

Interactive epistemology in games with payoff uncertainty

Pierpaolo Battigalli^{a,*}, Marciano Siniscalchi^{b,1}

^a *Università Bocconi, IEP and IGIER, 5 Via Salasco, 20136, Milano, Italy*

^b *Economics Department, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA*

Abstract

We adopt an interactive epistemology perspective to analyse dynamic games with partially unknown payoff functions. We consider solution procedures that iteratively delete strategies conditional on private information about the state of nature. In particular we focus on a weak and a strong version of the Δ -rationalizability solution concept, where Δ represents given restrictions on players' beliefs about state of nature and strategies [Battigalli, P., 2003. Rationalizability in infinite, dynamic games of incomplete information. *Research in Economics* 57, 1–38; Battigalli, P., Siniscalchi, M., 2003. Rationalization and incomplete information. *Advances in Theoretical Economics* 3 (Article 3). <http://www.bepress.com/bejte/advances/vol3/iss1/art3>]. We first show that *weak Δ -rationalizability* is characterized by initial common certainty of rationality and of the restrictions Δ , whereas *strong Δ -rationalizability* is characterized by common strong belief in rationality and the restrictions Δ (cf. [Battigalli, P., Siniscalchi, M., 2002. Strong belief and forward induction reasoning. *Journal of Economic Theory* 106, 356–391]). The latter result allows us to obtain an epistemic characterization of the iterated intuitive criterion. Then we use the framework to analyse the robustness of complete-information rationalizability solution concepts to the introduction of “slight” uncertainty about payoffs. If the set of conceivable payoff functions is sufficiently large, the set of strongly rationalizable strategies with slight payoff uncertainty coincides with the set of complete-information, weakly rationalizable strategies.

© 2007 University of Venice. Published by Elsevier Ltd. All rights reserved.

Keywords: Interactive epistemology; Rationalizability; Forward Induction; Iterated intuitive criterion; Robustness

1. Introduction

Games with payoff uncertainty are situations of strategic interaction where players' payoff functions are not common knowledge. Following Harsanyi's approach (Harsanyi, 1967–68), such situations are usually studied as Bayesian games. A Bayesian game adds to the primitives of the model (players set, rules of interactions, possible payoff functions, players' private information about such functions) a list of “types” for each player, whereby a type determines (implicitly) a whole hierarchy of beliefs about the unknown payoff parameters. In economic applications the analysis is often simplified by assuming that the correspondence from types to hierarchies of beliefs is trivial. For example, it is often assumed that each player's first-order beliefs, i.e. his beliefs about the payoff parameter vector, say θ , are solely determined by his private information about θ according to some mapping, and that all these mappings are common knowledge.

* Corresponding author. Tel.: +39 0258365402.

E-mail addresses: pierpaolo.battigalli@uni-bocconi.it (P. Battigalli), marciano@northwestern.edu (M. Siniscalchi).

¹ Tel.: +1 847 491 5398.




Fig. 1. A signalling game.

The signalling game depicted in Fig. 1 provides an example. Players 1 and 2 move sequentially: Player 1 chooses either u (up), thus terminating the game, or d (down); after d , Player 2 chooses between a , b , or c . There are three possible pairs of payoff functions (mappings from complete sequences of actions to payoff pairs) indexed by a state of nature $\theta \in \{\alpha, \beta, \theta\}$. Player 1 knows the true θ and Player 2 does not. This is represented by drawing three possible game trees with distinguished roots α , β and γ respectively and joining the three possible decision nodes of Player 2 with a dashed line signifying his ignorance of the true game tree. The model described so far is what we call a game with payoff uncertainty. A very simple way to obtain a Bayesian game is to specify an initial belief of the uninformed Player 2 (e.g. $\Pr_2(\alpha) = \frac{1}{2}$, $\Pr_2(\beta) = \frac{1}{3}$, $\Pr_2(\gamma) = \frac{1}{6}$ as in Fig. 1) and to assume that it is common knowledge that (Player 1 knows the true θ and) Player 2 holds this initial belief. But one could consider much more complicated Bayesian games. For example, if Player 1 were unsure about the initial belief about θ of Player 2, one should introduce several possible types of Player 2, corresponding to possible beliefs about θ , and a belief of Player 1 about such types. Were this belief not known to Player 2, one should multiply the possible types of Player 1, and so on.

Equilibria of Bayesian games are very sensitive to the precise specification of higher-order beliefs (see Weinstein and Yildiz, 2007). This is especially problematic because economic modellers find it hard to provide non-arbitrary specifications of the fine details of hierarchical beliefs.

Battigalli (1999, 2003), Battigalli and Siniscalchi (2003) propose a different approach to the analysis of dynamic games with payoff uncertainty: instead of specifying a (more or less complex) type space *à la* Harsanyi, they suggest to take as given some restrictions Δ on players' initial and updated beliefs about θ and their opponents' strategies, and then iteratively delete private information-strategy pairs that are inconsistent with progressively higher levels of mutual certainty of rationality and of the restrictions Δ . For example, in the signalling game above, the modeller may find it reasonable to assume that $\Pr_2(\beta)$ is larger than $\Pr_2(\gamma)$, and there is common certainty of this fact; Δ would then be the set of beliefs profiles such that $\Pr_2(\beta) > \Pr_2(\gamma)$. The resulting solution concept, called Δ -rationalizability, is therefore parametrized by the assumed restrictions Δ .² Battigalli and Siniscalchi (2003) specifically focus on a strong version of Δ -rationalizability, akin to extensive form rationalizability (Pearce, 1984), that also captures a forward induction principle. Indeed, they show that, when Δ reflects agreement of beliefs with a given probability distribution ζ on the terminal nodes of a signalling game, then non-emptiness of the strong Δ -rationalizability solution is equivalent to ζ passing the Iterated Intuitive Criterion of Cho and Kreps (1987). Battigalli (1999, 2003) also considers a weak version of the solution concept that only relies on initial common certainty of rationality and the restrictions Δ . These papers present several examples and economic applications of the approach and report about other applications in the literature.

To illustrate, independently of Player 2's initial beliefs, the signalling game depicted in Fig. 1 has a pooling (sequential) equilibrium where each type of Player 1 chooses u and Player 2 would choose a after d . To see this, note that the best response of Player 2 is the action whose label is the Latin equivalent of the ex-post most likely θ (i.e. a if α is most likely, etc.); since beliefs off the equilibrium path are not determined via Bayes' rule, we may have $\Pr_2(\alpha|d) \geq \max\{\Pr_2(\beta|d), \Pr_2(\theta|d)\}$, making action a a (sequential) best response. This outcome is *a fortiori* weakly Δ -rationalizable for every Δ that allows $\Pr_2(\alpha|d) \geq \max\{\Pr_2(\beta|d), \Pr_2(\theta|d)\}$. However, the pooling-equilibrium outcome is not strongly Δ -rationalizable; this follows from a forward-induction argument. Action d is dominated for

² For this reason it has been referred to as an "umbrella solution concept" by Dekel et al. (2007).

type α , and undominated for types β and γ . If Player 2 believes in the rationality of Player 1 whenever possible, then $\Pr_2(\alpha|d) = 0$ and a is deleted. Therefore the best response of types β and γ is d , which yields two for sure. (The Δ -restrictions may, or may not, pin down the choice between b and c . If Δ implies $\Pr_2(\beta|d) > \Pr_2(\gamma|d)$, then Player 2 best responds with b .)

The above-mentioned papers by Battigalli and Siniscalchi report that Δ -rationalizability is characterized by explicit and transparent assumptions on players' rationality and interactive beliefs. The present paper provides the interactive-epistemology analysis supporting these claims. Following Battigalli and Siniscalchi (1999, 2002), assumptions about rationality and interactive beliefs – including assumptions on how players revise their beliefs when they observe unexpected moves – are represented as events in a universal space of states of the world, whereby each state of the world ω specifies the state of nature θ and how players would behave and think conditional on every partial history of the game (including those that are counterfactual at ω). The weak and strong versions of Δ -rationalizability are shown to be the behavioural consequences of two different sets of assumptions on rationality and interactive beliefs. Given the above-mentioned equivalence result of Battigalli and Siniscalchi (2003), the epistemic characterization of strong Δ -rationalizability also yields an epistemic characterization of the Iterated Intuitive Criterion.

The interactive-epistemology framework is then used to explore the robustness of the complete information versions of these solution concepts to the introduction of “slight” payoff uncertainty. In particular, it is shown that if the set of conceivable payoff functions is sufficiently large, the set of strongly rationalizable strategies with slight payoff uncertainty coincides with the set of complete-information, weakly rationalizable strategies.

This paper is arranged in five sections, including the present one, plus an Appendix A. In Section 2 we introduce the elements of the analysis: dynamic games with payoff uncertainty, systems of conditional beliefs, epistemic models for dynamic games, (sequential) rationality and belief operators. Section 3 is devoted to the epistemic characterization of weak and strong Δ -rationalizability. Section 4 contains the analysis of robustness of the complete information versions of weak and strong rationalizability to the introduction of “slight payoff uncertainty”. Section 5 summarizes the results and discusses the related literature. The Appendix A collects all the proofs.

2. The framework

In this section we present the building blocks of our analysis: dynamic games with payoff uncertainty (Section 2.1), systems of conditional probabilities (Section 2.2), epistemic models (Section 2.3), our basic notion rationality as sequential best response (Section 2.4), and belief operators, i.e. the language to express our assumptions about interactive beliefs (Section 2.5).

2.1. Dynamic games with payoff uncertainty

In order to keep notation at a minimum we restrict our analysis to multistage games with *observable actions* (Fudenberg and Tirole (1991), §3.3; Osborne and Rubinstein (1994), Chap. 6), although most of our results can be extended to general games. We also note that the majority of dynamic games of interest in economics fits within this framework.

We describe the dynamic game under consideration using the following primitive objects: a set $I = \{1, \dots, |I|\}$ of players, a finite collection \mathcal{H} of *partial histories*,³ including the *empty history* ϕ , and a finite collection of *terminal histories* \mathcal{Z} . As the game progresses, the partial history h that has just occurred becomes public information. At some stages there can be simultaneous moves. The (finite) set of actions available to Player i at partial history $h \in \mathcal{H}$ is denoted $A_i(h)$ and, without loss of generality, it is assumed to be non-empty (this introduces some redundancies, but it simplifies the abstract notation). Player i is *active* if $|A_i(h)| \geq 2$ and *inactive* if $|A_i(h)| = 1$. If there is only one active player at each stage, we say that the game has *perfect information*.⁴ To allow for random moves, we also assume that a fictitious player, Player 0, at each partial history h can take an action a_0 in $A_0(h)$, a non-empty finite set (possibly a singleton).

³ Histories are sequences of consecutive action profiles.

⁴ This is not to be confused with “complete information”, which means that all the rules of the game and players' payoff functions (how payoffs depend on terminal histories) are *common knowledge*. Indeed we assume the opposite, since there is payoff uncertainty, there is incomplete information.

In order to model payoff uncertainty, we consider a set

$$\Theta = \Theta_0 \times \Theta_1 \times \cdots \times \Theta_{|I|},$$

of “conceivable” parameter values, where each set Θ_j , $j = 0, \dots, |I|$ is non-empty and finite. For every player $i = 1, \dots, |I|$, each element $\theta_i \in \Theta_i$ represents Player i ’s private information about the unknown payoff-relevant aspects of the game; we call it Player i ’s *information type*. Each element $\theta_0 \in \Theta_0$ (the information type of Player 0) represents any residual uncertainty about payoffs that remains after pooling the players’ private information. We often refer to profile $\theta = (\theta_0, \theta_1, \dots, \theta_{|I|})$ as the *state of nature*.

Observe that, with this formulation, the actions available to players at any given history do *not* depend upon their information-type. This restriction could be removed, at the expense of additional notation; our results do not depend upon it in a crucial way.⁵

Payoffs are associated with terminal nodes, and also depend on the state of nature $\theta \in \Theta$: formally, for each player $i \in I$, we consider a *parametrized payoff function*

$$u_i : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}.$$

For instance, in a pure “private-values” setting, Player i ’s payoff depends solely upon θ_i and the terminal history reached. At the opposite extreme, in a “common-values” environment, players’ payoffs at any terminal history depend only upon θ_0 ; in this case, each player’s information type θ_i is interpreted as a signal, which may be correlated with θ_0 according to i ’s (or someone else’s) subjective beliefs. We include “Player 0” and the residual uncertainty θ_0 in our framework (adding some notational complexity) because it is important in economic applications, and because it helps to relate our work to the recent literature on rationalizability in games with incomplete information. Note that, as long as (real) players believe that the choice of a_0 depends (deterministically or stochastically) on θ_0 , then a_0 can be interpreted as public signal about the residual uncertainty.

