Economic Theory (2019) 68:737–763
https://doi.org/10.1007/s00199-018-1142-8

RESEARCH ARTICLE

Interactive epistemology in simple dynamic games with a continuum of strategies

Pierpaolo Battigalli[1] · **Pietro Tebaldi**[2]

Received: 28 August 2017 / Accepted: 8 August 2018 / Published online: 17 August 2018

Abstract
We extend the epistemic analysis of dynamic games of Battigalli and Siniscalchi (J Econ Theory 88:188–230, 1999, J Econ Theory 106:356–391, 2002, Res Econ 61:165–184, 2007) from finite dynamic games to all *simple* games, that is, finite and infinite-horizon multistage games with finite action sets at nonterminal stages and compact action sets at terminal stages. We prove a generalization of Lubin's (Proc Am Math Soc 43:118–122, 1974) extension result to deal with conditional probability systems and strong belief. With this, we can provide a short proof of the following result: in every simple dynamic game, strong rationalizability characterizes the behavioral implications of rationality and common strong belief in rationality.

Keywords Epistemic game theory · Simple infinite dynamic game · Strong belief · Strong rationalizability

JEL Classification C72 · C73 · D82

We are very grateful to Gabriele Beneduci, for excellent research assistance and for proving a generalization of the main technical lemma of this paper. We thank Federico Bobbio, Roberto Corrao, Enrico De Magistris and Giacomo Lanzani for careful proof reading and Emiliano Catonini, Nicodemo De Vito, Michael Greinecker, the Editor, and an anonymous referee for helpful suggestions. Pierpaolo Battigalli gratefully acknowledges financial support from the European Research Council (Grant 324219).

✉ Pierpaolo Battigalli
pierpaolo.battigalli@unibocconi.it

Pietro Tebaldi
ptebaldi@uchicago.edu

[1] Department of Decision Sciences and IGIER, Bocconi University, Via Roentgen, 1, 20136 Milan, Italy

[2] Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637, USA

1 Introduction

Battigalli and Siniscalchi (henceforth BS) put forward the construction of a "canonical" type structure of all hierarchies of conditional probability systems (CPSs) that satisfy coherence and common full belief in coherence (Battigalli and Siniscalchi 1999). This construction applies to a large class of *simple* infinite dynamic games with compact metric strategy spaces, that is, finite and infinite-horizon multistage games with finite action sets at nonterminal stages and compact action sets at terminal stages (see Battigalli 2003). Simple games include the infinite discounted repetitions of finite games, and all finite-horizon games where feasible action sets are compact subsets of $\mathbb{R}^n$ in the last stage and finite in previous stages. Thus, for all such games one can construct a corresponding canonical structure that is complete (belief maps are onto), compact, and continuous.

Battigalli and Siniscalchi (2002) introduces the notion of "strong belief" to provide an epistemic foundation of solution concepts with a forward-induction flavor: a CPS μ (or a type t_i with such CPS) strongly believes an event E if $\mu(E|C) = 1$ for each conditioning event C that does not contradict E. Battigalli and Siniscalchi (2002) shows that, in a complete compact and continuous type structure based on a **finite** game Γ (hence, in the canonical structure based on Γ), strong rationalizability[1] characterizes the behavioral implications of rationality and common strong belief in rationality (RCSBR). The solution can be computed by maximal iterated deletion of conditionally dominated strategies (Shimoji and Watson 1998). In generic finite games with complete and perfect information, RCSBR implies the backward induction path, although not (always) the backward induction strategies (see Battigalli and Siniscalchi 2002). Furthermore, Battigalli and Siniscalchi (2002, 2003, 2007) show that similar epistemic assumptions that take as transparent some restrictions on players' beliefs justify the iterated intuitive criterion in signaling games, and similar equilibrium refinements in more general games.[2]

This leaves a gap between the class of games for which BS could provide epistemic justifications of solution concepts—i.e., finite games—and the class of simple dynamic games for which a complete compact and continuous type structure can be constructed, as in Battigalli and Siniscalchi (1999). We point out that many interesting applications of game theory consist of such simple, but **infinite** dynamic games. Here, we fill this gap. As a preliminary step, we prove a generalization of Lubin's (1974) extension result[3] (originally given for probability measures, here applied to CPSs): for every decreasing sequence of events $E_1, E_2, \ldots$ in the Polish state space $X \times T$ and every CPS ν on X that strongly believes each $\text{proj}_X E_m$, there exists a CPS μ on $X \times T$ that strongly believes each E_m and such that ν is the marginal of μ. This technical lemma

