff type strategy profile $s^*=\left(s_i^*\right)_{i=1}^I$ is an ex post equilibrium if for all i and all θ, we have

$$u_i\left(\left(s_i^*\left(\theta_i\right),s_{-i}^*\left(\theta_{-\{0,i\}}\right)\right),\theta\right)\geq u_i\left(\left(a_i,s_{-i}^*\left(\theta_{-\{0,i\}}\right)\right),\theta\right)$$

for all $a_i\in A_i$.

"Most" games will not have ex post equilibria. But the solution concept has been extensively studied in the mechanism design literature (where the game is constructed to have ex post equilibria). In particular, truthtelling is an ex post equilibrium in a direct mechanism if and only if it is ex post incentive compatible. Holmstrom and Myerson (1983) is an early reference dealing with ex post incentive compatibility (under the name "uniform incentive compatibility"). Ex post equilibrium has recently been studied in general game theory contexts (see, e.g., Kalai (2004) and Borgers and McQuade (2007)).

3 Compact Action and Type Spaces and an Example

Let the framework be as before except that each A_i and Θ_i are compact intervals of the real line and each u_i is continuous in a and θ. We let Θ_0 be a singleton and thus do not refer to it.

3.1 Solution Concepts

The definition of rationalizability becomes:

Definition 4 (Incomplete Information Rationalizability)
The incomplete information rationalizable actions $R = (R_i)_{i=1}^{I}$, each $R_i : \Theta_i \to 2^{A_i} / \varnothing$, are defined recursively as follows. Let $R_i^0(\theta_i) = A_i$,

$$R_i^{k+1}(\theta_i) = \left\{ a_i \in R_i^k(\theta_i) \;\middle|\; \begin{array}{l} \text{there exists } \mu_i \in \Delta(A_{-i} \times \Theta_{-i}) \text{ such that} \\ \text{(1) } \mu_i\left[\left\{(a_{-i}, \theta_{-i}) : a_j \in R_j^k(\theta_j) \text{ for each } j \neq i\right\}\right] = 1 \\ \text{(2) } a_i \in \arg\max\limits_{a_i'} \int\limits_{a_{-i}, \theta_{-i}} u_i\left((a_i', a_{-i}), (\theta_i, \theta_{-i})\right) d\mu_i \end{array} \right\}$$

for each $k = 1, 2, ...,$ and

$$R_i(\theta_i) = \bigcap_{k \geq 0} R_i^k(\theta_i).$$

The compactness and continuity assumptions ensure that $R_i(\theta_i)$ is well-defined and trans-finite iterations are not required.

Definition 5 (Incomplete Information Correlated Equilibrium)
A probability distribution $\mu \in \Delta(A \times \Theta)$ is an incomplete information correlated equilibrium (ICE) of u if for each i and each measurable $\phi_i : A_i \times \Theta_i \to A_i$

$$\int\limits_{a,\theta} u_i\left((a_i, a_{-i}), (\theta_i, \theta_{-i})\right) d\mu \geq \int\limits_{a,\theta} u_i\left((\phi_i(a_i, \theta_i), a_{-i}), (\theta_i, \theta_{-i})\right) d\mu.$$

We define $C_i(\theta_i)$ - the set of actions that can be played by type θ_i in an incomplete information correlated equilibrium of game u - for the compact action and type spaces case of this section. We will say that $a_i^* \in C_i(\theta_i^*)$ if for each $\varepsilon > 0$, there exists an ICE μ with $\mu\left[\{(a, \theta) \,| a_i \in [a_i^* - \varepsilon, a_i^* + \varepsilon] \text{ and } \theta_i \in [\theta_i^* - \varepsilon, \theta_i^* + \varepsilon]\}\right] > 0$.

3.2 Quadratic Example

Let $A_i = \Theta_i = [0, 1]$ for all i and let

$$\begin{aligned} u_i(a, \theta) &= v(a, \theta) \\ &= -\sum_{j=1}^{I} (a_j - \theta_j)\left[(a_j - \theta_j) + \gamma \sum_{k \neq j} (a_k - \theta_k)\right] \\ &= -\sum_{j=1}^{I} (a_j - \theta_j)^2 - \gamma \sum_{j=1}^{I} (a_j - \theta_j) \sum_{k \neq j} (a_k - \theta_k), \end{aligned}$$

for some $\gamma \in \mathbb{R}$. Note that this is a common interest game. Now suppose agent i has type θ_i and belief $\mu_i \in \Delta(A_{-i} \times \Theta_{-i})$. Then his expected utility from choosing action a_i is

$$\mathbb{E}_{\mu_i}(v|a_i,\theta_i) = -\int\limits_{a_{-i},\theta_{-i}} \left(\sum_{j=1}^{I} (a_j - \theta_j)^2 - \gamma \sum_{j=1}^{I} (a_j - \theta_j) \sum_{k \neq j} (a_k - \theta_k) \right) d\mu_i$$

Now

$$\frac{d\mathbb{E}_{\mu_i}(v|a_i,\theta_i)}{da_i} = -2(a_i - \theta_i) - 2\gamma \mathbb{E}_{\mu_i}\left(\sum_{j \neq i} (a_j - \theta_j) \right)$$

Setting this equal to zero gives agent i's best response:

$$a_i = \theta_i - \gamma \mathbb{E}_{\mu_i}\left(\sum_{j \neq i} (a_j - \theta_j) \right).$$

In this game, we have

$$R_i(\theta_i) = \begin{cases} \{\theta_i\}, \text{ if } -\frac{1}{I-1} < \gamma < \frac{1}{I-1}, \\ [0,1], \text{ otherwise,} \end{cases}$$

and

$$C_i(\theta_i) = \begin{cases} \{\theta_i\}, \text{ if } -\frac{1}{I-1} < \theta_i < 1, \\ [0,1], \text{ otherwise,} \end{cases}$$

and there is a unique ex post equilibrium s^* with $s_i^*(\theta_i) = \theta_i$ for all $\theta_i \in [0,1]$.

The rationalizability claim can be shown by the following inductive step, let

$$R_i^k(\theta_i) \triangleq \left[\max\left\{0, \theta_i - (|\gamma|(I-1))^k\right\}, \min\left\{1, \theta_i + (|\gamma|(I-1))^k\right\} \right],$$

and was shown (in a mechanism design application) as the leading example in Bergemann and Morris (2007a). If $\gamma < 0$, then the game has strategic complementarities and well known arguments imply that there will not be a gap between extremal rationalizable outcomes and correlated equilibria. Thus $-\frac{1}{I-1} < \gamma$ will remain a tight characterization for correlated equilibria for $\gamma \leq 0$. If $\gamma \geq 1$, it is easy to show that every action can be element of a correlated equilibrium. Later arguments will establish the claim of $C_i(\theta_i) = \{\theta_i\}$ if $-\frac{1}{I-1} < \gamma < 1$.