We call the structure

$$(I, \mathcal{H}, \Theta_0, A_0(\cdot), (\Theta_i, A_i(\cdot), u_i)_{i \in I})$$

described so far *game with payoff uncertainty*. It describes a situation of strategic interaction where players’ payoff functions are not common knowledge. We prefer to avoid “game with incomplete information” because this expression is mostly used to refer to Bayesian games. The difference between a Bayesian game and a game with payoff uncertainty is that the former specifies (implicitly) players’ hierarchies of beliefs about the state of nature. We shall introduce hierarchies of beliefs of a much richer kind later on in order to obtain a language that allows us to express assumptions about players’ rationality and interactive beliefs, and derive the behavioural implications of these assumptions.

To illustrate our notation with a simple example, consider the signalling game depicted in Fig. 1. Here $I = \{1, 2\}$, \mathcal{H} is isomorphic to $\{\phi, d\}$, \mathcal{Z} is isomorphic to $\{u, (d, a), (d, b), (d, c)\}$, $\Theta_0 = \{\hat{\theta}_0\}$, $\Theta_2 = \{\hat{\theta}_2\}$, $A_0(\phi) = A_0(d) = \{\hat{a}_0\}$, $A_2(\phi) = \{\hat{a}_2\}$ and $A_1(d) = \{\hat{a}_1\}$ are singletons, $\Theta_1 = \{\alpha, \beta, \gamma\}$, $A_1(\phi) = \{u, d\}$ (thus, to be pedantic, the set of partial histories according to our somewhat redundant notation is $\mathcal{H} = \{\phi, (\hat{a}_0, d, \hat{a}_2)\}$, $A_2((\hat{a}_0, d, \hat{a}_2)) = \{a, b, c\}$ (thus $\mathcal{Z} = \{(\hat{a}_0, u, \hat{a}_2), ((\hat{a}_0, d, \hat{a}_2), (\hat{a}_0, \hat{a}_1, a)), ((\hat{a}_0, d, \hat{a}_2), (\hat{a}_0, \hat{a}_1, b)), ((\hat{a}_0, d, \hat{a}_2), (\hat{a}_0, \hat{a}_1, c))\}$); the functions u_i are determined by the payoff pairs at the terminal nodes of the arborescence, (θ, z) .

Beside the basic structure described above, we shall make use of certain derived objects. First, for every $i \in \{0\} \cup I$, we shall denote by S_i the set of *strategies* available to Player i (including the fictitious Player 0). Formally, a strategy is a mapping $s_i : \mathcal{H} \rightarrow \bigcup_{h \in \mathcal{H}} A_i(h)$ such that $\forall h \in \mathcal{H}$, $s_i(h) \in A_i(h)$. Note that we do not model strategies as functions of θ_i . The reason is that we do not assume the existence of an *ex ante* stage at which symmetrically uninformed players plan their type-contingent behaviour. In keeping with standard game-theoretical notation, we let $S = \prod_{i=0}^{|I|} S_i$ and $S_{-i} = \prod_{j \neq i} S_j$; to clarify, here and in the following, both “ $-i$ ” and “ $j \neq i$ ” denote the index set $(\{0\} \cup I) \setminus \{i\}$. We then define

$$\Sigma_i = \Theta_i \times S_i, \quad \Sigma = \prod_{j=0}^{|I|} \Sigma_j, \quad \Sigma_{-i} = \prod_{j \neq i} \Sigma_j;$$

⁵ Battigalli (2003) allows for type-dependent actions sets.

every $\sigma_i \in \Sigma_i$ is an information type-strategy pair for Player i .

Second, for any history h , $S(h)$ denotes the set of strategy profiles which induce h ; its projections on S_i and S_{-i} are denoted by $S_i(h)$ and $S_{-i}(h)$, respectively. The correspondence $S(\cdot)$ provides a convenient strategic-form representation of the (non-payoff) information structure. To abbreviate notation we also write $\Sigma_i(h) = \Theta_i \times S_i(h)$, $\Sigma(h) = \prod_{j=0}^{|I|} \Sigma_j(h)$, $\Sigma_{-i}(h) = \prod_{j \neq i} \Sigma_j(h)$.

Using this notation, we can define a strategic-form payoff function $U_i : \Sigma_i \times \Sigma_{-i} \rightarrow \mathbb{R}$ in the usual way: let $\mathbf{z}(s)$ denote the terminal history induced by the strategy profile $s \in S$; for all $\sigma_i = (\theta_i, s_i) \in \Sigma_i$ and $\sigma_{-i} = (\theta_{-i}, s_{-i}) \in \Sigma_{-i}$,

$$U_i(\sigma_i, \sigma_{-i}) = u_i(\theta, \mathbf{z}(s)),$$

with $\theta = (\theta_0, \theta_1, \dots, \theta_{|I|})$, $s = (s_0, s_1, \dots, s_{|I|})$.

It is convenient to introduce two additional pieces of notation. For every strategy s_i , $\mathcal{H}(s_i) = \{h \in \mathcal{H} : s_i \in S_i(h)\}$ denotes the collection of partial histories consistent with s_i . For every partial history h and strategy s_i , s_i^h denotes the strategy consistent with h which coincides with s_i on the set of partial histories not preceding h (thus, $h \in \mathcal{H}(s_i)$ implies $s_i^h = s_i$).⁶

2.2. Conditional probability systems

As the game progresses, players update and/or revise their beliefs in light of newly acquired information. In order to account for this process, we represent beliefs by means of conditional probability systems (Myerson, 1986; R enyi, 1955).

Fix a player $i \in I$. For a given measure space (X_i, \mathcal{A}_i) , consider a non-empty, finite or countable collection $\mathcal{B}_i \subset \mathcal{A}_i$ of events such that $\emptyset \notin \mathcal{B}_i$. The interpretation is that Player i is uncertain about the “true” element $x \in X_i$, and \mathcal{B}_i is a collection of observable events – or “relevant hypotheses” – concerning a “discrete” component of x .

Definition 1. A conditional probability system (or CPS) on $(X_i, \mathcal{A}_i, \mathcal{B}_i)$ is⁷ a mapping

$$\mu(\cdot|\cdot) : \mathcal{A}_i \times \mathcal{B}_i \rightarrow [0, 1]$$

satisfying the following axioms:

Axiom 1. For all $B \in \mathcal{B}_i$, $\mu(B|B) = 1$.

Axiom 2. For all $B \in \mathcal{B}_i$, $\mu(\cdot|B)$ is a probability measure on (X_i, \mathcal{A}_i) .

Axiom 3. For all $A \in \mathcal{A}_i$, $B, C \in \mathcal{B}_i$, if $A \subset B \subset C$ then $\mu(A|B)\mu(B|C) = \mu(A|C)$.

The set of probability measures on (X_i, \mathcal{A}_i) will be denoted by $\Delta(X_i)$; we shall endow it with the topology of weak convergence of measures. The set of conditional probability systems on $(X_i, \mathcal{A}_i, \mathcal{B}_i)$ can be regarded as a subset of $[\Delta(X_i)]^{\mathcal{B}_i}$, endowed with the product topology.

Throughout this paper, we shall be interested solely in “relevant hypotheses” corresponding to the event that a certain partial history has occurred. Thus, Player i ’s first-order beliefs about the state of nature and his opponents’ behavior may be represented by taking $X_i = \Sigma_{-i} = \prod_{j \neq i} \Theta_j \times S_j$ and $\mathcal{B}_i = \{B \subset \Sigma_{-i} : B = \Sigma_{-i}(h) \text{ for some } h \in \mathcal{H}\}$. We denote the collection of CPSs on $(\Sigma_{-i}, \mathcal{B}_i)$ thus defined by $\Delta^{\mathcal{H}}(\Sigma_{-i})$. Since Σ_{-i} and \mathcal{H} are finite, $\Delta^{\mathcal{H}}(\Sigma_{-i})$ is easily seen to be a closed subset of Euclidean $|\mathcal{H}| \cdot |\Sigma_{-i}|$ -space.

To represent Player i ’s higher-order beliefs, we will consider a (finite or infinite) set of “possible worlds” $\Omega = \prod_{i=0}^{|I|} \Omega_i$, where $\Omega_0 = \Sigma_0$ and, for every player $i \in I$, $\Omega_i \subset \Sigma_i \times Y_i$ and $\text{proj}_{S_i} \Omega_i = S_i$. Elements of the sets Y_i will be interpreted as epistemic types of i . As will be clear momentarily, it is convenient to assume that

⁶ Recalling that we assume (without loss of generality) that each player takes an action immediately after every partial history, s_i^h is defined as follows: For all $h' \in \mathcal{H}$, if either $(h', (a_j)_{j \in I})$ comes before h or $(h', (a_j)_{j \in I}) = h$, then $s_i^h(h') = a_i$. Otherwise, $s_i^h(h') = s_i(h')$.

⁷ The tuple $(X_i, \mathcal{A}_i, \mathcal{B}_i, \mu)$ is called conditional probability space by R enyi (1955). When X_i is finite, $\mathcal{A}_i = 2^{X_i}$, $\mathcal{B}_i = 2^{X_i} \setminus \{\emptyset\}$, we obtain Myerson’s conditional probability systems (Myerson, 1986).

each Y_i is a Polish (i.e. separable and completely metrizable) space. Consistently with our previous notation, we let $\Omega_{-i} = \Omega_0 \times \prod_{j \in I \setminus \{i\}} \Omega_j$.

To represent Player i 's hierarchical beliefs about his opponents, we use the following structure: let $X_i = \Omega_{-i}$, let \mathcal{A}_i be the Borel sigma algebra on Ω_{-i} and

$$\mathcal{B}_i = \{B \in \mathcal{A}_i : B = \{((\theta_0, s_0), (\theta_j, s_j, y_j)_{j \in I \setminus \{i\}}) \in \Omega_{-i} : s_{-i} \in S_{-i}(h)\} \text{ for some } h \in \mathcal{H}\}.$$

The set of CPSs on $(\Omega_{-i}, \mathcal{B}_i)$ will be denoted by $\Delta^{\mathcal{H}}(\Omega_{-i})$. Similarly, to represent Player i 's hierarchical beliefs about the prevailing state of the world (including his own strategy and beliefs, as well as his opponents'), let $X_i = \Omega$, let \mathcal{A} be the Borel sigma algebra on Ω and

$$\mathcal{B} = \{B \in \mathcal{A} : B = \{((\theta_0, s_0), (\theta_j, s_j, y_j)_{j \in I}) \in \Omega : s \in S(h)\} \text{ for some } h \in \mathcal{H}\}.$$

The set of CPSs on (Ω, \mathcal{B}) is denoted $\Delta^{\mathcal{H}}(\Omega)$.

Note that Ω_{-i} and Ω are Polish spaces in the respective product topologies; also, the finite collections \mathcal{B}_i and \mathcal{B} consist of sets that are both *open and closed* in the respective topologies. Battigalli and Siniscalchi (1999) show that, under these conditions, $\Delta^{\mathcal{H}}(\Omega_{-i})$ and $\Delta^{\mathcal{H}}(\Omega)$ are closed subsets of the Polish spaces $[\Delta(\Omega_{-i})]^{\mathcal{B}_i}$ and, respectively, $[\Delta(\Omega)]^{\mathcal{B}}$. Hence, they are Polish spaces in the relative topology.

2.3. Epistemic models

We next introduce our basic representation of hierarchical conditional beliefs (Ben-Porath, 1997; Battigalli and Siniscalchi, 1999, 2002).

Definition 2. A *type space* on $(\mathcal{H}, \Theta, S(\cdot), I)$ is a tuple $\mathcal{T} = (\mathcal{H}, S(\cdot), I, \Omega_0, (\Omega_i, T_i, g_i)_{i \in I})$ such that, for every $i \in I$, T_i is a Polish space and

1. $\Omega_0 = \Sigma_0$;
2. Ω_i is a closed subset of $\Sigma_i \times T_i$ such that $\text{proj}_{\Sigma_i} \Omega_i = \Sigma_i$;
3. $g_i : T_i \rightarrow \Delta^{\mathcal{H}}(\Omega_{-i})$ is a continuous mapping.

For any $i \in I$, the elements of the set T_i are referred to as Player i 's *epistemic types*. A type space is *compact* if all the sets T_i , $i \in I$, are compact topological spaces.

Thus, at any “possible world” $\omega = ((\theta_0, s_0), (\theta_i, s_i, t_i)_{i \in I}) \in \Omega$, we specify a state of nature $(\theta_0, \theta_1, \dots, \theta_{|I|})$, as well as each player i 's *dispositions to act* (his strategy s_i) and (for $i \in I$) his *dispositions to believe* (his system of conditional probabilities $g_i(t_i) = (g_{i,h}(t_i))_{h \in \mathcal{H}}$). These dispositions also include what a player *would* do and think at histories that are inconsistent with ω .⁸

As is traditional in the epistemic analysis of games, we complete a player's system of conditional beliefs by assuming that he is certain of his true strategy, information type and epistemic type. More specifically, we assume that for every state of the world $((\theta_i, s_i, t_i), \omega_{-i})$ and every history h , Player $i \in I$ would be certain of t_i and θ_i given h , and would also be certain of s_i given h provided that s_i is consistent with h , i.e. $s_i \in S_i(h)$. We also assume that if $s_i \notin S_i(h)$ Player i would still be certain that his continuation strategy agrees with s_i . (The latter assumption is immaterial for our analysis, but we include it for completeness).