[1] In finite games of complete information without chance moves, strong rationalizability coincides with (the correlated version of) "extensive-form rationalizability," a solution concept put forward by Pearce (1984). Battigalli (2003) coined the term "strong rationalizability" to distinguish it from other legitimate, but weaker notions of rationalizability for dynamic games (e.g., Ben Porath 1997). For more on rationalizability in dynamic games, see the references therein.

[2] Of course, the general epistemic framework of BS allows for the analysis of different epistemic assumptions not involving forward-induction reasoning, as illustrated in Battigalli and Siniscalchi (1999).

[3] Actually, the extension result of Lubin (1974) follows from Lemma 2.2 of Varadarajan (1963).

allows us to provide streamlined proofs of results in epistemic game theory for simple dynamic games. In particular, we prove the following:

1.1 Characterization result

Fix a simple dynamic game Γ with complete or incomplete information and, for each player i, a closed set Δ_i of first-order CPSs representing some given transparent restrictions of i's beliefs; the profile $\Delta = (\Delta_i)$ of restricted sets of first-order beliefs gives raise to a solution procedure called strong Δ-rationalizability.[4] Let $[\Delta]$ denote the event that such belief restrictions hold; then, the behavioral implications of m-mutual strong belief in rationality and $[\Delta]$ are characterized by $m + 1$ steps of strong Δ-rationalizability, and the behavioral implications of common strong belief in rationality and $[\Delta]$ are characterized by strong Δ-rationalizability.

1.2 Related literature

We refer to Dekel and Siniscalchi (2015) for an up-to-date rigorous survey on epistemic game theory that adopts the approach of this paper. Here, we focus on earlier characterization results similar to the one stated above. We already cited most of the literature on dynamic epistemic game theory we build upon. The first characterization result for type structures where beliefs are CPSs is proved by Ben Porath (1997): in generic finite games with complete and perfect information, a strategy is consistent with rationality and common **initial** belief in rationality if and only if it survives one round of (maximal) deletion of weakly dominated strategies followed by the iterated deletion of strictly dominated strategies. Since initial belief is a monotone operator, Ben Porath can prove his result restricting attention to finite (hence, incomplete) type structures. As argued by Battigalli and Siniscalchi (2002), this standard line of proof cannot be applied to epistemic assumptions featuring strong belief, which is not monotone. Tan and Werlang (1988) consider the canonical type structure based on any compact continuous game with simultaneous moves and state that rationalizability characterizes the behavioral implications of rationality and common belief in rationality; but they prove only one direction of the characterization result, the one from the epistemic assumptions to the solution concept. The other—more difficult—direction is proved by Arieli (2010) for continuous games with simultaneous moves and Polish action spaces. Arieli's proof relies on Lubin's extension lemma. Our proof is inspired by this. Chen et al. (2016) analyzes the behavioral implications of rationality and common belief in rationality in arbitrary static games under general preferences.

The rest of the paper is structured as follows. Section 2 introduces preliminary mathematical concepts and results. Section 3 presents our generalization of Lubin's extension result for conditional probability systems and a sketch of proof. Section 4 presents our epistemic analysis of simple dynamic games, a short proof of the main game theoretic results (Theorem 1 and Corollary 1), and two illustrative economic

[4] See Battigalli (2003, 2006) and Battigalli and Siniscalchi (2003). Battigalli and Siniscalchi (2007), Battigalli and Friedenberg (2012), and Battigalli and Prestipino (2013) relate strong Δ-rationalizability to the epistemic analysis of forward-induction reasoning.

applications. Section 5 discusses some extensions and generalizations. All proofs but those of Theorem 1 and Corollary 1 are collected in 'Appendix".

2 Preliminaries

Following Battigalli and Siniscalchi (1999), we first carry out a preliminary analysis within an abstract probabilistic framework. Consider an individual who faces *primitive uncertainty* about the true state x in some Polish space X. Given any topological space Y (possibly, but not necessarily $Y = X$), we denote by $\mathcal{B}(Y)$ its Borel sigma-algebra, whose elements are called *events* (in Y). We endow the primitive uncertainty space X with a countable collection $\mathcal{C} \subseteq \mathcal{B}(X)$ of Borel subsets of X, which are the observables, or *conditioning* events.