4 Epistemic Foundations for the Solution Concepts

These belief free solution concepts are of interest not because we think that players don't have beliefs and higher order beliefs, but because we do not know what they are. In this section, we review results that explain why it make sense to use these solution concepts as a reduced form description of what might happen in more fully specified environments with beliefs and higher order beliefs. We return to the finite case to avoid technicalities.

4.1 Type Spaces

A type space $\mathcal{T}$ is defined as $\mathcal{T} \triangleq \left(T_i, \widehat{\pi}_i, \widehat{\theta}_i\right)_{i=1}^{I}$ where

1. T_i is a finite set of types
2. $\widehat{\pi}_i : T_i \to \Delta\left(T_{-i} \times \Theta_0\right)$ describes the beliefs of i's types
3. $\widehat{\theta}_i : T_i \to \Theta_i$ describes the payoff types of agent i's types

We write $T = T_1 \times \cdots \times T_I$ and $T_{-i} = T_1 \times \cdots \times T_{i-1} \times T_{i+1} \times \cdots \times T_I$; t_i is typical element of T_i, $t = (t_i)_{i=1}^{I}$ and $t_{-i} = (t_1, ..., t_{i-1}, t_{i+1}, ..., t_I)$; we let $\widehat{\theta} : T \to \Theta_{-0}$ be defined by $\widehat{\theta}\left(t\right) = \left(\widehat{\theta}_i\left(t_i\right)\right)_{i=1}^{I}$ and $\widehat{\theta}_{-i} : T_{-i} \to \Theta_{-\{0,i\}}$ be defined by $\widehat{\theta}_{-i}\left(t\right) = \left(\widehat{\theta}_j\left(t_j\right)\right)_{j \neq i}$.

A type space $\mathcal{T}$ has a *common prior* $\pi^* \in \Delta\left(T \times \Theta_0\right)$ if, for all i and $t_i \in T_i$,

$$\sum_{t_{-i},\theta_0} \pi^*\left(t_i, t_{-i}, \theta_0\right) > 0;$$

and

$$\widehat{\pi}_i\left(t_i\right)\left[t_{-i}, \theta_0\right] = \frac{\pi^*\left(t_i, t_{-i}, \theta_0\right)}{\sum\limits_{t'_{-i},\theta'_0} \pi^*\left(t_i, t'_{-i}, \theta'_0\right)}, \tag{2}$$

for all $t_{-i} \in T_{-i}$ and $\theta_0 \in \Theta_0$. A type space $\mathcal{T}$ is a *common prior* type space if there exists $\pi^* \in \Delta\left(T\right)$ such that $\mathcal{T}$ has common prior π^*. A type space $\mathcal{T}$ is a *payoff* type space if each $T_i = \Theta_i$ and each $\widehat{\theta}_i$ is the identity map. A type space $\mathcal{T}$ is a *full support* type space if $\widehat{\pi}_i\left(t_i\right)\left[t_{-i}, \theta_0\right] > 0$ for all i, t_i, t_{-i} and θ_0.

4.2 Interim Equilibrium

The belief free incomplete information game u and type space $\mathcal{T}$ together define an incomplete information game which may not have a common prior. A behavioral strategy of player i in type space $\mathcal{T}$ is given by a function $\sigma_i : T_i \to \Delta(A_i)$. We write Σ_i for the set of behavioral strategies of player i.

Definition 6 (Interim Equilibrium)
Strategy profile σ is an interim equilibrium of $(u, \mathcal{T})$ if for each i, $t_i \in T_i$, $a_i \in A_i$ with $\sigma_i(a_i|t_i) > 0$, and $a_i' \in A_i$,

$$\begin{aligned} & \sum_{a_{-i},t_{-i},\theta_0} \widehat{\pi}_i(t_i)[t_{-i},\theta_0] \left(\prod_{j\neq i} \sigma_j(a_j|t_j)\right) u_i\left((a_i, a_{-i}), \left(\theta_0, \widehat{\theta}_i(t_i), \widehat{\theta}_{-i}(t_{-i})\right)\right) \qquad (3) \\ \geq & \sum_{a_{-i},t_{-i},\theta_0} \widehat{\pi}_i(t_i)[t_{-i},\theta_0] \left(\prod_{j\neq i} \sigma_j(a_j|t_j)\right) u_i\left((a_i', a_{-i}), \left(\theta_0, \widehat{\theta}_i(t_i), \widehat{\theta}_{-i}(t_{-i})\right)\right). \end{aligned}$$

Definition 7 (Bayesian Nash Equilibrium)
Strategy profile σ is a Bayesian Nash equilibrium for type space $\mathcal{T}$ with common prior π^* if for each i and alternative strategy σ_i',

$$\begin{aligned} & \sum_{a,t,\theta_0} \pi^*(t,\theta_0) \left(\prod_{j=1}^{I} \sigma_j(a_j|t_j)\right) u_i\left(a, \left(\theta_0, \widehat{\theta}_i(t_i), \widehat{\theta}_{-i}(t_{-i})\right)\right) \\ \geq & \sum_{a,t,\theta_0} \pi^*(t,\theta_0)\, \sigma_i'(a_i|t_i) \left(\prod_{j\neq i} \sigma_j(a_j|t_j)\right) u_i\left(a, \left(\theta_0, \widehat{\theta}_i(t_i), \widehat{\theta}_{-i}(t_{-i})\right)\right). \end{aligned}$$

As is well known, this ex ante definition of a Bayesian Nash equilibrium is equivalent to interim equilibrium on common prior type spaces. But in the absence of a common prior, there is not a natural ex ante definition of the incomplete information equilibrium.

4.3 Epistemic Results

We denote the set of actions played by agent i with payoff type θ_i in some interim equilibrium on some type space $\mathcal{T}$ by $S_i(\theta_i)$. So

$$S_i(\theta_i) = \left\{ a_i \in A_i \;\middle|\; \begin{array}{l} \exists \text{ a type space } \mathcal{T}\text{, an interim equilibrium, } \sigma\text{, of } (u,\mathcal{T}), \\ \text{and a type } t_i \in T_i \text{ such that } \widehat{\theta}_i(t_i) = \theta_i \text{ and } \sigma_i(a_i|t_i) > 0 \end{array} \right\}.$$

We denote the set of actions played by agent i with payoff type θ_i in some interim equilibrium on some *common prior* type space $\mathcal{T}$ by $S_i^{CP}(\theta_i)$. So

$$S_i^{CP}(\theta_i) = \left\{ a_i \in A_i \left| \begin{array}{l} \exists \text{ a common prior type space } \mathcal{T}, \text{ an interim equilibrium, } \sigma, \text{ of } (u, \mathcal{T}), \\ \text{and a type } t_i \in T_i \text{ such that } \widehat{\theta}_i(t_i) = \theta_i \text{ and } \sigma_i(a_i | t_i) > 0 \end{array} \right. \right\}.$$

The following proposition records the straightforward incomplete information generalizations of the epistemic foundations for rationalizability and correlated equilibrium, respectively, from Brandenburger and Dekel (1987) and Aumann (1987).