Formally, Player i 's conditional beliefs on (Ω, \mathcal{B}) are given by a continuous mapping

$$g_i^* = (g_{i,h}^*)_{h \in \mathcal{H}} : \Omega_i \rightarrow \Delta^{\mathcal{H}}(\Omega)$$

derived from g_i by the following formula: for all $(\theta_i, s_i, t_i) \in \Omega_i$, $h \in \mathcal{H}$, $E \in \mathcal{A}$,

$$g_{i,h}^*(\theta_i, s_i, t_i)(E) = g_{i,h}(t_i) \left(\left\{ \omega_{-i} \in \Omega_{-i} : ((\theta_i, s_i^h, t_i), \omega_{-i}) \in E \right\} \right). \quad (1)$$

Type spaces encode a collection of infinite hierarchies of CPSs for each epistemic type of each player. It is natural to ask whether there exists a type space which encodes *all* “conceivable” hierarchical beliefs. Mertens and Zamir (1985)

⁸ History h is inconsistent with (or counterfactual at) $\omega = (\theta, s, t)$ if $s \notin S(h)$.

and **Brandenburger and Dekel (1993)** answered this question in the affirmative when beliefs are represented by probability measures on a compact Hausdorff or Polish space; **Battigalli and Siniscalchi (1999)** provide a counterpart of these results in the present “dynamic” setting where beliefs are represented by CPSs.

Consider the following definition:

Definition 3. A belief-complete type space on $(\mathcal{H}, \Theta, S(\cdot), I)$ is a type space $\mathcal{T} = (\mathcal{H}, S(\cdot), I, \Omega_0, (\Omega_i, T_i, g_i)_{i \in I})$ such that, for every $i \in I$, $\Omega_i = \Sigma_i \times T_i$ and $g_i : T_i \rightarrow \Delta^{\mathcal{H}}(\Omega_{-i})$ is onto.⁹

It is shown in **Battigalli and Siniscalchi (1999)** that a belief-complete type space may always be constructed (for all finite games, and also a large class of infinite games) by taking the sets of epistemic types to be the collection of all possible hierarchies of conditional probability systems that satisfy certain intuitive coherency conditions.¹⁰ Also, every type space may be viewed as a belief-closed subspace of the space of infinite hierarchies of conditional beliefs.¹¹ Finally, since we assume that the set of external states Σ is finite and hence compact, the belief-complete type space thus constructed is also compact.

2.4. Sequential rationality

We take the view that a strategy $s_i \in S_i$ for Player i should be optimal, given Player i 's beliefs and payoff type, conditional upon any history consistent with s_i ; we do not impose restrictions on the action specified as histories that cannot obtain if Player i follows the strategy s_i . This is a sequential best-response property which applies to plans of actions¹² as well as strategies (see, for example, **Rubinstein (1991)** and **Reny (1992)**).

Definition 4. Fix a player $i \in I$, a CPS $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ and an information type $\theta_i \in \Theta_i$. A strategy $s_i \in S_i$ is a sequential best reply to μ_i for θ_i if and only if, for every $h \in \mathcal{H}(s_i)$ and every $s'_i \in S_i(h)$,

$$\sum_{\sigma_{-i}} [U_i(\theta_i, s_i, \sigma_{-i}) - U_i(\theta_i, s'_i, \sigma_{-i})] \mu_i(\{\sigma_{-i}\} | \Sigma_{-i}(h)) \geq 0.$$

For any CPS $\mu_i \in \Delta^{\mathcal{H}}(S_{-i})$, let $r_i(\mu_i)$ denote the set of pairs (s_i, θ_i) such that s_i is a sequential best reply to μ_i for θ_i .

It can be shown by standard arguments that r_i is a non-empty valued and upper-hemicontinuous correspondence. It is convenient to introduce the following additional notation: Fix a type space $(\mathcal{H}, S(\cdot), I, \Omega_0, (\Omega_i, T_i, g_i)_{i \in I})$. For every player $i \in I$, let $f_i = (f_{i,h})_{h \in \mathcal{H}} : T_i \rightarrow [\Delta(\Sigma_{-i})]^{\mathcal{H}}$ denote his first-order belief mapping, that is, for all $t_i \in T_i$ and $h \in \mathcal{H}$,

$$f_{i,h}(t_i) = \text{marg}_{\Sigma_{-i}} g_{i,h}(t_i)$$

(recall that $\text{proj}_{\Sigma_{-i}} \Omega_{-i} = \Sigma_{-i}$). It is easy to see that $f_i(t_i) \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ for every $t_i \in T_i$; also, f_i is continuous.

Finally, we can introduce our key behavioural axiom. We say that Player i is rational at a state $\omega = (\theta, s, t)$ in \mathcal{T} if and only if $(\theta_i, s_i) \in r_i(f_i(t_i))$. Then the event

$$R_i = \{\omega = (\theta, s, t) \in \Omega : (\theta_i, s_i) \in r_i(f_i(t_i))\}$$

corresponds to the statement “Player i is rational”. (Note that R_i is closed because the correspondence $r_i \circ f_i$ is upper hemicontinuous.) We shall also refer to the events $R = \bigcap_{i \in I} R_i$ (“every player is rational”) and $R_{-i} = \bigcap_{j \in I \setminus \{i\}} R_j$ (“every opponent of Player i is rational”).

A word of caution. Events are defined with reference to a specific type space, but this dependence is immaterial in the following analysis.

⁹ We use “complete” in the same sense as **Brandenburger (2003)**, who shows (in a different framework) that a (belief-) complete, filter-theoretical type of space does not exist. Of course, this notion of completeness is not to be confused with the topological one.

¹⁰ **Battigalli and Siniscalchi (1999)** uses a slightly different definition of type space. But all the arguments in **Battigalli and Siniscalchi (1999)** can be easily adapted to the present framework.

¹¹ The representation of a type space as a belief-closed subspace of hierarchies “coalesces” duplicate types, i.e. types that yield the same hierarchy of CPSs. Duplicate types do not play any role in our analysis.

¹² Intuitively, a plan of action for player i is silent about which actions would be taken by i if i did not follow that plan. Formally, a plan of action is a class of realization-equivalent strategies. In generic extensive games, a plan of action is a strategy of the reduced normal form.

2.5. Belief operators

The next building block is the epistemic notion of (conditional) *probability one belief*, or (conditional) *certainty*. Recall that an epistemic type encodes the beliefs a player would hold, should any one of the possible non-terminal histories occur. This allows us to formalize statements such as, “Player i would be certain that Player j is rational, were he to observe history h ”.

Given a type space $\mathcal{T} = (\mathcal{H}, S(\cdot), I, \Omega_0, (\Omega_i, T_i, g_i)_{i \in I})$, for every $i \in I$, $h \in \mathcal{H}$, and $E \in \mathcal{A}$, define the event¹³

$$B_{i,h}(E) = \{(\sigma, t) \in \Omega : g_{i,h}^*(\sigma_i, t_i)(E) = 1\}$$

which corresponds to the statement “Player i would be certain of E , were he to observe history h ”. Observe that this definition incorporates the natural requirement that a player only be certain of events which are consistent with her own (continuation) strategy and epistemic type (recall how g_i^* was derived from g_i). Recalling that ϕ is the empty history, $B_{i,\phi}(E)$ is the event “Player i believes E at the beginning of the game”.

For each player $i \in I$ and history $h \in \mathcal{H}$, the definition identifies a set-to-set operator $B_{i,h} : \mathcal{A} \rightarrow \mathcal{A}$ which satisfies the usual properties of falsifiable beliefs (see, for example, Chapter 3 of Fagin et al. (1995)); in particular, it satisfies

- *Conjunction*: For all events $E, F \in \mathcal{A}$, $B_{i,h}(E \cap F) = B_{i,h}(E) \cap B_{i,h}(F)$;
- *Monotonicity*: For all events $E, F \in \mathcal{A}$: $E \subset F$ implies $B_{i,h}(E) \subset B_{i,h}(F)$.

We shall say that Player i *strongly believes* that an event $E \neq \emptyset$ is true (i.e. adopts E as a “working hypothesis”) if and only if he is certain of E at all histories consistent with E . Formally, for any type space $\mathcal{T} = (\mathcal{H}, S(\cdot), I, \Omega_0, (\Omega_i, T_i, g_i)_{i \in I})$, for each $i \in I$ define the operator $SB_i : \mathcal{A} \rightarrow \mathcal{A}$ by $SB_i(\emptyset) = \emptyset$ and

$$SB_i(E) = \bigcap_{h \in \mathcal{H}: E \cap (\Sigma(h) \times T) \neq \emptyset} B_{i,h}(E)$$

for all events $E \in \mathcal{A} \setminus \{\emptyset\}$. Note that $SB_i(E) \subset B_{i,\phi}(E)$ for all $E \in \mathcal{A}$; that is, strong belief implies initial certainty (since ϕ is the empty history, which represents the beginning of the game, $B_{i,\phi}$ is the initial certainty operator for Player i). Also note that in non-trivial games SB_i fails Conjunction and Monotonicity (for more on this see Battigalli and Siniscalchi (2002)).

The corresponding mutual belief operators are defined in the obvious way:

$$B_h(E) = \bigcap_{i \in I} B_{i,h}(E),$$

$$SB(E) = \bigcap_{i \in I} SB_i(E).$$

3. Epistemic characterization of Δ -rationalizability

In this section we introduce restrictions Δ on first-order beliefs (Section 3.1). These restrictions are taken as parametrically given in the definition of two solution procedures, weak Δ -rationalizability (Section 3.2) and strong Δ -rationalizability (Section 3.3). Finally we show how these solution concepts are implied by assumptions about players rationality and interactive beliefs, including assumptions on how players revise their beliefs (Section 3.4).

3.1. Restrictions on beliefs

A player’s beliefs may be assumed to satisfy some restrictions that are not implied by mutual, common, or strong belief in rationality. We may distinguish between (i) *restrictions on exogenous* (first-order) *beliefs*, that is, beliefs about the state of nature and chance moves, and (ii) *restrictions on endogenous* (first-order) *beliefs*, that is, beliefs about behaviour; our general theory considers both (i) and (ii). An example of restrictions on exogenous beliefs is that chance moves have given “objective” probabilities. Another example is that players’ beliefs about the state of

¹³ For any measurable subset $E \subset \Omega$, $B_{i,h}(E)$ is closed, hence measurable; this follows from the continuity of $g_{i,h}^*$, via an application of the portmanteau theorem.

nature conform to a given prior on Θ . An example of restrictions on endogenous beliefs is that players’ beliefs satisfy certain independence or conditional independence properties (see Dekel et al. (2007) and Battigalli et al. (2007)). Furthermore, we allow for the possibility that such restrictions may depend upon each player’s information type. Other examples of restrictions on beliefs, as well as applications and additional discussion, can be found in Battigalli and Siniscalchi (2003) and Battigalli (2003).

Formally, for every player $i \in I$, every information type $\theta_i \in \Theta_i$, and every history $h \in \mathcal{H}$, consider a set $\Delta_{\theta_i,h} \subseteq \Delta(\Theta_{-i} \times S_{-i}(h))$; say that a collection Δ_{θ_i} of CPSs on Σ_{-i} is *regular* if there are non-empty and closed subsets $\Delta_{\theta_i,h}$ ($h \in \mathcal{H}$) such that

$$\Delta_{\theta_i} = \Delta^{\mathcal{H}}(\Sigma_{-i}) \cap \prod_{h \in \mathcal{H}} \Delta_{\theta_i,h}.$$

To simplify notation, we denote the whole set of restrictions by Δ : $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$. That is: Player i ’s first-order conditional beliefs are required to be represented by a CPS, but also to satisfy the history – and information-type – dependent restrictions encoded in the sets $\Delta_{\theta_i,h}$. Observe that $\Delta^{\mathcal{H}}(\Sigma_{-i})$ is a subset of $\prod_{h \in \mathcal{H}} [\Delta(\Theta_{-i} \times S_{-i}(h))]$, so the above intersection makes sense.

As we made clear, the restrictions Δ only concern first-order beliefs, but in our epistemic analysis we will assume that there is initial common certainty that the restrictions Δ hold, thus yielding restrictions on hierarchies of beliefs.