Besides the primitive uncertainty space, we also posit a Polish space T to be interpreted as the set of possible hierarchies of beliefs, or types, or profiles of types, and we regard the Polish space $X \times T$ as the relevant *uncertainty space*. From now on, we reserve the name *event* for the Borel subsets of $X \times T$. We assume that the individual can only obtain direct information about X, that is, the set of observable events is the collection of "cylinders" of the form $C \times T$ for some C in $\mathcal{C}$, denoted by $\mathcal{C} \times T$. There is an obvious identification of the elements of $\mathcal{C}$ with the elements of $\mathcal{C} \times T$, and they are all referred to as conditioning events. By the same identification, we can write, with a small abuse of notation, $(X \times T, \mathcal{C})$ for the pair $(X \times T, \mathcal{C} \times T)$. Furthermore, $(X, \mathcal{C})$ can be identified with $(X \times T, \mathcal{C} \times T)$ if T is a singleton.

2.1 Conditional probability systems and strong belief

For every topological space Y, let $\Delta(Y)$ denote the set of probability measures on $\mathcal{B}(Y)$. We consider arrays of probability measures on Y indexed by elements of a collection $\mathcal{C} \subseteq \mathcal{B}(Y)$: i.e., $\mu = (\mu_C)_{C\in\mathcal{C}} \in [\Delta(Y)]^{\mathcal{C}}$. When we want to stress the interpretation of μ_C as a conditional probability given C, we write it as $\mu(\cdot|C)$.

Definition 1 Let Y be a Polish space and $\mathcal{C}$ be a countable collection of Borel subsets of Y. A *conditional probability system* (*CPS*) on $(Y, \mathcal{C})$ is an array of probability measures $(\mu_C)_{C\in\mathcal{C}}$ in $[\Delta(Y)]^{\mathcal{C}}$ such that, for all E in $\mathcal{B}(Y)$, and C, D in $\mathcal{C}$, $\mu_C(C) = 1$ and

$$E \subseteq D \subseteq C \Rightarrow \mu_C(E) = \mu_D(E)\,\mu_C(D). \tag{1}$$

Using the conditional probability notation, (1) can be written as follows: if $E \subseteq D \subseteq C$ then

$$\mu\,(D|C) > 0 \Rightarrow \mu\,(E|D) = \frac{\mu\,(E|C)}{\mu\,(D|C)}.$$

Recall the setting introduced in the previous subsection, where conditioning events $C \subseteq X$ concern primitive uncertainty and correspond to "cylinders" $C \times T$ in the overall uncertainty space $Y = X \times T$. Then, beliefs of the players are given by CPSs on $(X \times T, \mathcal{C})$:

Remark 1 Let $\mathcal{C} \subseteq \mathcal{B}(X)$ be countable; a CPS on $(X \times T, \mathcal{C})$ is an element of $[\Delta(X \times T)]^{\mathcal{C}}$, say μ, that satisfies the following conditions: $\mu_C(C \times T) = 1$, and

$$\mu_C(E) = \mu_D(E)\,\mu_C(D \times T)$$

for all E in $\mathcal{B}(X \times T)$ and C, D in $\mathcal{C}$ such that $E \subseteq D \times T \subseteq C \times T$.

The set $\Delta(X \times T)$ is endowed with the topology of weak convergence of measures. If X and T are Polish, then also $\Delta(X \times T)$ is a Polish space (see Aliprantis and Border 2006, Theorem 15.15), and so is $[\Delta(X \times T)]^{\mathcal{C}}$ because $\mathcal{C}$ is countable. By Remark 1, the set of CPSs on $(X \times T, \mathcal{C})$ is a subset of $[\Delta(X \times T)]^{\mathcal{C}}$ and is denoted by $\Delta^{\mathcal{C}}(X \times T)$.