Proposition 1 (Epistemic Foundations)
For all i and for all θ_i,

1. $R_i(\theta_i) = S_i(\theta_i)$;
2. $C_i(\theta_i) = S_i^{CP}(\theta_i)$.

Proof. For part (1), fix (i) a type space $\mathcal{T} = \left(T_i, \widehat{\theta}_i, \widehat{\pi}_i\right)_{i=1}^{I}$; (ii) an interim equilibrium σ of $(u, \mathcal{T})$ and (iii) a type $t_i^* \in T_i$ with (a) $\widehat{\theta}_i(t_i^*) = \theta_i^*$; and (b) $\sigma_i(t_i^*)[a_i] > 0$. Let

$$S_i(\theta_i) = \left\{ a_i \in A_i | \exists t_i \in T_i \text{ s.t. } \sigma_i(t_i)[a_i] > 0 \text{ and } \widehat{\theta}_i(t_i) = \theta_i \right\}.$$

Now for each i, θ_i in the range of $\widehat{\theta}_i$ and $a_i \in S_i(\theta_i)$, let

$$\lambda_i^{\theta_i, a_i}(\theta_{-i}, a_{-i}) \triangleq \sum_{\left\{t_{-i} \in T_{-i} \,:\, \widehat{\theta}_{-i}(t_{-i}) = \theta_{-\{0,i\}}\right\}} \widehat{\pi}_i(t_i)[t_{-i}, \theta_0] \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(t_j)[a_j] \right). \quad (4)$$

Now because σ is an equilibrium,

$$a_i \in \underset{a_i'}{\arg\max} \sum_{\theta_{-i}, a_{-i}} \lambda_i^{\theta_i, a_i}(\theta_{-i}, a_{-i}) \, u_i\left(\left(a_i', a_{-i}\right), (\theta_i, \theta_{-i})\right).$$

Now we show by induction on k that $S_i(\theta_i) \subseteq R_i^k(\theta_i)$ for all i, θ_i and k. This is true for $k = 0$ by definition. Suppose that it is true for k. Now $\lambda_i^{\theta_i, a_i}(\theta_{-i}, a_{-i}) > 0$ implies that $a_{-i} \in S_{-i}\left(\theta_{-\{0,i\}}\right)$, by construction, which implies $a_{-i} \in R_{-i}^k\left(\theta_{-\{0,i\}}\right)$ by the inductive hypothesis. Together with (4), this establishes $a_i \in R_i^{k+1}(\theta_i)$. This proves the induction. Now $a_i^* \in S_i(\theta_i^*) \subseteq R_i(\theta_i^*)$, proving the "if" claim of the proposition.

Conversely, suppose that $a_i^* \in R_i(\theta_i^*)$. Observe that for each i, θ_i and $a_i \in R_i(\theta_i)$, there exists $\lambda_i^{\theta_i, a_i} \in \Delta(\Theta_{-i} \times A_{-i})$ such that:

(a) $\lambda_i^{\theta_i, a_i}(\theta_{-i}, a_{-i}) > 0 \Rightarrow a_{-i} \in R_{-i}\left(\theta_{-\{0,i\}}\right)$

and

$$\text{(b) } a_i \in \arg\max_{a_i'} \sum_{\theta_{-i},a_{-i}} \lambda_i^{\theta_i,a_i}(\theta_{-i},a_{-i})\, u_i\left(\left(a_i',a_{-i}\right),(\theta_i,\theta_{-i})\right).$$

Now construct (i) a type space $\mathcal{T}$ with

$$\begin{aligned}
T_i &= \{(\theta_i,a_i) \in \Theta_i \times A_i \,|a_i \in R_i(\theta_i)\}, \\
\widehat{\theta}_i((\theta_i,a_i)) &= \theta_i, \text{ and} \\
\widehat{\pi}_i((\theta_i,a_i))\left[(\theta_j,a_j)_{j\neq i},\theta_0\right] &= \lambda_i^{\theta_i,a_i}(\theta_{-i},a_{-i});
\end{aligned}$$

where (a) above ensures that this is well-defined, and (ii) a strategy profile σ with

$$\sigma_i((\theta_i,a_i))\left[a_i'\right] = \begin{cases} 1, \text{ if } a_i' = a_i, \\ 0, \text{ otherwise.} \end{cases}$$

Now (b) ensures that σ is an equilibrium and by construction $t_i^* = (\theta_i^*,a_i^*) \in T_i$ with $\widehat{\theta}_i(t_i^*) = \theta_i$; and $\sigma_i(t_i^*)[a_i^*] > 0$. This establishes the "only if" part.

For part (2), first suppose that $a_i^* \in C_i(\theta_i^*)$. Thus there exists an ICE $\mu \in \Delta(A \times \Theta)$ and $(a_{-i},\theta_{-i}) \in A_{-i} \times \Theta_{-i}$ such that $\mu((a_i^*,a_{-i}),(\theta_i^*,\theta_{-i})) > 0$. Now we construct a type space. Let

$$T_i = \left\{ (a_i,\theta_i) \left| \sum_{a_{-i},\theta_{-i}} \mu((a_i,a_{-i}),(\theta_i,\theta_{-i})) > 0 \right. \right\}.$$

Define $\pi^* \in \Delta(T)$ by

$$\pi^*\left((a_i,\theta_i)_{i=1}^I,\theta_0\right) = \mu(a,\theta),$$

define each $\widehat{\pi}_i$ by

$$\widehat{\pi}_i(t_i)[t_{-i},\theta_0] \sum_{t_{-i}',\theta_0'} \pi^*\left(t_i,t_{-i}',\theta_0'\right) = \pi^*(t_i,t_{-i},\theta_0)$$

and define each $\widehat{\theta}_i$ by

$$\widehat{\theta}_i((a_i,\theta_i)) = \theta_i.$$

Observe that π^* is a common prior for $\mathcal{T}$ by construction. Now consider the strategy profile σ for the game $(u,\mathcal{T})$ defined by

$$\sigma_i\left(a_i'\middle|(a_i,\theta_i)\right) = \begin{cases} 1, \text{ if } a_i' = a_i \\ 0 \text{ otherwise} \end{cases}$$

By construction, σ is an interim equilibrium of $(u,\mathcal{T})$. Now consider the type $t_i = (a_i^*,\theta_i^*)$. We have constructed a common prior type space $\mathcal{T}$, an interim equilibrium, σ, of $(u,\mathcal{T})$, such that $\widehat{\theta}_i(t_i) = \theta_i^*$ and $\sigma_i(a_i^*\,|t_i) > 0$. Thus $a_i^* \in S_i^{CP}(\theta_i^*)$.