We point out that our approach is sufficiently flexible to encompass the restrictions on infinite hierarchies of exogenous beliefs implicit in Bayesian games. To clarify this point, recall that information types θ_i need not be payoff relevant, they may just represent information that is perceived to be correlated with the payoff-relevant state. More generally, θ_i might encode both payoff relevant and payoff-irrelevant information. But, to simplify our discussion of the relationship with Bayesian games, we consider here the extreme “common value” setting whereby only θ_0 is payoff relevant. According to Harsanyi’s approach (Harsanyi, 1967–68) the game model must specify belief mappings $\pi_i : \Theta_i \rightarrow \Delta(\Theta_{-i})$, $i \in I$ (which may, but need not be derived from a common prior $\Pi \in \Delta(\Theta)$). Appending the belief mappings $(\pi_i)_{i \in I}$ to the given game with payoff uncertainty we obtain a Bayesian game. Taking the marginal of each $\pi_i(\theta_i)$ on Θ_0 we obtain the first-order beliefs on the payoff relevant parameter θ_0 . Next, using the mappings $\theta_j \mapsto \text{marg}_{\Theta_0} \pi_j(\theta_j)$ ($j \neq 0, i$) it is possible to derive the second-order beliefs of each type θ_i over $\Theta_0 \times [\Delta(\Theta_0)]^{I \setminus \{i\}}$. Proceeding in a similar fashion it is possible to obtain a Θ_0 -based infinite hierarchy of beliefs for each type θ_i of each player i . This is the representation of hierarchical beliefs about the payoff parameter implicit in Bayesian games.

Now, for given belief mappings $(\pi_i)_{i \in I}$ we can consider the following restrictions:

$$\Delta_{\theta_i} = \{\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i}) : \text{marg}_{\Theta_0} \mu_i(\cdot | \Sigma_{-i}) = \pi_i(\theta_i)\}.$$

Assuming initial common certainty of the restrictions $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$, we obtain the same Θ_0 -based hierarchies of initial belief as in the Bayesian game determined by $(\pi_i)_{i \in I}$.¹⁴

3.2. Weak Δ -rationalizability

The first solution concept we consider is a generalization of a notion of rationalizability for game with complete and perfect information due to Ben-Porath (1997).¹⁵

Definition 5. Consider the following procedure.

- (Step 0) For every $i \in \{0\} \cup I$, let $\mathbf{W}_{i,\Delta}^0 = \Sigma_i$. Also, let $\mathbf{W}_{-i,\Delta}^0 = \prod_{j \neq i} \mathbf{W}_{i,\Delta}^0$ and $\mathbf{W}_{\Delta}^0 = \prod_{i \in \{0\} \cup I} \mathbf{W}_{i,\Delta}^0$.
- (Step $n > 0$) Let $\mathbf{W}_{0,\Delta}^n = \Sigma_0$. Then, for every $i \in I$, and for every $\sigma_i = (\theta_i, s_i) \in \Sigma_i$, let $\sigma_i \in \mathbf{W}_{i,\Delta}^n$ if and only if $\sigma_i \in \mathbf{W}_{i,\Delta}^{n-1}$ and there exists a CPS $\mu \in \Delta_{\theta_i}$ such that

¹⁴ For more on this in the context of static games see Battigalli et al. (2007). Later on we will comment on the possibility to interpret Δ as “common knowledge restrictions” on first-order beliefs.

¹⁵ Recall that in our framework a game has *complete information* if Θ is a singleton, and has *perfect information* if there is only one active player at every non-terminal history. Thus, according to our definition these two properties are logically independent.

- (a) $\sigma_i \in r_i(\mu)$;
 (b) $\mu(\mathbf{W}_{-i,\Delta}^{n-1} | \Sigma_{-i}) = 1$.
 Also let $\mathbf{W}_{-i,\Delta}^n = \prod_{j \neq i} \mathbf{W}_{j,\Delta}^n$ and $\mathbf{W}_{\Delta}^n = \prod_{i \in \{0\} \cup I} \mathbf{W}_{i,\Delta}^n$.

Note that the above definition is silent about how the players would change their beliefs if they observed a history h which they previously believed impossible, even if h is consistent with rationality (and Δ) or mutual certainty of rationality (and Δ) of any order. Therefore weak rationalizability satisfies only a very weak form of backward induction (e.g. in two-stage games with complete and perfect information) and can not capture any kind of forward induction reasoning. This is what makes weak rationalizability different from strong rationalizability, the concept we consider next.

3.3. Strong Δ -rationalizability

The following solution concept is a generalization of a notion of extensive form rationalizability put forward by Pearce (1984) and further analysed by Battigalli (1996, 1997).

Definition 6. Consider the following procedure.

- (Step 0) For every $i \in \{0\} \cup I$, let $\mathbf{S}_{i,\Delta}^0 = \Sigma_i$. Also, let $\mathbf{S}_{-i,\Delta}^0 = \prod_{j \neq i} \mathbf{S}_{j,\Delta}^0$ and $\mathbf{S}_{\Delta}^0 = \prod_{i \in \{0\} \cup I} \mathbf{S}_{i,\Delta}^0$.
 (Step $n > 0$) Let $\mathbf{S}_{0,\Delta}^n = \Sigma_0$. Then, for every $i \in I$, and for every $\sigma_i = (\theta_i, s_i) \in \Sigma_i$, let $\sigma_i \in \mathbf{S}_{i,\Delta}^n$ if and only if $\sigma_i \in \mathbf{S}_{i,\Delta}^{n-1}$ and there exists a CPS $\mu \in \Delta_{\theta_i}$ such that
 (a) $\sigma_i \in r_i(\mu)$;
 (b) for all $h \in \mathcal{H}$, if $\mathbf{S}_{-i,\Delta}^{n-1} \cap \Sigma_{-i}(h) \neq \emptyset$, then $\mu(\mathbf{S}_{-i,\Delta}^{n-1} | \Sigma_{-i}(h)) = 1$.
 Also let $\mathbf{S}_{-i,\Delta}^n = \prod_{j \neq i} \mathbf{S}_{j,\Delta}^n$ and $\mathbf{S}_{\Delta}^n = \prod_{i \in \{0\} \cup I} \mathbf{S}_{i,\Delta}^n$.

Finally, let $\mathbf{S}_{\Delta}^{\infty} = \bigcap_{n \geq 0} \mathbf{S}_{\Delta}^n$.

3.4. Characterization results

We adopt a uniform notation for the n -fold composition of operators. Formally, let \mathcal{A} be the sigma-algebra of events in a state space Ω , and fix a map $\mathcal{O} : \mathcal{A} \rightarrow \mathcal{A}$; then, for any event $E \in \mathcal{A}$, let $\mathcal{O}^0(E) = E$ and, for $n \geq 1$, let $\mathcal{O}^n(E) = \mathcal{O}(\mathcal{O}^{n-1}(E))$.

Denote by $[\Delta]$ the event that the players' first-order beliefs satisfy the regular restrictions given by $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$; that is,

$$[\Delta_{\theta_i}] = \{(\theta_i, s_i, t_i, \omega_{-i}) \in \Omega : f_i(t_i) \in \Delta_{\theta_i}\}, \quad [\Delta] = \bigcap_{i \in I, \theta_i \in \Theta_i} [\Delta_{\theta_i}].$$

Note that these are a well-defined events because the sets Δ_{θ_i} (by regularity) are closed and function f_i is continuous. Recall that \mathbf{B}_{ϕ} is the mutual initial belief operator.

Proposition 1. Fix a collection $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of regular subsets of CPSs. Then, for any belief-complete type space,

- (i) for every $n \geq 0$, $\mathbf{W}_{\Delta}^{n+1} = \text{proj}_{\Sigma} \bigcap_{k=0}^n \mathbf{B}_{\phi}^k(R \cap [\Delta])$;
 (ii) if the type space is also compact, then $\mathbf{W}_{\Delta}^{\infty} = \text{proj}_{\Sigma} \bigcap_{k \geq 0} \mathbf{B}_{\phi}^k(R \cap [\Delta])$.

Proposition 1 says that weak Δ -rationalizability is characterized by the following assumptions:

- (a) Players are rational and their first-order beliefs satisfy restrictions Δ ,
 (b) there is initial common certainty of (a).

To characterize strong Δ -rationalizability, consider the auxiliary “correct mutual strong belief” operator $\text{CSB} : \mathcal{A} \rightarrow \mathcal{A}$ defined by

$$\text{CSB}(E) = E \cap \text{SB}(E)$$

for any $E \in \mathcal{A}$ (recall that SB is the mutual strong belief operator). Note that by definition $\text{CSB}^n(E) = \text{CSB}^{n-1}(E) \cap \text{SB}(\text{CSB}^{n-1}(E)) \subset \text{CSB}^{n-1}(E)$. Therefore we obtain a decreasing sequence of events, and it makes sense to define $\text{CSB}^{\infty}(E) = \bigcap_{n \geq 0} \text{CSB}^n(E)$.

Proposition 2. Fix a collection $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of regular subsets of CPSs. Then, for any belief-complete type space,

- (i) for every $n \geq 0$, $S_{\Delta}^{n+1} = \text{proj}_{\Sigma} \text{CSB}^n(R \cap [\Delta])$;
- (ii) if the type space is also compact, then $S_{\Delta}^{\infty} = \text{proj}_{\Sigma} \text{CSB}^{\infty}(R \cap [\Delta])$.

Proposition 2 says that strong Δ -rationalizability is characterized by the following assumptions:

- (a) players are rational and their first-order beliefs satisfy restrictions Δ ,
 - (b.1) players strongly believe in (a)
 - (b.2) players strongly believe in (a) and (b.1),
 - ...
 - (b.k + 1) players strongly believe in (a), (b.1), ..., (b.k),

etc.

This is similar to the best rationalization principle described by Battigalli (1996) and analysed epistemically by Battigalli and Siniscalchi (2002), but here the rationalization is consistent with the restrictions Δ . Thus, Player i 's own first-order beliefs, when his information type is θ_i , are an element of Δ_{θ_i} ; he adopts the working assumption that her opponents are rational and their beliefs are elements of $\prod_{j \in I \setminus \{i\}, \theta_j \in \Theta_j} \Delta_{\theta_j}$; and so on.

A subtle issue pertaining to interactive beliefs about the restrictions Δ deserves further discussion. The epistemic assumptions of Propositions 1 and 2 only imply that there is initial common certainty of event $[\Delta]$, but in a dynamic game this is different from assuming “common knowledge” of $[\Delta]$. The latter would imply common certainty of $[\Delta]$ conditional on every history; on the other hand, according to our epistemic assumptions, common certainty of $[\Delta]$ may well fail at some history $h \neq \phi$.

However, we conjecture that the following versions of our epistemic characterization results hold: let $B_{\mathcal{H}}$ be the mutual belief operator defined by the formula $B_{\mathcal{H}}(E) = \bigcap_{h \in \mathcal{H}} B_h(E)$. The event $\bigcap_{k \geq 0} B_{\mathcal{H}}^k([\Delta])$ is essentially equivalent to “common knowledge of $[\Delta]$ ”. We then conjecture that

$$W_{\Delta}^{\infty} = \text{proj}_{\Sigma} \left(\bigcap_{k \geq 0} B_{\phi}^k(R) \cap B_{\mathcal{H}}^k([\Delta]) \right),$$

$$S_{\Delta}^{\infty} = \text{proj}_{\Sigma} \left(\text{CSB}^{\infty}(R) \cap \bigcap_{k \geq 0} B_{\mathcal{H}}^k([\Delta]) \right).$$

The intuition is that players never get any direct information about (opponents') beliefs, and their best responses only depend on first-order beliefs; therefore, it should be possible to modify higher-order beliefs at any state of the world $\omega = (\theta, s, t) \in \bigcap_{k \geq 0} B_{\phi}^k(R \cap [\Delta])$ (or $\omega \in \text{CSB}^{\infty}(R \cap [\Delta])$) so as to obtain a corresponding state $\omega' = (\theta, s, t') \in \bigcap_{k \geq 0} B_{\phi}^k(R) \cap B_{\mathcal{H}}^k([\Delta])$ (or $\omega' \in \text{CSB}^{\infty}(R) \cap \bigcap_{k \geq 0} B_{\mathcal{H}}^k([\Delta])$).¹⁶

We also remark that Proposition 2 can be used to provide an epistemic characterization of the Iterated Intuitive Criterion (IIC) for signalling games (Cho and Kreps, 1987). In such games, $I = 2$, Θ_0 and Θ_2 are singletons, and Θ_1 corresponds to the sender's private information; terminal histories are of the form (m, a) , where m is the message sent by the sender and a is the action taken by the receiver. Now fix a probability distribution $\zeta \in \Delta(\Theta \times Z)$ over terminal nodes of the game arborescence, interpreted as a candidate equilibrium type-dependent outcome ($\text{marg}_{\Theta} \zeta$ is the prior, that is, the commonly-known initial belief of the receiver); then, let Δ represent the assumption that both players initially believe that the outcome distribution ζ will obtain. Note that, although Cho and Kreps proposed the intuitive criterion and IIC as refinements of Bayes–Nash equilibrium (indeed of sequential equilibrium), it makes sense to apply these criteria to any self-confirming equilibrium¹⁷ distribution (Kohlberg, 1990). Battigalli and Siniscalchi (2003) show that ζ is a self-confirming equilibrium outcome that passes the IIC if and only if $S_{\Delta}^{\infty} \neq \emptyset$ (see Sobel et al. (1990) for a related result). It then follows from Proposition 2 that ζ is a self-confirming equilibrium outcome satisfying IIC if and only if $\text{CSB}^{\infty}(R \cap [\Delta]) \neq \emptyset$.