Lemma 1 (Battigalli and Siniscalchi 1999) *If all the elements of the countable collection $\mathcal{C}$ are clopen (closed and open), the set of CPSs $\Delta^{\mathcal{C}}(X \times T)$ is closed in $[\Delta(X \times T)]^{\mathcal{C}}$. Therefore, $\Delta^{\mathcal{C}}(X \times T)$ is a Polish space.*

Given spaces Y and Q, for every probability measure ν on their Cartesian product $Y \times Q$, we let $\mathrm{mrg}_Y\, \nu$ denote the *marginal* of ν on Y. If $Q = R \times W$, it makes sense to consider the marginal on $Y \times R$ of a CPS μ on $(Y \times Q, \mathcal{C})$.

Definition 2 Let Y be a Polish space and $\mathcal{C}$ a countable subset of $\mathcal{B}(Y)$, let $Q = R \times W$, where R and W are Polish spaces. For any $\mu = (\mu_C)_{C \in \mathcal{C}}$ in $[\Delta(Y \times Q)]^{\mathcal{C}}$, the *marginal* of μ on $Y \times R$, denoted by $\mathrm{mrg}_{Y \times R}\, \mu$, is the array of probability measures $\nu = (\nu_C)_{C \in \mathcal{C}}$ in $[\Delta(Y \times R)]^{\mathcal{C}}$ such that $\nu_C = \mathrm{mrg}_{Y \times R}\, \mu_C$ for each C in $\mathcal{C}$.

Remark 2 Under the assumptions of Definition 2, we have $\mathrm{mrg}_{Y \times R}\, \mu \in \Delta^{\mathcal{C}}(Y \times R)$ for every CPS μ on $(Y \times R \times W, \mathcal{C})$.

We conclude this paragraph introducing the notion of strong belief, which is central to the epistemic analysis of dynamic games.

Definition 3 Given an event E in $X \times T$ and an array of probability measures $\mu = (\mu_C)_{C \in \mathcal{C}}$ in $[\Delta(X \times T)]^{\mathcal{C}}$, we say that μ *strongly believes* E if, for every C in $\mathcal{C}$,

$$E \cap (C \times T) \neq \emptyset \Rightarrow \mu_C(E) = 1.$$

Let $\mathcal{E}$ be a decreasing (finite or infinite) sequence of events in $X \times T$. We say that μ in $[\Delta(X \times T)]^{\mathcal{C}}$ *strongly believes* $\mathcal{E}$ if μ strongly believes each element of $\mathcal{E}$.

Simple Conditional Spaces. For some dynamic games, the information structure satisfies nice properties that allow for a simpler analysis. In line with our momentary abstraction from the game theoretic setting, we introduce these properties with the following definition.

Definition 4 A pair $(Y, \mathcal{C})$ with $\mathcal{C} \subseteq \mathcal{B}(Y)$ is a *simple conditional space* if Y is a compact metrizable space, $\mathcal{C}$ is a countable collection of clopen subsets of Y, and

(i) $Y \in \mathcal{C}$;

(ii) for all C', C'', D in $\mathcal{C}$, if $C' \supseteq D$ and $C'' \supseteq D$, then either $C' \supseteq C''$ or $C'' \supseteq C'$;
(iii) for all D in $\mathcal{C}$, the set $\{C \in \mathcal{C} : C \supseteq D\}$ is finite.

If we interpret the elements of $\mathcal{C}$ as "nodes," $C \supseteq D$ can be read as "node C weakly precedes node D," and each node in $(\mathcal{C}, \supseteq)$ has a finite number of predecessors, including Y. Thus, we refer to $\supseteq$ as a (weak) precedence relation and we can regard $(\mathcal{C}, \supseteq)$ as a tree with distinguished root Y.

Notice that if $(X, \mathcal{C})$ is a simple conditional space and T is a compact metric space, then the pair $(X \times T, \mathcal{C} \times T)$ is also a simple conditional space.[5] The following example illustrates the relevance of this concept.