Conversely, suppose that $a_i^* \in S_i^{CP}(\theta_i^*)$. Thus there exists a common prior type space $\mathcal{T}$ with prior π^*, an interim equilibrium, σ, of $(u, \mathcal{T})$, such that $\widehat{\theta}_i(t_i) = \theta_i^*$ and $\sigma_i(a_i^* | t_i) > 0$. Define $\mu \in \Delta(A \times \Theta)$ by

$$\mu(a,\theta) = \sum_{\left\{t \middle| \widehat{\theta}(t)=\theta\right\}} \pi^*(t,\theta_0) \left[\prod_{i=1}^{I} \sigma_i(a_i|t_i)\right].$$

By construction, μ is an ICE. Since the prior assigns positive probability to every type, we have that $\mu((a_i^*, a_{-i}), (\theta_i^*, \theta_{-i})) > 0$ for some $(a_{-i}, \theta_{-i}) \in A_{-i} \times \Theta_{-i}$ and thus $a_i^* \in C_i(\theta_i^*)$. ■

Part (1) is a special case of propositions 4.2 and 4.3 in Battigalli and Siniscalchi (2003), and is a straightforward generalization of the complete information argument in Brandenburger and Dekel (1987). We recorded this result earlier as proposition 6 in the appendix of Bergemann and Morris (2007b). Part (2) is a straightforward generalization of the complete information argument of Aumann (1987); while Forges (1993) does not state a result in exactly this form, this argument captures the idea of the incomplete information generalization of Aumann's analysis in Forges (1993) section 6.

The next proposition describes how we can formalize the idea that the solution concept of ex post equilibrium makes sense if we want to identify behavior that will constitute an equilibrium whatever players' beliefs and higher order beliefs about others' payoff types and unknown payoff relevant states. Write $\sigma^{s,\mathcal{T}}$ for the strategy profile in $(u, \mathcal{T})$ induced by s, so that

$$\sigma^{s,\mathcal{T}}(s_i(\theta_i) | t_i) = \begin{cases} 1, \text{ if } \widehat{\theta}_i(t_i) = \theta_i, \\ 0, \text{ if } \widehat{\theta}_i(t_i) \neq \theta_i. \end{cases}$$

Proposition 2 (Ex Post Equilibrium)
The following are equivalent:

1. *s is an ex post equilibrium*
2. *$\sigma^{s,\mathcal{T}}$ is an interim equilibrium of $(u, \mathcal{T})$ for all type spaces $\mathcal{T}$*
3. *$\sigma^{s,\mathcal{T}}$ is an interim equilibrium of $(u, \mathcal{T})$ for all full support common prior payoff type spaces $\mathcal{T}$*

Proof. (1) $\Rightarrow$ (2). We verify that $\sigma^{s,\mathcal{T}}$ satisfies the definition of an interim equilibrium (see definition 6) on an arbitrary type space: for each i, t_i and a_i,

$$
\begin{aligned}
& \sum_{a_{-i},t_{-i},\theta_0} \widehat{\pi}_i(t_i)[t_{-i},\theta_0]\left(\prod_{j\neq i}\sigma_j^{s,\mathcal{T}}(a_j|t_j)\right)u_i\left(\left(s_i\left(\widehat{\theta}_i(t_i)\right),a_{-i}\right),\left(\theta_0,\widehat{\theta}_i(t_i),\widehat{\theta}_{-i}(t_{-i})\right)\right)\\
= & \sum_{\theta_{-i}}\left(\sum_{\left\{t_{-i}:\widehat{\theta}_{-i}(t_{-i})=\theta_{-\{0,i\}}\right\}}\widehat{\pi}_i(t_i)[t_{-i},\theta_0]\right)u_i\left(\left(s_i\left(\widehat{\theta}_i(t_i)\right),s_{-i}\left(\theta_{-\{0,i\}}\right)\right),\left(\widehat{\theta}_i(t_i),\theta_{-i}\right)\right),\\
& \text{by definition of } \sigma^{s,\mathcal{T}}\\
\geq & \sum_{\theta_{-i}}\left(\sum_{\left\{t_{-i}:\widehat{\theta}_{-i}(t_{-i})=\theta_{-\{0,i\}}\right\}}\widehat{\pi}_i(t_i)[t_{-i},\theta_0]\right)u_i\left(\left(a_i,s_{-i}\left(\theta_{-\{0,i\}}\right)\right),\left(\widehat{\theta}_i(t_i),\theta_{-i}\right)\right),\\
& \text{since } s \text{ is an ex post equilibrium}\\
= & \sum_{a_{-i},t_{-i},\theta_0} \widehat{\pi}_i(t_i)[t_{-i},\theta_0]\left(\prod_{j\neq i}\sigma_j^{s,\mathcal{T}}(a_j|t_j)\right)u_i\left(\left(a_i,a_{-i}\right),\left(\theta_0,\widehat{\theta}_i(t_i),\widehat{\theta}_{-i}(t_{-i})\right)\right).
\end{aligned}
$$

(2) $\Rightarrow$ (3). This is true by definition.

(3) $\Rightarrow$ (1). Let each $T_i = \Theta_i$ and $\widehat{\theta}_i$ be the identity map. For any $\pi^* \in \Delta_{++}(\Theta)$, consider the type space $\mathcal{T}_{\pi^*} = \left(T_i, \widehat{\pi}_i, \widehat{\theta}_i\right)_{i=1}^{I}$, where each $\widehat{\pi}_i$ is derived from π^* by Bayes rule as in equation (2) on page 10. This type space is a full support common prior payoff type space. Fix $\theta^* \in \Theta$ and let π^k be any sequence of full support priors with $\pi^k(\theta^*) \to 1$ as $k \to \infty$. Now suppose that (3) holds, so $\sigma^{s,\mathcal{T}}$ is an interim equilibrium of $(u, \mathcal{T}_{\pi^k})$ for each $k = 1, 2, ...$. Then for each i,

$$
\begin{aligned}
& \sum_{\theta_{-i}} \pi^k(\theta_i^*,\theta_{-i})\,u_i\left(\left(s_i(\theta_i^*),s_{-i}\left(\theta_{-\{0,i\}}\right)\right),(\theta_i,\theta_{-i})\right)\\
\geq & \sum_{\theta_{-i}} \pi^k(\theta_i^*,\theta_{-i})\,u_i\left(\left(a_i,s_{-i}\left(\theta_{-\{0,i\}}\right)\right),(\theta_i^*,\theta_{-i})\right)
\end{aligned}
$$

for all a_i and $k = 1, 2, ...$. Now $\pi^k(\theta^*) \to 1$ implies:

$$
u_i\left(\left(s_i(\theta_i^*),s_{-i}\left(\theta^*_{-\{0,i\}}\right)\right),(\theta_i^*,\theta_{-i}^*)\right) \geq u_i\left(\left(a_i,s_{-i}\left(\theta^*_{-\{0,i\}}\right)\right),(\theta_i^*,\theta_{-i}^*)\right).
$$