¹⁶ We present a similar result, Proposition 3, in the next section, with $[\Delta]$ replaced by the event that the true state of nature is some given $\hat{\theta}$. However, the proof of Proposition 3 cannot be adapted to prove the above conjecture.

¹⁷ We mean self-confirming equilibrium with unitary beliefs (Fudenberg and Levine, 1993) of the “interim game” where each information type of the sender is a separate player.

4. Robustness of rationalizability with respect to payoff uncertainty

We conclude with a collection of results pertaining to the robustness of forward-induction reasoning to “slight” payoff uncertainty. Our analysis is similar in spirit to that of Fudenberg et al. (1988) and, especially, Dekel and Fudenberg (1990); however, our arguments do not involve payoff perturbations and limiting arguments. Rather, we relate robustness (or lack thereof) to specific assumptions about belief revision.

As in the first reference cited above, we embed a complete information game within a richer one featuring payoff uncertainty. Specifically, fix a game IG with payoff uncertainty, a state of nature $\hat{\theta} \in \Theta$, and denote by $G_{\hat{\theta}}$ the complete information game corresponding to $\hat{\theta}$. Since a complete information game is a special case of our framework (with $\Theta = \{\hat{\theta}\}$, a singleton), we can apply the procedures defined in Section 3 to $G_{\hat{\theta}}$. In particular, we shall focus on the weak and strong solution procedures obtained for $G_{\hat{\theta}}$ when no belief restrictions are assumed; in order to emphasize the dependence on $\hat{\theta}$, we shall use the notation $\{\mathbf{W}_{\hat{\theta}}^n\}_{n \geq 0}$ and $\{\mathbf{S}_{\hat{\theta}}^n\}_{n \geq 0}$ respectively. Note that $\{\mathbf{S}_{\hat{\theta}}^n\}_{n \geq 0}$ coincides with (the correlated version of) Pearce’s (1984) extensive form rationalizability for $G_{\hat{\theta}}$.

Our objective is to relate weak and strong rationalizability in $G_{\hat{\theta}}$ with assumptions about rationality and belief revision in IG . As a preliminary observation, intuition suggests that analysing the complete information game $G_{\hat{\theta}}$ should be equivalent to analysing the game IG focusing on states where (0) the state of nature is $\hat{\theta}$, (1) every player $i \in I$ would be certain of (0) conditional on every history $h \in \mathcal{H}$, ... (k + 1) every player i would be certain of (k) conditional on every history $h \in \mathcal{H}$...

The following result validates this intuition and derives its implications for weak and strong (extensive-form) rationalizability. To capture assumptions (0), (1), ... we consider the iterations of operator $B_{\mathcal{H}}$ defined by

$$B_{\mathcal{H}}(E) = \bigcap_{h \in \mathcal{H}} B_h(E)$$

and we denote by $[\hat{\theta}]$ the set of states of the world in which the state of nature is $\hat{\theta}$: that is, $[\hat{\theta}] = \{(\theta, s, t) \in \Omega : \theta = \hat{\theta}\}$.

Proposition 3. *Let IG be a game with payoff uncertainty and fix $\hat{\theta} \in \Theta$. Then, in any belief-complete type space, for all $n \geq 0$,*

- (i) $\text{proj}_S \left(\bigcap_{k=0}^n B_{\phi}^k(R) \cap \bigcap_{k \geq 0} B_{\mathcal{H}}^k([\hat{\theta}]) \right) = \mathbf{W}_{\hat{\theta}}^{n+1}$;
- (ii) $\text{proj}_S \left(\text{CSB}^n(R) \cap \bigcap_{k \geq 0} B_{\mathcal{H}}^k([\hat{\theta}]) \right) = \mathbf{S}_{\hat{\theta}}^{n+1}$.

In the setting of Proposition 3, the assumption that the state of nature is $\hat{\theta}$ is accorded the highest “epistemic priority”: it is maintained throughout the game, even at histories where the event $R \cap [\hat{\theta}]$ is falsified (furthermore, this is common certainty).

We now focus on games with distributed knowledge of the state of nature ($\Theta_0 = \{\hat{\theta}_0\}$, a singleton) and *private values*: that is, we assume that, for all $i \in I$, u_i is independent of θ_{-i} . (Note that the private values assumption makes θ_0 redundant. Thus there is no essential loss of generality in assuming that Θ_0 is a singleton. We will drop θ_0 from our notation in what follows.) Our next result may be interpreted as stating that, in such games, assigning the same epistemic priority to the events R and $[\hat{\theta}]$ (as well as to assumptions concerning the players’ beliefs about them) is actually sufficient to obtain a characterization of weak and strong rationalizability.

Proposition 4. *Let IG be a game with private values and fix $\hat{\theta} \in \Theta$. Then, in any belief-complete type space, for all $n \geq 0$,*

- (i) $\text{proj}_S \left(\bigcap_{k=0}^n B_{\phi}^k(R \cap [\hat{\theta}]) \right) = \mathbf{W}_{\hat{\theta}}^{n+1}$.
- (ii) $\text{proj}_S \left(\text{CSB}^n(R \cap [\hat{\theta}]) \right) = \mathbf{S}_{\hat{\theta}}^{n+1}$.

However, the assumption that the profile of payoff types is $\hat{\theta}$ (and that this is common certainty) may conceivably be accorded a *low* epistemic priority. In the next proposition we only assume that there is initial common certainty of $[\hat{\theta}]$. We interpret this as a form of “slight” payoff uncertainty: even though there is initial common certainty of $\hat{\theta}$ players may change their beliefs about the state of nature after they observe unexpected moves. To elaborate, consider the following, alternative set of assumptions:

- (0) every player initially assigns positive probability to every profile θ_{-i} , and probability at least $p \in (0, 1)$ to $\hat{\theta}_{-i}$;
- (1) every player initially assigns probability at least p to (0);
- ...
- (k) every player initially assigns probability at least p to $(k - 1)$.

In other words, there is *initial common p-belief* in the profile $\hat{\theta}$. The assumption of initial common certainty of $[\hat{\theta}]$ may then be seen as a limiting case of this alternative hypothesis as $p \rightarrow 1$.

It is then natural to ask whether forward-induction reasoning generally retains its bite in this setting. The main result of this subsection shows that, if the range of conceivable payoff functions is large, then the answer is negative.

More specifically, we analyse the implications of iterated correct strong belief in rationality at states in which there is common certainty *conditional on the initial history alone* that the state of nature is $\hat{\theta}$.

We say that the game *IG* is *rich* if for all $j \in I$, $s_j \in S_j$ and $\mu \in \Delta^{\mathcal{H}}(\Sigma_{-j})$ there is some $\theta_j \in \Theta_j$ such that s_j is a sequential best response to μ for θ_j .¹⁸ Embedding a complete information game within a rich game with payoff uncertainty is similar in spirit to considering “elaborations” of a given extensive game, as in Fudenberg et al. (1988).¹⁹

Proposition 5. *Let IG be a game with private values, and fix $\hat{\theta} \in \Theta$. If IG is rich, then, in any belief-complete type space, for all $n \geq 0$,*

$$\text{proj}_S \left(\text{CSB}^n(R) \cap \left(\bigcap_{k=0}^n \text{B}_{\hat{\theta}}^k([\hat{\theta}]) \right) \right) = \mathbf{W}_{\hat{\theta}}^{n+1}.$$

As was noted above, this result is related to Dekel and Fudenberg’s analysis of the robustness of iterated weak dominance with respect to “slightly incomplete information” (Dekel and Fudenberg, 1990). Indeed, the procedure they characterize coincides with $\left\{ \mathbf{W}_{\hat{\theta}}^n \right\}_{n \geq 1}$ if $G_{\hat{\theta}}$ is a perfect information game without ties between payoffs at distinct terminal histories (cf. Ben-Porath (1997)).

5. Concluding remarks

This paper contributes to the literature on rationalizability in games of incomplete information. As in some previous work of ours (Battigalli and Siniscalchi, 1999, 2002), we take an interactive-epistemology perspective whereby solutions concepts are characterized by looking at the behavioural consequences of explicit and transparent assumptions about rationality and interactive beliefs. These assumptions correspond to events in a belief-complete space of states of the world, such as the universal space of infinite hierarchies of conditional beliefs analysed by Battigalli and Siniscalchi (1999).²⁰ For this reason, we do not represent players’ possible hierarchies of beliefs about the state of nature in the game model, as is done (implicitly) in Bayesian games. Our primitives only comprise a set Θ of states of nature, players’ private information about the state of nature, a set $\mathcal{H} \cup \mathcal{Z}$ of feasible (partial and terminal) histories of action profiles (representing feasible moves and their order), and parametrized utility functions $u_i : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}$. Simultaneous-moves games (with unknown payoffs) and complete-information games are special cases. We call such a structure “game with payoff uncertainty”. Any assumption about players’ (exogenous and endogenous) beliefs is represented as an event in the belief-complete space of states of the world.

We consider two versions of Δ -rationalizability, an “umbrella solution concept” put forward and analysed by Battigalli (1999, 2003) and Battigalli and Siniscalchi (2003). These solutions concepts, weak and strong Δ -rationalizability, take as parametrically given some restrictions Δ on players’ first-order beliefs.²¹ The two solution

¹⁸ Note that a sufficient condition for richness is that each Θ_j contains an indifferent payoff type, i.e. some θ_j^* such that $u_j(\theta_j^*, \cdot)$ is constant. Alternatively, it is sufficient to assume that, for every player j and $s_j \in S_j$, there is a type $\theta_j(s_j)$ such that s_j weakly dominates for $\theta_j(s_j)$ every s'_j not realization equivalent to s_j . Weinstein and Yildiz (2007) use a similar assumption.

¹⁹ On the other hand, we emphasize that our assumptions require that players initially assign probability zero to states of nature other than $\hat{\theta}$, and there is initial common certainty of this.

²⁰ Tan and Werlang (1988) is the first paper adopting this approach, but it only covers simultaneous moves games with complete information.

²¹ Of course, we cover as a special case the possibility that there are no restrictions at all. The epistemic analysis of this special case was provided by Battigalli and Siniscalchi (1999, 2002).

concepts differ in how players are supposed to revise their beliefs upon observing unexpected moves: strong rationalizability features a forward-induction principle, weak rationalizability does not. Therefore the two notions coincide for simultaneous moves games. We prove that weak Δ -rationalizability is characterized by initial (and correct) common certainty of rationality and of the restrictions Δ , whereas strong Δ -rationalizability is characterized by (correct) common strong belief in rationality and Δ . From the latter characterization result and an equivalence result due to Battigalli and Siniscalchi (2003), we also obtain an epistemic characterization of the Iterated Intuitive Criterion for signalling games (Cho and Kreps, 1987).

Next, we use our framework to provide an epistemic analysis of the robustness of forward-induction reasoning in situations where there is initial common certainty of the true state of nature, as in complete information games, but (unlike complete information games) players may revise their beliefs about the state of nature when they face unexpected moves. We obtain epistemic characterizations of the complete-information versions of weak and strong rationalizability under alternative sets of assumptions concerning the epistemic priority (as determined by players' belief revision policies) of hypotheses concerning rationality versus hypotheses about the state of nature. In particular, we show that in a private-values setting, if the parameter space Θ is sufficiently rich, initial common certainty of the true state of nature (payoff parameter) $\hat{\theta}$ and common strong belief in rationality are equivalent to initial common certainty of rationality and $\hat{\theta}$, and yield the weakly rationalizable strategies of the complete information game corresponding to $\hat{\theta}$. This shows that forward-induction reasoning is not robust to the introduction of "slight payoff uncertainty" if the set of conceivable payoff functions is very large. This confirms a similar result obtained by Dekel and Fudenberg (1990) with a different approach.

We now discuss the related literature. For the applications of rationalizability to economic models with incomplete information, we refer to Battigalli (1999, 2003) and Battigalli and Siniscalchi (2003) and the references therein. Here we consider the theoretical literature on this topic.

The standard way to analyse incomplete-information games is to adopt Harsanyi's approach (Harsanyi, 1967–68): represent them as Bayesian games and then analyse the Bayes–Nash equilibria or the rationalizable strategies of such games. The types in a Bayesian game do not only specify a player's information, they also parametrize hierarchical beliefs about payoff-relevant parameters. Therefore, an information type may correspond to a *set* of types in the Bayesian game. Here we refer to the latter types as types *à la* Harsanyi, or simply "Harsanyi types".