Example 1 The elements of $X = A^{\mathbb{N}}$ are the paths (infinite histories) of an infinitely repeated game with perfect monitoring and a finite set of action profiles A in the stage game. The set A is endowed with the discrete metric d_A (i.e., $d_A(x, y) = 1$ if $x \neq y$, $d_A(x, y) = 0$ if $x = y$), and $A^{\mathbb{N}}$ is endowed with the "discounting metric," defined for $x = (a_n)_{n \in \mathbb{N}}$ and $x' = (a'_n)_{n \in \mathbb{N}}$ as

$$d(x, x') = \sum_{n \in \mathbb{N}} \frac{1}{2^n} d_A\left(a_n, a'_n\right).$$

With this, X is a compact metric space. Let $H = A^{<\mathbb{N}_0}$ be the set of finite histories (including the empty history) and let $\preceq$ denote the "(weak) prefix of" precedence relation on $H \cup X$. In particular, given h in H and x in X, write $h \preceq x$ if there exists x' in X such that $x = \left(h, x'\right)$, i.e., x is the concatenation of h and x'. Then, each cylinder $X(h) = \{x \in X : h \preceq x\}$ is clopen, H is countable, and the class of cylinders

$$\mathcal{C} = \{C \in 2^X : \exists h \in H, C = X(h))\}$$

has a tree-like structure. Therefore, $(X, \mathcal{C})$ is a simple conditional space, where each node in $\mathcal{C}$ is interpreted as the event that a given finite sequence of action profiles is observed. ▲

3 A generalization of Lubin's lemma

For any Cartesian product $X \times T$ and subset $E \subseteq X \times T$, we let $\mathrm{proj}_X\, E$ denote the projection of E onto X, that is,

$$\mathrm{proj}_X\, E = \{x \in X : \exists t \in T, (x, t) \in E)\}.$$

Lubin (1974) proved an "extension" result for probability measures, essentially stating the existence of a probability measure concentrated on a given subset of a product space and with a given marginal. This is the content of the following lemma, which adapts Lubin's result to our simple topological assumptions.

[5] Note that if C is clopen, then so is $C \times T$.

Lemma 2 (Lubin 1974) *Let X and T be compact metrizable spaces and let $E \subseteq X \times T$ be closed. For each ν in $\Delta(X)$ such that $\nu(\text{proj}_X E) = 1$, there exists some μ in $\Delta(X \times T)$ such that $\mu(E) = 1$ and $\text{mrg}_X \mu = \nu$.*

We extend this result in two ways: first, we show that it holds for CPSs; second, while the above result just looks at a measure that assigns probability one to an event, we look at a CPS that strongly believes a chain (decreasing sequence) of events.

Lemma 3 *Let $(X, \mathcal{C})$ be a simple conditional space and let T be compact metrizable. Fix a decreasing sequence of closed events $\mathcal{E} = (E_1, \ldots, E_n)$ in $X \times T$. For each CPS ν on $(X, \mathcal{C})$ that strongly believes $(\text{proj}_X E_1, \ldots, \text{proj}_X E_n)$, there is a CPS μ on $(X \times T, \mathcal{C})$ that strongly believes $\mathcal{E}$ and such that $\text{mrg}_X \mu = \nu$.*

The proof is in "Appendix," here we give only a sketch. Starting from a CPS ν on $(X, \mathcal{C})$ that strongly believes $(\text{proj}_X E_1, \ldots, \text{proj}_X E_n)$, Lemma 2 applied to each measure ν_C gives an array of probability measures $\bar{\mu} = (\bar{\mu}_C)_{C \in \mathcal{C}}$ in $[\Delta(X \times T)]^{\mathcal{C}}$ such that (i) $\text{mrg}_X \bar{\mu} = \nu$, $\bar{\mu}_C(C \times T) = 1$ for all C in $\mathcal{C}$, and (ii) $\bar{\mu}$ strongly believes $\mathcal{E}$. However, $\bar{\mu}$ need not be a CPS because it may fail to satisfy the chain rule (1). To obtain the required CPS μ, we need to carefully select some of the measures in $\bar{\mu}$ and construct μ from such measures so that the chain rule is satisfied. Certainly, we want to derive μ_C from the initial belief $\bar{\mu}_X$ for each conditioning event C with $\bar{\mu}_X(C \times T) > 0$ (i.e., $\nu_X(C) > 0$): for every event E in $X \times T$,

$$\mu_C(E) = \frac{\bar{\mu}_X(E \cap (C \times T))}{\bar{\mu}_X(C \times T)}.$$

Now, think of $\mathcal{C} \times T$ as a tree with root $X \times T$, where $C \times T$ (weakly) precedes $D \times T$ if and only if $C \times T \supseteq D \times T$ (i.e., $C \supseteq D$). For each maximal (hence possibly infinite) chain of events $\mathcal{D} \subseteq \mathcal{C} \times T$ starting with the root $X \times T$, pick the first $C \times T$ in $\mathcal{D}$ such that $\bar{\mu}_X(C \times T) = 0$ (if it exists) and derive μ_D from $\bar{\mu}_C$ for every $D \times T$ in $\mathcal{D}$ with $\bar{\mu}_C(D \times T) > 0$: for every event E in $X \times T$,