This conclusion holds for each $\theta^* \in \Theta$, i and a_i. This proves (1). ■

If we restrict attention to the special case where the game is a direct revelation mechanism and a planner is trying to implement a social choice function, this is a special case of results in our earlier work on robust mechanism design, Bergemann and Morris (2005): the arguments there do not depend on the mechanism design application and the arguments there prove the result in a general game theoretic setting.[3] In a private values environment, this result relates to earlier observation in the mechanism design literature showing the equivalence between "Bayesian equilibrium for all beliefs" and dominant strategies equilibrium, e.g., Ledyard (1979).

Borgers and McQuade (2007) have stated results along these lines for general games. In particular, they define s to be a *strongly information invariant equilibrium* if claim (2) of the proposition holds. Thus their proposition 1 establishes the equivalence of (1) and (2). They define s to be a *weakly information invariant equilibrium* if $\sigma^{s,\mathcal{T}}$ is an interim equilibrium of $(u, \mathcal{T})$ for all type spaces $\mathcal{T}$ where each type of every player puts positive probability on his opponents having any payoff type profile. Their proposition 2 shows that s is an ex post equilibrium if and only if it is a weakly information invariant equilibrium. Note that claim (3) in the above proposition is in principle a stronger claim than that s is a weakly information invariant equilibrium.

5 Relations Between Solution Concepts

We want to collect together some results on the relation between the belief free solution concepts.

Write $\psi \in \Delta(\Theta)$ for a distribution over payoff type profiles. Write ψ_μ for the distribution over payoff types generated by μ, i.e.,

$$\psi_\mu(\theta) \triangleq \sum_{a \in A} \mu(a, \theta).$$

For any $\psi \in \Delta(\Theta)$ and payoff type strategy profile s, we write $\mu^{\psi,s}$ for the induced probability distribution over $A \times \Theta$, i.e.,

$$\mu^{\psi,s}(a, \theta) = \begin{cases} \psi(\theta), \text{ if } a = s(\theta), \\ 0, \text{ otherwise.} \end{cases}$$

Lemma 1
If $\psi^* \in \Delta_{++}(\Theta)$ and $a_i \in C_i(\theta_i)$, then there exists an ICE μ such that:

1. $\mu((a_i, a_{-i}), (\theta_i, \theta_{-i})) > 0$ *for some* $(a_{-i}, \theta_{-i}) \in A_{-i} \times \Theta_{-i}$; *and*

[3] It is the restriction to a social choice *function* (rather than a correspondence) that requires the mechanism to depend only on players' reported payoff types. In particular, the step showing (1) $\Rightarrow$ (2) in the above proposition 2 is implied by proposition 1 in Bergemann and Morris (2005) while the step showing (3) $\Rightarrow$ (1) is implied by proposition 3 in Bergemann and Morris (2005).

2. $\psi_\mu = \psi^*$.

Proof. If $a_i \in C_i(\theta_i)$, then by definition there exists an ICE μ' such that $\mu'((a_i, a_{-i}), (\theta_i, \theta_{-i})) > 0$ for some $(a_{-i}, \theta_{-i}) \in A_{-i} \times \Theta_{-i}$. Now let

$$\widetilde{\psi}(\theta) = \frac{1}{1-\varepsilon}\left(\psi^*(\theta) - \varepsilon\psi_{\mu'}(\theta)\right),$$

choosing ε positive but sufficiently small so that $\widetilde{\psi} \in \Delta(\Theta)$. Let $\widetilde{\mu}$ be any ICE with $\psi_{\widetilde{\mu}} = \widetilde{\psi}$. Now let $\mu(a,\theta) = (1-\varepsilon)\widetilde{\mu}(a,\theta) + \varepsilon\mu'(a,\theta)$. By construction, μ has the two properties of the lemma. ■

By definition:

Lemma 2
If s^* is an ex post equilibrium of u, then, for any $\psi \in \Delta(\Theta)$, μ^{ψ,s^*} is an ICE of u.

And an immediate corollary is:

Corollary 1
Suppose that s is an ex post equilibrium of u and, for every $\psi \in \Delta(\Theta)$, there is at most one ICE μ with $\psi_\mu = \psi$. Then μ is an ICE if and only if $\mu = \mu^{\psi,s}$.

We can also record some natural inclusions.

Lemma 3
For all i and $\theta_i \in \Theta_i$,

1. *$C_i(\theta_i) \subseteq R_i(\theta_i)$; and*

2. *if s^* is an ex post equilibrium, then $s_i^*(\theta_i) \in C_i(\theta_i)$.*

Proof. (1) follows immediately from definitions; (2) follows from lemma 2. ■

Proposition 3
Suppose that each C_i is single valued, i.e., there exists payoff strategy profile $s = (s_i)_{i=1}^I$ such that for all i and all θ_i, $C_i(\theta_i) = \{s_i(\theta_i)\}$. Then

1. *if μ is an ICE, then there exists $\psi \in \Delta(\Theta)$ such that $\mu = \mu^{\psi,s}$; and*

2. s is an ex post equilibrium.