Battigalli and Siniscalchi (2003) show that, in any given game with payoff uncertainty IG , Δ -rationalizability yields the set of equilibrium outcomes of all the Bayesian games consistent with IG and with the restrictions Δ . Battigalli (1999) proves a similar result relating weak rationalizability and weakly perfect Bayesian equilibrium outcomes. Bergemann and Morris (2007) obtain results of the same kind for static games with payoff uncertainty (which they call "belief-free incomplete information games") without considering restrictions on first-order beliefs. In particular, they show that the set of incomplete-information correlated-equilibrium outcomes of a game IG coincides with the set of Bayes–Nash equilibria ("interim equilibria" in their terminology) of all the common prior Bayesian games consistent with IG .

Pearce (1984) put forward a notion of extensive form rationalizability (EFR). Static and dynamic Bayesian games can be represented as extensive form games with imperfect information about an initial chance move; therefore, EFR can be applied to Bayesian games. In static Bayesian games, EFR coincides with *interim rationalizability*, i.e. rationalizability in the interim strategic form.²² Ely and Peski (2006) and Dekel et al. (2007) noticed that interim rationalizability implicitly rests on a conditional-independence assumption: conditional on the (Harsanyi) type of every opponent, each player regards θ and his opponents' actions as mutually independent. This assumption makes the rationalizable actions depend on features of Harsanyi types not captured by the corresponding hierarchical beliefs about the payoff-relevant parameters. This problem, which arises only when there is residual uncertainty about the payoff parameter (i.e. when Θ_0 is not a singleton) is addressed by these authors in different ways. Ely and Peski (2006) put forward an alternative notion of hierarchical beliefs that is sufficient to identify interim-rationalizable actions, but they do not provide an epistemic characterization of interim rationalizability *via* expressible assumptions on players' rationality and interactive beliefs. Dekel et al. (2007) instead remove the conditional independence assumption to obtain a weaker notion of *interim correlated rationalizability*, which turns out to depend only on hierarchical beliefs about the payoff parameters. They also provide an epistemic characterization of this solution concept.

²² For the definition of the interim strategic form of a Bayesian game see, for example, §2.6 in Osborne and Rubinstein (1994). A Bayes–Nash equilibrium is a Nash equilibrium of the interim strategic form.

We can relate Δ -rationalizability to EFR and interim rationalizability for Bayesian games under the assumption that the types *à la* Harsanyi correspond to our information types (which implies that information types map *onto* the set of infinite hierarchies of beliefs implicitly specified by the Bayesian game). In this case it can be shown that, in static games, interim rationalizability/EFR is equivalent to Δ -rationalizability, provided that Δ represents the belief mappings of the Bayesian Game and the conditional-independence hypothesis. Interim correlated rationalizability is equivalent to Δ -rationalizability if Δ represents the belief mappings. Battigalli et al. (2007) elaborates on the relationships between these solution concepts and their epistemic characterizations. We conjecture that a similar equivalence result holds for dynamic Bayesian games with two players: with appropriate restrictions Δ , EFR on the extensive form of the Bayesian game should be equivalent to *strong* Δ -rationalizability.²³ In particular, this would follow from our conjecture in Section 3.4 that strong Δ -rationalizability is characterized by common strong belief in rationality and “common knowledge” of the restrictions Δ .

Appendix A. Proofs

A.1. Main characterization results

Proposition 4 does *not* follow from Propositions 1 and 2, but the proofs are very similar. We shall emphasize the proof of Propositions 1 and 2, and note the modifications required to establish Proposition 4.

We begin with two preliminary results. In the first, recall that $T_{-i} = \prod_{j \in I \setminus \{i\}} T_j$ whereas $\Sigma_{-i} = \prod_{j \in \{0\} \cup I \setminus \{i\}} S_j$, but this is immaterial.

Lemma 6. Fix a map $\tau_{-i} : \Sigma_{-i} \rightarrow T_{-i}$. Also, fix a first-order CPS $\delta_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$. Then there exists an epistemic type $t_i \in T_i$ such that, for each $h \in \mathcal{H}$, $g_{i,h}(t_i)$ has finite support and

$$g_{i,h}(t_i)((\sigma_{-i}, \tau_{-i}(\sigma_{-i}))) = \delta_i(\sigma_{-i} | \Sigma_{-i}(h))$$

for all $\sigma_{-i} \in \Sigma_{-i}$.

Proof. Define a candidate CPS μ_i on $\Sigma_{-i} \times T_{-i}$ by setting

$$\mu_i(\{(\sigma_{-i}, \tau_{-i}(\sigma_{-i}))\} | \Sigma_{-i}(h) \times T_{-i}) = \delta_i(\sigma_{-i} | \Sigma_{-i}(h))$$

for every $h \in \mathcal{H}$, and extending the assignments by additivity. Axioms 1 and 2 follow immediately from the observation that the map $\sigma_{-i} \mapsto (\sigma_{-i}, \tau_{-i}(\sigma_{-i}))$ yields an embedding of $\bigcup_{h \in \mathcal{H}} \text{supp}[\delta_i(\cdot | \Sigma_{-i}(h))] \subset \Sigma_{-i}$ (a finite set) in $\Sigma_{-i} \times T_{-i}$, so that, for every $h \in \mathcal{H}$, $\mu_i(\cdot | \Sigma_{-i}(h) \times T_{-i})$ is indeed a probability measure on $\Sigma_{-i} \times T_{-i}$. By the same argument, μ_i must also satisfy Axiom 3, i.e. it must be a CPS; of course, each $\mu_i(\cdot | \Sigma_{-i}(h) \times T_{-i})$ has finite support by construction. Since g_i is onto, there exists a type $t_i \in T_i$ such that

$$g_{i,h}(t_i)((\sigma_{-i}, \tau_{-i}(\sigma_{-i}))) = \mu_i((\sigma_{-i}, \tau_{-i}(\sigma_{-i})) | \Sigma_{-i}(h) \times T_{-i}) = \delta_i(\sigma_{-i} | \Sigma_{-i}(h))$$

for all $\sigma_{-i} \in \Sigma$ and $h \in \mathcal{H}$. ■

The next lemma provides an alternative characterization of $\{\mathbf{S}_{\Delta}^n\}_{n=0}^{\infty}$, where $\Delta = (\Delta_{\theta_i})_{i \in I, \theta_i \in \Theta_i}$ is any regular collection of subsets of CPSs.

Lemma 7. Suppose Δ is regular. For every $i \in I$ and $n \geq 1$, $(\theta_i, s_i) \in \mathbf{S}_{i,\Delta}^n$ if and only if there exists a CPS $\mu \in \Delta_{\theta_i}$ such that $(\theta_i, s_i) \in r_i(\mu)$ and

$$\forall m = 0, \dots, n - 1, \forall h \in \mathcal{H} : \mathbf{S}_{-i,\Delta}^m \cap \Sigma_{-i}(h) \neq \emptyset \Rightarrow \mu(\mathbf{S}_{-i,\Delta}^m | \Sigma_{-i}(h)) = 1. \tag{A.2}$$

Proof. The statement is obvious for $n = 1$. Now pick $n \geq 2$ and assume it is true for $m = 0, \dots, n - 1$. If $\sigma_i = (\theta_i, s_i) \in r_i(\mu)$ for some $\mu \in \Delta_{\theta_i}$ satisfying (A.2), then $\sigma_i \in \mathbf{S}_{i,\Delta}^{n-1}$ by the induction hypothesis, because

²³ Our conjecture considers only two-person games because Battigalli (1996) shows that, in n -person games, Pearce’s EFR relies on an ill-conceived structural consistency property. A result for n -person games should therefore refer to amended versions of EFR: either one that allows for correlation in beliefs about opponents, or one that incorporates an appropriate independence assumption.

$\mathbf{S}_{-i,\Delta}^m \cap \Sigma_{-i}(h) \neq \emptyset \Rightarrow \mu(\mathbf{S}_{-i,\Delta}^m | \Sigma_{-i}(h)) = 1$ for $m = 0 \dots n - 2$; moreover, since also $\mathbf{S}_{-i,\Delta}^{n-1} \cap \Sigma_{-i}(h) \neq \emptyset \Rightarrow \mu(\mathbf{S}_{-i,\Delta}^{n-1} | \Sigma_{-i}(h)) = 1$, and $\sigma_i \in r_i(\mu)$, we conclude $\sigma_i \in \mathbf{S}_{i,\Delta}^n$.

In the other direction, suppose $\sigma_i = (\theta_i, s_i) \in \mathbf{S}_{i,\Delta}^n$. Then also $\sigma_i \in \mathbf{S}_{i,\Delta}^m$ for $m = 0, \dots, n - 1$, so we can find CPSs $\mu^m \in \Delta_{\theta_i}$, $m = 0, \dots, n - 1$, such that, for each such m , $\sigma_i \in r_i(\mu^m)$ and, for any $h \in \mathcal{H}$, $\mathbf{S}_{-i,\Delta}^m \cap \Sigma_{-i}(h) \neq \emptyset$ implies $\mu^m(\mathbf{S}_{-i,\Delta}^m | \Sigma_{-i}(h)) = 1$. Now construct a new CPS μ as follows: for any $h \in \mathcal{H}$, let $m(h) = \max\{m = 0, \dots, n - 1 : \mathbf{S}_{-i,\Delta}^m \cap \Sigma_{-i}(h) \neq \emptyset\}$, and define $\mu(\cdot | \Sigma_{-i}(h)) = \mu^{m(h)}(\cdot | \Sigma_{-i}(h))$. It is easy to verify that this is a well-defined CPS, i.e. $\mu \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ (for a similar construction, see e.g. Battigalli (1997)).

By construction, $\mu(\cdot | \Sigma_{-i}(h)) \in \Delta_{\theta_i,h}$ for all h . By definition of regularity, $\Delta_{\theta_i} = \Delta^{\mathcal{H}}(\Sigma_{-i}) \cap \prod_{h \in \mathcal{H}} \Delta_{\theta_i,h}$. Therefore $\mu \in \Delta_{\theta_i}$. Moreover, clearly $\sigma_i \in r_i(\mu)$. Finally, μ satisfies (A.2), which concludes the proof. ■

Note that Lemma 7 also applies to games with complete information (take Θ_i to be a singleton for each i); hence, in the setting of Section 4, it applies to the game $G_{\hat{\theta}}$ and the sets $\mathbf{S}_{\hat{\theta}}^n$, $n = 0, 1, \dots$

We can finally prove our main result.

Proof of Proposition 2. To prove (i), we proceed by induction, assuming first that the sets appearing in the statement are non-empty.

(Step 0.) Fix $(\sigma, t) = (\theta_0, s_0, \theta_1, s_1, \dots, \theta_{|I|}, s_{|I|}) \in \text{CSB}^0(R \cap [\Delta]) = R \cap [\Delta]$. Then by definition $\sigma_i = (\theta_i, s_i) \in r_i(f_i(t_i))$ and $f_i(t_i) \in \Delta_{\theta_i}$ for every $i \in I$, which implies that $\sigma \in \mathbf{S}_{\Delta}^1$.

Conversely, for each $i \in I$ and $\sigma_i \in \Sigma_i$, pick $\tau_i^0(\sigma_i) \in T_i$ arbitrarily. Now fix $\sigma \in \mathbf{S}_{\Delta}^1$, and for each player $i \in I$ and $\theta_i \in \Theta_i$, let $\mu_i \in \Delta_{\theta_i}$ be such that $\sigma_i = (\theta_i, s_i) \in r_i(\mu_i)$. Now Lemma 6 yields a type $\tau_i^1(\sigma_i) \in T_i$ such that $g_{i,h}(\tau_i^1(\sigma_i))(\{\sigma'_j, \tau_j^0(\sigma_j)\}_{j \neq i}) = \mu_i(\sigma'_i | \Sigma_{-i}(h))$ for every $\sigma'_i \in \Sigma_{-i}$, and hence $f_i(\tau_i^1(\sigma_i)) = \mu_i$. Thus, $(\sigma_i, \tau_i^1(\sigma_i))_{i \in I} \in R \cap [\Delta]$.

Finally, for each $i \in I$, we complete the definition of the function $\tau_i^1(\cdot)$ by letting $\tau_i^1(\sigma_i) = \tau_i^0(\sigma_i)$ for $\sigma_i \in \Sigma_i \setminus \mathbf{S}_{i,\Delta}^1$.

(Step $n > 0$.) Now assume that Part (i) has been shown to hold for $m = 0, \dots, n - 1$, and that, for each such m , we have defined functions $\tau_i^{m+1} : \Sigma_i \rightarrow T_i$ such that $(\sigma_i, \tau_i^{m+1}(\sigma_i))_{i \in I} \in \text{CSB}^m(R \cap [\Delta])$ whenever $\sigma \in \mathbf{S}_{\Delta}^{m+1}$. Finally, let the functions $\tau_i^0(\cdot)$ be defined as above.