$$\mu_D(E) = \frac{\bar{\mu}_C(E \cap (D \times T))}{\bar{\mu}_C(D \times T)}.$$

Next, for each maximal chain $\mathcal{D}' \subseteq \mathcal{C} \times T$ starting with $C \times T$, pick the first $D \times T$ in $\mathcal{D}'$ such that $\bar{\mu}_C(D \times T) = 0$ (if it exists) and derive μ_E from $\bar{\mu}_D$ for every $E \times T$ in $\mathcal{D}'$ such that $\bar{\mu}_D(E \times T) > 0$. Proceeding in this way, one can assign a probability measure concentrated on $C \times T$ for each conditioning event C in $\mathcal{C}$. The construction implies that the chain rule (1) holds and that μ strongly believes $\mathcal{E}$.

4 Interactive epistemology in infinite dynamic games

We consider a class of dynamic games, with finite or infinite horizon, in which each player has a finite set of feasible actions at each stage except (possibly) the last one, where they are only required to have a compact set of feasible actions (see Battigalli

2003). The length of the game, the set of active players and their sets of feasible actions may be history dependent. We allow for incomplete information about payoff functions, but we assume for simplicity that there is perfect monitoring of past actions and that the feasibility constraints do not depend on private information. For such games, we characterize the behavioral implications of rationality and common strong belief in rationality.

4.1 Simple infinite dynamic games

Our definition of game takes as primitive terms the actions and private information of each player, e.g., the output and cost function of a firm; we refer to the latter as "information type."[6] Histories are sequences of action profiles. With this, nodes of the game tree are constructed as pairs of states of nature and histories. We use the following notation on sequences: for any range A, $A^{<\mathbb{N}_0} = \bigcup_{n \in \mathbb{N}_0} A^n$ is the set of sequences $(a^1, \ldots, a^n)$ of finite (but otherwise arbitrary) length.[7] In particular, the set $A^{<\mathbb{N}_0}$ contains the empty sequence $\varnothing$, which is the only element of A^0 (i.e., it is the unique sequence of length 0), and every sequence (a) of length 1, which we simply identify with its only element a; $A^{\mathbb{N}}$ is the set of infinite sequences $(a^n)_{n \in \mathbb{N}}$, that is, the set of functions from $\mathbb{N}$ to A, and $A^{\leq \mathbb{N}_0} = A^{<\mathbb{N}_0} \cup A^{\mathbb{N}}$. For all h' in $A^{<\mathbb{N}_0}$ and h'' in $A^{\leq \mathbb{N}_0}$, (h', h'') denotes the concatenation of h' with h'' (in particular, $h = (h, h') = (h', h)$ if $h' = \varnothing$), and we say that h' is a *prefix* of h, written $h' \preceq h$, if $h = (h', h'')$ for some h''. With this, the empty sequence is a prefix of every sequence. A nonempty set $\bar{H} \subseteq A^{\leq \mathbb{N}_0}$ of finite or infinite sequences in A is a *tree* if it is closed with respect to prefixes, that is, for every h in $\bar{H}$ and every prefix h' of h, $h' \in \bar{H}$; in particular, $\varnothing \in \bar{H}$.

Fix a tree $\bar{H} \subseteq A^{\leq \mathbb{N}_0}$; a sequence in $\bar{H}$ is *terminal* if it is not the prefix of any other nonempty sequence in $\bar{H}$, and it is *nonterminal* otherwise. Thus, terminal sequences may have infinite length and nonterminal sequences are necessarily finite. For each finite sequence h in $\bar{H}$, we let $\mathcal{A}(h) = \{a \in A : (h, a) \in \bar{H}\}$. The sets of nonterminal and terminal sequences are, respectively, $H = \{h \in \bar{H} \cap A^{<\mathbb{N}_0} : \mathcal{A}(h) \neq \emptyset\}$ and $Z = \bar{H} \setminus H$. A nonterminal sequence h in tree $\bar{H}$ is *preterminal* if $(h, a) \in Z$ for *all* a in $\mathcal{A}(h)$.