Proof. Suppose that $C_i(\theta_i) = \{s_i(\theta_i)\}$ for all i and θ_i. Let μ be any ICE. So $\mu((a_i, a_{-i}), (\theta_i, \theta_{-i})) > 0$ implies $a_i \in C_i(\theta_i) = \{s_i(\theta_i)\}$ for each i and thus $a = s(\theta_{-0})$. Thus

$$\mu^{\psi,s}(a,\theta) = \begin{cases} \psi_\mu(\theta), \text{ if } a = s(\theta_{-0}), \\ 0, \text{ otherwise.} \end{cases}$$

This proves part (1). Recall that for each $\psi \in \Delta(\Theta)$, there exists ICE μ with $\psi_\mu = \psi$. By part (1), $\mu = \mu^{\psi,s}$. So

$$u_i(s(\theta), \theta) \geq u_i\left(\left(a_i, s_{-i}\left(\theta_{-\{0,i\}}\right)\right), \theta\right)$$

for all i and $a_i \in A_i$. But since this argument holds for each θ, we have that s is an ex post equilibrium. ■

6 Supermodular Games

In this section, we let Θ_0 be a singleton and suppress reference to θ_0. Now suppose that each action set and type set is complete ordered; there are increasing differences in actions, so that for each $\theta \in \Theta$,

$$u_i((a_i, a_{-i}), \theta) - u_i\left(\left(a_i', a_{-i}\right), \theta\right) \geq u_i\left(\left(a_i, a_{-i}'\right), \theta\right) - u_i\left(\left(a_i', a_{-i}'\right), \theta\right)$$

if $a_i \geq a_i'$ and $a_{-i} \geq a_{-i}'$; and there are increasing differences in own action and states, so that for each $a_{-i} \in A_{-i}$,

$$u_i((a_i, a_{-i}), \theta) - u_i\left(\left(a_i', a_{-i}\right), \theta\right) \geq u_i\left((a_i, a_{-i}), \theta'\right) - u_i\left(\left(a_i', a_{-i}\right), \theta'\right)$$

if $a_i \geq a_i'$ and $\theta \geq \theta'$.

Say that a game u is *generic* if

$$\underset{a_i}{\arg\max} u_i((a_i, a_{-i}), (\theta_i, \theta_{-i}))$$

is single value for each a_{-i} and θ. Write $\underline{a}_i$ and $\overline{a}_i$ for the smallest and largest actions in A_i and $\underline{\theta}_i$ and $\overline{\theta}_i$ for the smallest and largest types in Θ_i. Iteratively define $\underline{r}_i^k : \Theta_i \to A_i$ and $\overline{r}_i^k : \Theta_i \to A_i$ as follows:

$$\underline{r}_i^0(\theta_i) = \underline{a}_i \text{ and } \overline{r}_i^0(\theta_i) = \overline{a}_i,$$

and for step $k+1$:

$$\underline{r}_i^{k+1}(\theta_i) = \min\left(\arg\max_{a_i} \ u_i\left(\left(a_i, \underline{r}_{-i}^k(\underline{\theta}_{-i})\right), (\theta_i, \underline{\theta}_{-i})\right)\right)$$

and

$$\overline{r}_i^{k+1}(\theta_i) = \max\left(\arg\max_{a_i} \ u_i\left(\left(a_i, \overline{r}_{-i}^k(\overline{\theta}_{-i})\right), (\theta_i, \overline{\theta}_{-i})\right)\right),$$

and the limit points are given by:

$$\underline{r}_i^*(\theta_i) = \lim_{k\to\infty} \underline{r}_i^k(\theta_i) \quad \text{and} \quad \overline{r}_i^*(\theta_i) = \lim_{k\to\infty} \overline{r}_i^k(\theta_i).$$

Lemma 4

For each i and $\theta_i \in \Theta_i$,

$$R_i(\theta_i) = \{a_i \in A_i \,|\underline{r}_i^*(\theta_i) \le a_i \le \overline{r}_i^*(\theta_i)\},$$

where

$$\underline{r}_i^*(\theta_i) = \min\left(\arg\max_{a_i} \ u_i\left(\left(a_i, \underline{r}_{-i}^*(\underline{\theta}_{-i})\right), (\theta_i, \underline{\theta}_{-i})\right)\right),$$

and

$$\overline{r}_i^*(\theta_i) = \min\left(\arg\max_{a_i} \ u_i\left(\left(a_i, \overline{r}_{-i}^*(\overline{\theta}_{-i})\right), (\theta_i, \overline{\theta}_{-i})\right)\right).$$

Proposition 4

In a generic game, each C_i is single valued if and only if each R_i is single valued.

Proof. "If" follows from lemma 3. To prove "only if", we use lemma 4. Now the proof of the proposition is completed as follows. Suppose R_i is not single valued. Then there exists i and θ_i^* with $\overline{r}_i^*(\theta_i^*) > \underline{r}_i^*(\theta_i^*)$. We will construct an ICE where type θ_i^* plays $\overline{r}_i^*(\theta_i^*)$. Let

$$\mu(a,\theta) = \begin{cases} 1-\varepsilon, & \text{if } \theta = \overline{\theta} \text{ and } a = \overline{r}^*\left(\overline{\theta}\right), \\ \varepsilon, & \text{if } \theta = \left(\theta_i^*, \overline{\theta}_{-i}\right) \text{ and } a = \left(\overline{r}_i^*(\theta_i^*), \overline{r}_{-i}^*\left(\overline{\theta}_{-i}\right)\right), \\ 0, & \text{otherwise.} \end{cases}$$

For sufficiently small $\varepsilon > 0$, this will be an ICE by construction: note that type θ_i^* of player i puts probability 1 on $\left(\overline{r}_{-i}^*\left(\overline{\theta}_{-i}\right), \overline{\theta}_{-i}\right)$, so $\overline{r}_i^*(\theta_i^*)$ is a best response while type $\overline{\theta}_i$ of player i puts probability 1 on $\left(\overline{r}_{-i}^*\left(\overline{\theta}_{-i}\right), \overline{\theta}_{-i}\right)$, so $\overline{r}_i^*\left(\overline{\theta}_i\right)$ is a best response; genericity ensures that for each $j \ne i$, $\overline{r}_j^*\left(\overline{\theta}_j\right)$ is unique maximizer of $u_j\left(\left(a_j, \overline{r}_{-j}^*\left(\overline{\theta}_{-j}\right)\right), \left(\theta_j, \overline{\theta}_{-j}\right)\right)$ and therefore it remains a maximizer if he puts probability $1-\varepsilon$ on $\left(\overline{r}_{-j}^*\left(\overline{\theta}_{-j}\right), \overline{\theta}_{-j}\right)$ and probability ε on $\left(\left(\overline{r}_i^*(\theta_i^*), \overline{r}_{-i,j}^*\left(\overline{\theta}_{-j}\right)\right), \left(\theta_i^*, \overline{\theta}_{-i,j}\right)\right)$. ■

In a generic supermodular game, we observe the following additional results.