Note that, for any event $E \in \mathcal{A}$ and $n \geq 1$,

$$\text{CSB}^n(E) = E \cap \bigcap_{i \in I} \left\{ \bigcap_{m=0}^{n-1} \text{SB}_i(\Omega_i \times [\text{proj}_{\Omega_{-i}} \text{CSB}^m(E)]) \right\}. \tag{A.3}$$

Also note that, for any $i \in I$, $h \in \mathcal{H}$ and event E such that $\text{proj}_{\Omega_i} E = \Omega_i$,

$$E \cap (\Sigma(h) \times T) \neq \emptyset \Leftrightarrow [\text{proj}_{\Sigma_{-i}} E] \cap \Sigma_{-i}(h) \neq \emptyset. \tag{A.4}$$

Now consider $(\sigma, t) \in \text{CSB}^n(R \cap [\Delta])$ and fix $i \in I$. By Eq. (A.3) (taking $E = R \cap [\Delta]$) we conclude that $\sigma_i = (\theta_i, s_i) \in r_i(f_i(t_i))$ and $f_i(t_i) \in \Delta_{\theta_i}$; also, for any $m = 0, \dots, n - 1$, the induction hypothesis and Eq. (A.4) imply that, for any $h \in \mathcal{H}$, $\mathbf{S}_{-i,\Delta}^{m+1} \cap \Sigma_{-i}(h) = [\text{proj}_{\Sigma_{-i}} \text{CSB}^m(R \cap [\Delta])] \cap \Sigma_{-i}(h) \neq \emptyset$ if and only if $[\Omega_i \times \text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])] \cap (\Sigma(h) \times T) \neq \emptyset$. Now Eq. (A.3) and the definition of strong belief implies that, for any $h \in \mathcal{H}$ satisfying the latter condition for some $m = 0, \dots, n - 1$, $g_{i,h}(t_i)(\text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])) = 1$. This implies $f_i(t_i)(\text{proj}_{\Sigma_{-i}} \text{CSB}^m(R \cap [\Delta]) | \Sigma_{-i}(h)) = 1$; in turn, the induction hypothesis implies $f_i(t_i)(\mathbf{S}_{-i,\Delta}^{m+1} | \Sigma_{-i}(h)) = 1$. Hence, Lemma 7 implies that $\sigma_i \in \mathbf{S}_{i,\Delta}^{n+1}$.

For the converse implication, begin by defining

$$m_i(\sigma_i) = \max\{m = 0, \dots, n : \sigma_i \in \mathbf{S}_{i,\Delta}^m\}$$

for every $i \in I$ and $\sigma_i \in \Sigma_i$; recall that $\mathbf{S}_{i,\Delta}^0 = \Sigma_i$, so $m_i(\cdot)$ is well-defined for every $\sigma_i \in \Sigma_i$. Now consider $\sigma \in \mathbf{S}_{\Delta}^{n+1}$ and fix a player $i \in I$. By Lemma 7, we can find a CPS $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ satisfying Eq. (A.2). By Eq. (A.4) and the induction hypothesis, for $h \in \mathcal{H}$ and $m = 0, \dots, n - 1$, $[\Omega_i \times \text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])] \cap (\Sigma(h) \times T) \neq \emptyset$ if and only if $\mathbf{S}_{-i,\Delta}^{m+1} \cap \Sigma_{-i}(h) \neq \emptyset$. But if the latter inequality holds, $\mu_i(\mathbf{S}_{-i,\Delta}^{m+1} | \Sigma_{-i}(h)) = 1$ by Eq. (A.2).

Now define $\tau_{-i} : \Sigma_{-i} \rightarrow T_{-i}$ by letting

$$\tau_{-i}(\sigma'_{-i}) = (\tau_j^{m_j(\sigma'_j)}(\sigma'_j))_{j \neq i} \quad \forall \sigma'_{-i} \in \Sigma_{-i};$$

Lemma 6 now yields a type $\tau_i^{n+1}(\sigma_i) \in T_i$ such that

$$g_{i,h}(\tau_i^{n+1}(\sigma_i))(\{\sigma'_j, \tau_j^{m_j(\sigma'_j)}(\sigma'_j)\}_{j \neq i}) = \mu_i(\{\sigma'_{-i}\} | \Sigma_{-i}(h))$$

for all $h \in \mathcal{H}$ and $\sigma'_{-i} \in \Sigma_{-i}$. Now note that, for $m = 0, \dots, n - 1$,

$$\sigma'_{-i} \in \mathbf{S}_{-i,\Delta}^{m+1} \Rightarrow (\sigma'_j, \tau_j^{m_j(\sigma'_j)}(\sigma'_j))_{j \neq i} \in \text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])$$

because, for all $j \neq i$: (a) $m_j(\sigma'_j) \geq m + 1$ if $\sigma'_{-i} \in \mathbf{S}_{-i,\Delta}^{m+1}$; (b) if $m_j(\sigma'_j) \geq 1$ then, by the induction hypothesis,

$$(\sigma'_j, \tau_j^{m_j(\sigma'_j)}(\sigma'_j)) \in \text{proj}_{\Omega_j} \text{CSB}^{m_j(\sigma'_j)-1}(R \cap [\Delta]);$$

and finally (c) the sets $(\text{CSB}^m(R \cap [\Delta]))_{m \geq 0}$ are monotonically decreasing. But then

$$g_{i,h}^*(\sigma_i, \tau_i^{n+1}(\sigma_i))(\Omega_i \times \text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])) = 1$$

for any $m = 0 \dots n - 1$ and $h \in \mathcal{H}$ such that $[\Omega_i \times \text{proj}_{\Omega_{-i}} \text{CSB}^m(R \cap [\Delta])] \cap (\Sigma(h) \times T) \neq \emptyset$, because by the argument above $\text{supp } \mu(\cdot | \Sigma_{-i}(h)) \subset \mathbf{S}_{-i,\Delta}^{m+1}$ at any such history.

Moreover, since by construction $f_i(\tau_i^{n+1}(\sigma_i)) = \mu_i$, we also have $\sigma_i = (\theta_i, s_i) \in r_i(f_i(\tau_i^{n+1}(\sigma_i)))$ and $f_i(\tau_i^{n+1}(\sigma_i)) \in \Delta_{\theta_i}$.

Repeating the argument for every $i \in I$ yields a profile of types $(\tau_i^{n+1}(\sigma_i))_{i \in I}$ which, by Eq. (A.3), satisfies $(\sigma_i, \tau_i^{n+1}(\sigma_i))_{i \in I} \in \text{CSB}^n(R \cap [\Delta])$. To complete the induction step, for each $i \in I$ we now define the function $\tau_i^{n+1}(\cdot)$ for $\sigma_i \in \Sigma_i \setminus \mathbf{S}_{i,\Delta}^{n+1}$ by letting $\tau_i^{n+1}(\sigma_i) = \tau_i^n(\sigma_i)$ for any such strategy σ_i .

The argument just given shows that if one of the sets appearing in the statement of (i) is non-empty, so is the other one. Hence, the proof of (i) is complete.

For Part (ii), assume first that $\mathbf{S}_{\Delta}^{\infty} \neq \emptyset$. Then $\mathbf{S}_{\Delta}^n \neq \emptyset$ for all $n \geq 0$; hence $\text{CSB}^n(R \cap [\Delta]) \neq \emptyset$ for $n \geq 0$ by Part (i). Then $\text{CSB}^{\infty}(R \cap [\Delta])$ is non-empty, because T is compact by assumption and the nested, non-empty closed sets $\{\text{CSB}^n(R \cap [\Delta])\}_{n \geq 0}$ form a family with the finite intersection property.

Now suppose $(\sigma, t) \in \text{CSB}^{\infty}(R \cap [\Delta])$. Since, by Part (i), $\mathbf{S}_{\Delta}^{n+1} = \text{proj}_{\Sigma} \text{CSB}^n(R \cap [\Delta])$ for any $n \geq 0$, we conclude that $\sigma \in \mathbf{S}_{\Delta}^n$ for every $n \geq 1$; so $\sigma \in \bigcap_{n \geq 1} \mathbf{S}_{\Delta}^n = \mathbf{S}_{\Delta}^{\infty}$. Hence $\text{proj}_{\Sigma} \text{CSB}^{\infty}(R \cap [\Delta]) \subset \mathbf{S}_{\Delta}^{\infty}$.

Next, let N be the smallest integer such that $\mathbf{S}_{\Delta}^N = \mathbf{S}_{\Delta}^{\infty}$ (which must exist because Σ is finite). Pick any $\sigma \in \mathbf{S}_{\Delta}^N = \mathbf{S}_{\Delta}^{\infty}$ and consider the sequence of sets $M(m, \sigma) = \text{CSB}^{(N-1)+m}(R \cap [\Delta]) \cap (\{\sigma\} \times T)$, $m \geq 0$ (let $M(0, \sigma) = \{\sigma\} \times T$ if $N = 0$). Each set $M(m, \sigma)$ is non-empty and closed; also, the sequence of sets $M(m, \sigma)$ is decreasing, and hence has the finite intersection property. Then $\emptyset \neq \bigcap_{m \geq 0} M(m, \sigma) \subset \text{CSB}^{\infty}(R \cap [\Delta])$, so $\mathbf{S}_{\Delta}^{\infty} \subset \text{proj}_{\Sigma} \text{CSB}^{\infty}(R \cap [\Delta])$.

If $\mathbf{S}_{\Delta}^{\infty} = \emptyset$, let N be the smallest integer such that $\mathbf{S}_{\Delta}^N = \emptyset$. Since $\mathbf{S}_{\Delta}^N = \text{proj}_{\Sigma} \text{CSB}^{N-1}(R \cap [\Delta])$, we conclude that $\text{CSB}^{N-1}(R \cap [\Delta]) = \emptyset$, so $\text{CSB}^{\infty}(R \cap [\Delta]) = \emptyset$, and again $\mathbf{S}_{\Delta}^{\infty} = \text{proj}_{\Sigma} \text{CSB}^{\infty}(R \cap [\Delta])$. ■

The proof of Proposition 1 is analogous, but simpler; we only provide a sketch.²⁴ To elaborate, Step (0) is identical. In Step n , for any $(\sigma, t) \in \bigcap_{k=0}^n \mathbf{B}_{\phi}^k(R \cap [\Delta])$, it is enough to verify that $g_{i,\phi}(t_i)(\text{marg}_{\Omega_{-i}} \bigcap_{k=0}^{n-1} \mathbf{B}_{\phi}^k(R \cap [\Delta])) = 1$, so that, by the induction hypothesis, $f_{i,\phi}(t_i)(\mathbf{W}_{-i,\Delta}^n) = 1$; the arguments given above then imply that $\sigma \in \mathbf{W}_{\Delta}^{n+1}$, as required. Conversely, for $\sigma \in \mathbf{W}_{\Delta}^{n+1}$ and $i \in I$, Definition 5 implies that $\sigma_i = (s_i, \theta_i) \in r_i(\mu)$ for some $\mu \in \Delta_{\theta_i}$ such that $\mu(\mathbf{W}_{-i,\Delta}^n | \Sigma_{-i}) = 1$; invoking the induction hypothesis, for every $\sigma'_{-i} \in \text{supp } \mu$ one can find $\tau_{-i}(\sigma'_{-i}) \in T_{-i}$ such that $(\sigma'_{-i}, \tau_{-i}(\sigma'_{-i})) \in \text{proj}_{\Omega_{-i}} \bigcap_{k=0}^n \mathbf{B}_{\phi}^k(R \cap [\Delta])$. It is then sufficient to construct a type $t_i \in T_i$ such that $f_i(t_i) = \mu$ and $g_{i,\phi}(t_i)(\{\sigma'_{-i}, \tau_{-i}(\sigma'_{-i})\}) = \mu(\{\sigma'_{-i}\} | \Sigma_{-i})$ for all $\sigma'_{-i} \in \text{supp } \mu$.

Finally, the following observations allow one to modify the proof of Proposition 2 to prove Part (ii) of Proposition 4. Fix a player $i \in I$ and a profile of payoff types $\hat{\theta} \in \Theta$.

First, note that, for any CPS $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$, one can define a “marginal” CPS $\mu_i^S \in \Delta^{\mathcal{H}}(S_{-i})$ by letting $\mu_i^S(\{s_{-i}\} | \Sigma_{-i}(h)) = \mu_i(\Theta_{-i} \times \{s_{-i}\} | \Sigma_{-i}(h))$ for each $h \in \mathcal{H}$; conversely, for any CPS $v_i^S \in \Delta^{\mathcal{H}}(S_{-i})$, one can define a CPS $v_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ by letting $v_i(\{\hat{\theta}_{-i}, s_{-i}\} | \Sigma_{-i}(h)) = v_i^S(\{s_{-i}\} | S_{-i}(h))$ for each $h \in \mathcal{H}$.

²⁴ An analogous result is proved in Battigalli and Siniscalchi (1999), without exogenous restrictions on beliefs.