Corollary 2
If each R_i is single valued, then:

1. $\underline{r}_i^* = \overline{r}_i^* = r^*$;

2. *the best response $r_i^*(\theta_i)$ is given by:*

$$r_i^*(\theta_i) \in \underset{a_i}{\arg\max} \ u_i\left(\left(a_i, r_{-i}^*(\theta_{-i})\right), (\theta_i, \theta_{-i})\right) \text{ for all } \theta_{-i} \in \Theta_{-i};$$

3. *μ is an ICE if and only if $\mu = \mu^{\psi, r^*}$ for some $\psi \in \Delta(\Theta)$.*

Proof. (1) follows from lemma 4. (2) follows from the definitions of $\underline{r}_i^*$ and $\overline{r}_i^*$ and the finiteness of each A_i and Θ_i. For (3), (2) implies that every μ of this form is an ICE and Proposition 4 implies that every ICE must be of this form. ■

7 Potential Games and Unique ICE

We return to the compact continuous case. We say that a game u has *weighted potential* $v : A \times \Theta \to \mathbb{R}$ if there exist $w \in \mathbb{R}_{++}^I$ such that

$$u_i((a_i, a_{-i}), (\theta_i, \theta_{-i})) - u_i\left(\left(a_i', a_{-i}\right), (\theta_i, \theta_{-i})\right) = w_i\left[v((a_i, a_{-i}), (\theta_i, \theta_{-i})) - v\left(\left(a_i', a_{-i}\right), (\theta_i, \theta_{-i})\right)\right]$$

for all i, $a_i, a_i' \in A_i$, $a_{-i} \in A_i$ $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$. This is a belief free incomplete information generalization of the definition of a weighted potential in Monderer and Shapley (1996);[4] in particular, it is equivalent to requiring that each incomplete information game $(u_i(\cdot, \theta))_{i=1}^I$ is a weighted potential game in the sense of Monderer and Shapley (1996), using the same weights for each $\theta \in \Theta$. Game u has *a best response potential* $v : A \times \Theta \to \mathbb{R}$ if for each i, $\theta_i \in \Theta_i$ and $\lambda_i \in \Delta(A_{-i} \times \Theta_{-i})$,

$$\underset{a_i \in A_i}{\arg\max} \int\limits_{a_{-i}, \theta_{-i}} u_i((a_i, a_{-i}), (\theta_i, \theta_{-i}))\, d\lambda_i = \underset{a_i \in A_i}{\arg\max} \int\limits_{a_{-i}, \theta_{-i}} v((a_i, a_{-i}), (\theta_i, \theta_{-i}))\, d\lambda_i.$$

This is a an incomplete information generalization of a best response potential in Morris and Ui (2004). Note in particular that if v is a weighted potential for u, it is also a best response potential for u. We say that v is a strictly concave potential if $v(\cdot, \theta)$ is a strictly concave function of a for all $\theta \in \Theta$.

[4] See Heumen, Peleg, Tjis, and Borm (1996) and Ui (2004) for definitions of Bayesian potentials with prior probability distributions.

Proposition 5 (Uniqueness)

If u has a strictly concave smooth potential function and an ex post equilibrium s, then μ is an incomplete information correlated equilibrium of u if and only if there exists $\psi \in \Delta(\Theta)$ such that $\mu = \mu^{\psi,s}$.

Proof. Neyman (1997) shows that if a complete information game $v(\cdot,\theta)$ (for fixed θ) has a strictly concave potential, then the unique correlated equilibrium is the unique pure strategy profile maximizing $v(\cdot,\theta)$. Thus we have that

$$s(\theta) = \left\{ \arg\max_a v(a,\theta) \right\} \tag{5}$$

for all $\theta \in \Theta$. We adapt the proof by Neyman (1997) to our belief free incomplete information environment. We consider an arbitrary correlated equilibrium given by $\mu \in \Delta(A \times \Theta)$ and show that - if it is not generated by the ex post equilibrium - there exists an improvement for at least one agent i. At any $(a,\theta) \in A \times \Theta$ with $a \neq s(\theta)$, we know that

$$\begin{aligned}\lim_{\varepsilon \to 0^+} \frac{v((1-\varepsilon)a + \varepsilon s(\theta),\theta) - v(a,\theta)}{\varepsilon} &> v(s(\theta),\theta) - v(a,\theta), \text{ by the strict concavity of } v \\ &> 0, \text{ by (5).}\end{aligned}$$

The smoothness of v implies that that

$$\begin{aligned}&\lim_{\varepsilon \to 0^+} \frac{v((1-\varepsilon)a + \varepsilon s(\theta),\theta) - v(a,\theta)}{\varepsilon} \\ = \;&\sum_{i=1}^{I} \lim_{\varepsilon \to 0^+} \frac{v((1-\varepsilon)a + \varepsilon(s_i(\theta_i), a_{-i}),\theta) - v(a,\theta)}{\varepsilon},\end{aligned}$$

and we denote the partial derivative by

$$v_{s_i(\theta_i)-a_i}(a,\theta) \triangleq \lim_{\varepsilon \to 0^+} \frac{v((1-\varepsilon)a + \varepsilon(s_i(\theta_i), a_{-i}),\theta) - v(a,\theta)}{\varepsilon}.$$

We know that for all $a \neq s(\theta)$, we have

$$\sum_{i=1}^{I} v_{s_i(\theta_i)-a_i}(a,\theta) > 0.$$

Now suppose that $\mu \neq \mu^{\psi_\mu,s}$. Then taking expectations over (a,θ) we have

$$\int_{a,\theta} \sum_{i=1}^{I} v_{s_i(\theta_i)-a_i}(a,\theta)\, d\mu > 0$$

and thus

$$\sum_{i=1}^{I} \int\limits_{a,\theta} v_{s_i(\theta_i)-a_i}(a,\theta)\, d\mu > 0.$$

So there exists i such that

$$\int\limits_{a,\theta} v_{s_i(\theta_i)-a_i}(a,\theta)\, d\mu > 0.$$

Thus the deviation $\phi_i(a_i,\theta_i) = \varepsilon s_i(\theta_i) + (1-\varepsilon)a_i$ would strictly increase agent i's ex ante utility for sufficiently small $\varepsilon > 0$. ■

We can apply this result to derive our characterization of $C_i(\theta_i)$ in the quadratic example of Section 3.2.

Corollary 3 (Quadratic Game)
The quadratic common interest game has a smooth concave potential if and only if

$$\gamma \in \left[-\frac{1}{I-1}, 1\right].$$

Proof. Since it was a game of common interests and the payoffs were quadratic, we have the existence and smoothness of the potential function. Now we establish the conditions for concavity. The first derivative is

$$\frac{dv}{da_j}(a,\theta) = -2(a_j-\theta_j) - 2\gamma\left(\sum_{k\neq j}(a_k-\theta_k)\right),$$

and hence the second derivatives are:

$$\frac{1}{2}\frac{dv}{da_i da_j}(a,\theta) = \begin{cases} -1, \text{ if } j = i, \\ -\gamma, \text{ if } j \neq i. \end{cases}$$

Now if M is the (constant) Hessian and $x \in \mathbb{R}^I$, then the quadratic form is:

$$x^T M x = -\gamma\left(\sum_{i=1}^{I} x_i\right)^2 - (1-\gamma)\left(\sum_{i=1}^{I} x_i^2\right)$$

Now if $\gamma < -\frac{1}{I-1}$, and $x_i = z > 0$ for all i, then

$$\begin{aligned} x^T M x &= -\gamma z^2 I^2 - (1-\gamma) z^2 I \\ &= -I z^2 (\gamma I + (1-\gamma)) \\ &= -I z^2 (\gamma(I-1)+1) \\ &= -I(I-1) z^2 \left(\gamma + \frac{1}{I-1}\right) \\ &> 0. \end{aligned}$$

If $\gamma > 1$, and $x_1 = z > 0$, $x_2 = -z$ and $x_3 = \cdots = x_I = 0$, then

$$\begin{aligned} x^T M x &= -\gamma \left(\sum_{i=1}^{I} x_i \right)^2 - (1-\gamma) \left(\sum_{i=1}^{I} x_i^2 \right) \\ &= -2(1-\gamma) z^2 \\ &> 0. \end{aligned}$$

But if $0 \leq \gamma \leq 1$, then

$$\left(\sum_{i=1}^{I} x_i \right)^2 \geq 0 \text{ and } \sum_{i=1}^{I} x_i^2 \geq 0$$

implies that:

$$x^T M x = -\gamma \left(\sum_{i=1}^{I} x_i \right)^2 - (1-\gamma) \left(\sum_{i=1}^{I} x_i^2 \right) \leq 0$$

and if $-\frac{1}{I-1} \leq \gamma \leq 0$, then

$$\left(\sum_{i=1}^{I} x_i \right)^2 \leq I \sum_{i=1}^{I} x_i^2$$

implies that

$$\begin{aligned} x^T M x &= -\gamma \left(\sum_{i=1}^{I} x_i \right)^2 - (1-\gamma) \left(\sum_{i=1}^{I} x_i^2 \right) \\ &\leq -\gamma I \sum_{i=1}^{I} x_i^2 - (1-\gamma) \left(\sum_{i=1}^{I} x_i^2 \right) \\ &= -(I-1) \sum_{i=1}^{I} x_i^2 \left(\gamma + \frac{1}{I-1} \right) \\ &\leq 0 \end{aligned}$$

Thus v is M is negative semi-definite if and only if $-\frac{1}{I-1} \leq \gamma \leq 1$ and is negative definite if $-\frac{1}{I-1} < \gamma < 1$. ■

8 Conclusion

The objective of this paper was to collect and compare belief free solution concepts in games of incomplete information. Among these three concepts under consideration, the notion of ex post equilibrium has received the widest attention in the context of mechanism design. By comparing and relating these solution concepts, it was our objective to emphasize the properties common to these belief free notions in a general game theoretic environment rather than the special setting of mechanism design.

It is important to emphasize that all of these notions do not impose any restrictions on the distributions over the payoff types. A natural question which we hope to address in the future is how the predictions, in particular of the incomplete information correlated equilibrium, would be refined if we were to consider a given prior over payoff types, yet allow for all possible belief type spaces which could be generated by a common prior type space. This intermediate scenario is interesting as the players (and the outside observer) may have learned or otherwise acquired information about the frequency of the payoff types, yet have very little information about the current beliefs and higher order beliefs of the agents.

References

Aumann, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.

Battigalli, P. (1999): "Rationalizability in Incomplete Information Games," Discussion Paper ECO 99/17, European University Institute.

Battigalli, P., and M. Siniscalchi (2003): "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, 3, Article 3.

Bergemann, D., and S. Morris (2001): "Robust Mechanism Design," Discussion Paper http://www.princeton.edu/ smorris/pdfs/robustmechanism2001.pdf, Cowles Foundation, Yale University.

——— (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

——— (2007a): "Robust Implementation: The Case of Direct Mechanisms," Discussion Paper 1561R, Cowles Foundation, Yale University.

——— (2007b): "Strategic Distinguishability with an Application to Robust Virtual Implementation," Discussion Paper 1609, Cowles Foundation, Yale University.

Bergemann, D., and J. Välimäki (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70, 1007–1033.

Borgers, T., and T. McQuade (2007): "Information Invariant Equilibria of Extensive Games," Discussion paper, University of Michigan.

Brandenburger, A., and E. Dekel (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391–1402.

Cotter, K. (1991): "Correlated Equilibrium in Games with Type Dependent Strategies," *Journal of Economic Theory*, 54, 48–68.

——— (1994): "Type Correlated Equilibrium with Payoff Uncertainty," *Economic Theory*, 4, 617–627.

Cremer, J., and R. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," *Econometrica*, 53, 345–361.

Dasgupta, P., and E. Maskin (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341–388.

D'Aspremont, C., and L. Gerard-Varet (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25–45.

Dekel, E., D. Fudenberg, and S. Morris (2006): "Topologies on Types," *Theoretical Economics*, 1, 275–309.

Ely, J. C., and M. Peski (2006): "Hierarchies of Belief and Interim Rationalizability," *Theoretical Economics*, 1.

Forges, F. (1993): "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information," *Theory and Decision*, 35, 277–310.

——— (2006): "Correlated Equilibrium in Games with Incomplete Information Revisited," *Theory and Decision*, 61, 329–344.

Heumen, R. V., B. Peleg, S. Tjis, and P. Borm (1996): "Axiomatic Characterizations of Solutions for Bayesian Games," *Theory and Decision*, 40, 103–130.

Holmstrom, B., and R. Myerson (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.

Jehiel, P., and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

Kalai, E. (2004): "Large Robust Games," *Econometrica*, 72, 1631–1666.

Ledyard, J. (1979): "Dominant Strategy Mechanisms and Incomplete Information," in *Aggregation and Revelation of Preferences*, ed. by J.-J. Laffont, chap. 17, pp. 309–319. North-Holland, Amsterdam.

Monderer, D., and L. Shapley (1996): "Potential Games," *Games and Economic Behavior*, 14, 124–143.

Morris, S., and T. Ui (2004): "Best Response Equivalence," *Games and Economic Behavior*, 14, 260–287.

Neyman, A. (1997): "Correlated Equilibrium and Potential Games," *International Journal of Game Theory*, 26, 223–227.

Ui, T. (2004): "Bayesian Potentials and Information Structures: Team Decision Problems Revisited," Discussion paper, Yokohama National University.