With these definitions, for any strategy $s_i \in S_i$, by the private values assumption $(\hat{\theta}_i, s_i) \in r_i(\mu_i)$ implies $s_i \in r_{i,\hat{\theta}}(\mu_i^S)$, and conversely $s_i \in r_{i,\hat{\theta}}(\nu_i^S)$ implies $(\hat{\theta}_i, s_i) \in r_i(\nu_i)$, where $r_{i,\hat{\theta}}(\cdot)$ denotes Player i 's best response correspondence in the game $G_{\hat{\theta}}$. This allows one to adapt Step 0 in the above proof of Part (i) and show that $S_{\hat{\theta}}^1 = W_{\hat{\theta}}^1 = \text{proj}_S R \cap [\hat{\theta}]$. We leave the rest of the proof of Part (i) to the reader.

As for the proof of Part (ii), note that, since $\Sigma(h) = \Theta \times S(h)$ for all $h \in \mathcal{H}$, $E \cap (\Sigma(h) \times T) \neq \emptyset$ if and only if $[\text{proj}_{S_{-i}} E] \cap S_{-i}(h) \neq \emptyset$ for any event E such that $\text{proj}_{S_i} E = S_i$. The inductive step in the proof of Part (i) of Proposition 2 may then be easily adapted to the present context. Again we leave the details to the reader.

A.2. Other proofs

Proof of Proposition 3. Given a belief-complete type space for game IG we derive a belief-complete type space for game $G_{\hat{\theta}}$ as follows:

For all $k = 0, 1, \dots$ and $i \in I$, let $T_{\hat{\theta},i}^0 = T_i$,

$$T_{\hat{\theta},i}^{k+1} = \left\{ t_i \in T_{\hat{\theta},i}^k : \forall h \in \mathcal{H}, g_{i,h}(t_i) \left(\prod_{j \neq i} \{\hat{\theta}_j\} \times S_j \times T_{\hat{\theta},j}^k \right) = 1 \right\}$$

and

$$T_{\hat{\theta},i} = \bigcap_{k \geq 0} T_{\hat{\theta},i}^k.$$

We take $T_{\hat{\theta},i}$ to be Player i 's space of epistemic types in game $G_{\hat{\theta}}$ and define the belief mapping $g_{\hat{\theta},i} : T_{\hat{\theta},i} \rightarrow \Delta^{\mathcal{H}}$ ($\prod_{j \neq i} S_j \times T_{\hat{\theta},j}$) so that, for all $t_i \in T_{\hat{\theta},i}$, $g_{\hat{\theta},i}(t_i)$ is the CPS satisfying

$$\begin{aligned} & \forall h \in \mathcal{H}, \forall s_{-i} \in S_{-i}, \forall K_{-i} \subset T_{\hat{\theta},-i} \text{ (measurable),} \\ & g_{\hat{\theta},i,h}(t_i) (\{s_{-i}\} \times K_{-i}) = g_{i,h}(t_i) (\{\hat{\theta}_{-i}, s_{-i}\} \times K_{-i}). \end{aligned}$$

(we abuse notation in writing ordered tuples and Cartesian products: the meaning is obvious). By construction $(T_{\hat{\theta},i}, g_{\hat{\theta},i})_{i \in I}$ defines a belief-complete type space for game $G_{\hat{\theta}}$ and for all $i \in I$, $(s, t) \in S \times T_{\hat{\theta}}$, $h \in \mathcal{H}$, $E \subset S \times T_{\hat{\theta}}$ (measurable),

$$(s, t) \in B_{\hat{\theta},i,h}(E) \Leftrightarrow (\hat{\theta}, s, t) \in B_{i,h}(\{\hat{\theta}\} \times E)$$

and

$$(s, t) \in R_{\hat{\theta}} \Leftrightarrow (\hat{\theta}, s, t) \in R \cap [\hat{\theta}],$$

where $B_{\hat{\theta},i,h}$ and $R_{\hat{\theta}}$ denote the (i, h) -belief operator and the rationality event in the type space for game $G_{\hat{\theta}}$. By Propositions 1 and 2 these equivalences imply the thesis. ■

Proof of Proposition 5. The statement is obviously true for $n = 0$. Suppose it is true for index $n - 1$. It can be easily shown by induction that, for every event E , $\text{CSB}^n(E) \subset \bigcap_{k=0}^n B_{\phi}^k(E)$, which implies

$$\text{CSB}^n(R) \cap \left(\bigcap_{k=0}^n B_{\phi}^k([\hat{\theta}]) \right) \subset \bigcap_{k=0}^n \left(B_{\phi}^k(R) \cap B_{\phi}^k([\hat{\theta}]) \right) = \bigcap_{k=0}^n \left(B_{\phi}^k(R \cap [\hat{\theta}]) \right).$$

Therefore, by Proposition 4, $\text{proj}_S(\text{CSB}^n(R) \cap (\bigcap_{k=0}^n B_{\phi}^k([\hat{\theta}]))) \subset W_{\hat{\theta}}^{n+1}$.

Assume that $s \in W_{\hat{\theta}}^{n+1}$. Fix $i \in I$. By assumption there exists a CPS $\nu \in \Delta^{\mathcal{H}(S_{-i})}$ such that $s_i \in r_{\hat{\theta},i}(\nu)$ ($r_{\hat{\theta},i}$ is Player i 's best-response correspondence in $G_{\hat{\theta}}$) and $\nu(W_{\hat{\theta},-i}^n | S_{-i}) = 1$. We now construct a CPS $\mu \in \Delta^{\mathcal{H}(\Omega_{-i})}$ having ν as marginal CPS on S_{-i} .

By the induction hypothesis, $\text{proj}_{S_{-i}}(\text{CSB}^{n-1}(R) \cap (\bigcap_{k=0}^{n-1} B_{\phi}^k([\hat{\theta}]))) = W_{\hat{\theta},-i}^n$. Hence, for any $s_{-i} \in W_{\hat{\theta},-i}^n$ we can find $\theta_{-i}(s_{-i}) \in \Theta_{-i}$ and $t_{-i}(s_{-i}) \in T_{-i}$ such that $(\theta_{-i}(s_{-i}), s_{-i}, t_{-i}(s_{-i})) \in \text{proj}_{\Omega_{-i}} \text{CSB}^{n-1}(R) \cap (\bigcap_{k=0}^{n-1} B_{\phi}^k([\hat{\theta}]))$.

Since the game IG is rich, $\text{proj}_{S_{-i}} \Sigma^n = S_{-i}$, where Σ^n is the result of the procedure in Definition 6 when there are no restrictions on first-order beliefs. By Proposition 2 $\text{proj}_{\Sigma} \text{CSB}^{n-1}(R) = \Sigma$. Therefore $\text{proj}_{S_{-i}} \text{CSB}^{n-1}(R) = S_{-i}$, and for every $s_{-i} \in S_{-i} \setminus W_{\hat{\theta}, -i}^n$ we can find $\theta_{-i}(s_{-i}) \in \Theta_{-i}$ and $t_{-i}(s_{-i}) \in T_{-i}$ such that $(\theta_{-i}(s_{-i}), s_{-i}, t_{-i}(s_{-i})) \in \text{proj}_{\Omega_{-i}} \text{CSB}^{n-1}(R)$.

We have thus defined a map $s_{-i} \mapsto (\theta_{-i}(s_{-i}), s_{-i}, t_{-i}(s_{-i}))$ which provides an embedding of S_{-i} into Ω_{-i} . As in the proof of Lemma 6, we can then construct a CPS $\mu \in \Delta^{\mathcal{H}}(\Omega_{-i})$ such that, for all $s_{-i} \in S_{-i}$ and $h \in \mathcal{H}$, $\mu(\{\theta_{-i}(s_{-i}), s_{-i}, t_{-i}(s_{-i})\} | \Sigma_{-i}(h) \times T_{-i}) = v(s_{-i} | S_{-i}(h))$. Therefore,

$$\mu \left(\text{proj}_{\Omega_{-i}} \text{CSB}^{n-1}(R) \cap \left(\bigcap_{k=0}^{n-1} B_{\phi}^k([\hat{\theta}]) \right) | \Sigma_{-i}(\phi) \times T_{-i} \right) = v(W_{\hat{\theta}, -i}^n | S_{-i}(\phi)) = 1$$

and $\mu(\text{proj}_{\Omega_{-i}} \text{CSB}^{n-1}(R) | \Sigma_{-i}(h) \times T_{-i}) = 1$ for all $h \in \mathcal{H}$.

Since we are considering a belief-complete space there is an epistemic type $t_i \in T_i$ such that $g_i(t_i) = \mu$. By the private values assumption $(\hat{\theta}_i, s_i) \in r_i(f_i(t_i))$.

Repeat the same construction for each player and let (s, t) be the tuple of strategies and epistemic types thus obtained. It now follows that

$$(\hat{\theta}, s, t) \in \text{CSB}^n(R) \cap \left(\bigcap_{k=0}^n B_{\phi}^k([\hat{\theta}]) \right).$$

This concludes the proof. ■

References

- Battigalli, P., 1996. Strategic rationality orderings and the best rationalization principle. *Games and Economic Behavior* 13, 178–200.
- Battigalli, P., 1997. On rationalizability in extensive games. *Journal of Economic Theory* 74, 40–61.
- Battigalli, P., 1999. Rationalizability in incomplete-information games. European University Institute, Department of Economics, .wp. 99/11.
- Battigalli, P., 2003. Rationalizability in infinite, dynamic games of incomplete information. *Research in Economics* 57, 1–38.
- Battigalli, P., Di Tillio, A., Grillo, E., 2007. Interactive epistemology and solution concepts for games with asymmetric information. Mimeo, Bocconi University.
- Battigalli, P., Siniscalchi, M., 1999. Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *Journal of Economic Theory* 88, 188–230.
- Battigalli, P., Siniscalchi, M., 2002. Strong belief and forward induction reasoning. *Journal of Economic Theory* 106, 356–391.
- Battigalli, P., Siniscalchi, M., 2003. Rationalization and incomplete information. *Advances in Theoretical Economics* 3 (Article 3). <http://www.bepress.com/bejte/advances/vol3/iss1/art3>.
- Ben-Porath, E., 1997. Rationality, Nash equilibrium and backwards induction in perfect information games. *Review of Economic Studies* 64, 23–46.
- Bergemann, D., Morris, S., 2007. Belief free incomplete information games, Cowles Foundation Discussion Paper 1629.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* 59, 189–198.
- Brandenburger, A., 2003. On the existence of a ‘complete’ possibility structure. In: Basili, M., Dimitri, N., Gilboa, I. (Eds.), *Cognitive Processes and Economic Behavior*. Routledge, London.
- Cho, I.K., Kreps, D., 1987. Signalling games and stable equilibria. *Quarterly Journal of Economics* 102, 179–221.
- Dekel, E., Fudenberg, D., 1990. Rational play under payoff uncertainty. *Journal of Economic Theory* 52, 243–267.
- Dekel, E., Fudenberg, D., Morris, S., 2007. Interim correlated rationalizability. *Theoretical Economics* 2, 15–40.
- Ely, J., Peski, M., 2006. Hierarchies of belief and interim rationalizability. *Theoretical Economics* 1, 19–65.
- Fagin, R., Halpern, J., Moses, Y., Vardi, M., 1995. *Reasoning About Knowledge*. MIT Press, Cambridge, MA.
- Fudenberg, D., Kreps, D., Levine, D.K., 1988. On the robustness of equilibrium refinements. *Journal of Economic Theory* 44, 354–380.
- Fudenberg, D., Levine, D.K., 1993. Self-confirming equilibrium. *Econometrica* 61, 523–545.
- Fudenberg, D., Tirole, J., 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Harsanyi, J., 1967–68. Games of incomplete information played by bayesian players. Parts I, II, III. *Management Science* 14, 159–182; 320–334; 486–502.
- Kohlberg, E., 1990. Refinement of Nash equilibrium: The main ideas. In: Ichiishi, T., Neyman, A., Tauman, Y. (Eds.), *Game Theory and Applications*. Academic Press, San Diego.
- Mertens, J.F., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14, 1–29.
- Myerson, R., 1986. Multistage games with communication. *Econometrica* 54, 323–358.
- Osborne, M., Rubinstein, A., 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.
- Pearce, D., 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1050.

- Reny, P., 1992. Backward induction, normal-form perfection and explicable equilibria. *Econometrica* 60, 626–649.
- Rényi, A., 1955. On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungaricae* 6, 285–335.
- Rubinstein, A., 1991. Comments on the interpretation of game theory. *Econometrica* 59, 909–904.
- Sobel, J., Stole, L., Zapater, I., 1990. Fixed-equilibrium rationalizability in signaling games. *Journal of Economic Theory* 52, 304–331.
- Tan, T., Werlang, S., 1988. The Bayesian foundation of solution concepts of games. *Journal of Economic Theory* 45, 370–391.
- Weinstein, J., Yildiz, M., 2007. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica* 75, 365–400.