Definition 5 A *simple dynamic game structure* is a tuple

$$\langle I, \bar{H}, (\Theta_i, A_i)_{i \in I} \rangle$$

where I is a countable set of *players*, Θ_i is a nonempty compact metrizable space of *information types*, each A_i is a nonempty compact metrizable space of *actions* $(i \in I)$, $\bar{H}$ is a tree of sequences of elements of $A = \prod_{i \in I} A_i$, called *histories*, such

[6] Information types may be different from types in the sense of Harsanyi (1967–1968), that implicitly determine hierarchies of exogenous initial beliefs, and from the epistemic types introduced later, that represent hierarchies of conditional probability systems.

[7] Note that A^n is the set of all sequences of elements of A with length n. In the rest of the paper, it will be clear from the context whether a superscript stands for a product (as in this case) or as a mere index (as, for example, in the statement of Lemma 3).

that, for every nonterminal history h in H, (i) $\mathcal{A}(h) = \prod_{i \in I} \text{proj}_{A_i} \mathcal{A}(h)$, (ii) if h is not preterminal, $\mathcal{A}(h)$ is finite, (iii) if h is preterminal, $\mathcal{A}(h)$ is compact.

Condition (i) means that what is feasible for a player cannot depend on actions simultaneously chosen by other players. Thus, $\mathcal{A}_i(h) = \text{proj}_{A_i} \mathcal{A}(h)$ is the (nonempty) set of feasible actions for player i given history h in H. Player i is *active* at history h if $|\mathcal{A}_i(h)| > 1$ and *inactive* otherwise (that is, if $|\mathcal{A}_i(h)| = 1$). An information type of player i, $\theta_i \in \Theta_i$, describes the private information of i: for example, θ_i may be a preference, or a productivity parameter. It is informally assumed that there are no other sources of private information; in particular, the actions chosen in previous stages are public information. In other words, we consider multistage game structures *with observable actions*. We make this assumption only for notational simplicity. The fact that $\mathcal{A}_i$ does not depend on θ_i means that, at each stage, the history-dependent set of feasible actions is common knowledge. It is easy to extend our results to the case where $\mathcal{A}_i$ depends on a discrete component of θ_i. We can also add an explicit common source of uncertainty θ_0, which is a payoff-relevant parameter unknown to all the players; this is conceptually useful for some purposes, but not here. Similarly, we could easily add chance moves.[8] We avoid these generalizations only to simplify some aspects of the mathematical notation.[9] The game tree can be derived as follows: let $\Theta = \prod_{i \in I} \Theta_i$ be the set of *states of nature*. The nonterminal nodes of the game tree are the pairs (θ, h) in $\Theta \times H$, the terminal nodes are pairs in $\Theta \times Z$; node (θ, h) precedes node (θ', h') if $\theta = \theta'$ and $h \preceq h'$; the information sets of player i have the form

$$\{(\theta', h') \in \Theta \times H : \theta'_i = \theta_i,\ h' = h\},$$

for θ_i in Θ_i and h in H (where in the expression above, θ'_i denotes the ith component of the profile θ'). In words, all the nodes featuring the same information type of i and the same history are indistinguishable for i.

We let payoffs be functions of the terminal nodes (θ, z) in $\Theta \times Z$ and informally assume that such functions are common knowledge. All uncertainty and private information about payoffs is therefore captured by states of nature $\theta = (\theta_i)_{i \in I}$ with the understanding that each i knows his information type θ_i. With this, we consider simple conditional spaces determined by the game structure. Let the *strategies* of player i be given by elements of $S_i = \prod_{h \in H} \mathcal{A}_i(h)$. Note that we take an "interim perspective": i is "born" with knowledge of his information type θ_i, and s_i only describes how i behaves as a function of h. Player i does not know the types of the other players, nor what strategies they are implementing; thus, the space of primitive uncertainty for i is $\prod_{j \neq i}(\Theta_j \times S_j)$. To ease notation, we write $\Sigma_i = \Theta_i \times S_i$ and $\Sigma_{-i} = \prod_{j \neq i} \Sigma_j$. The relevant conditioning events are determined as follows: each strategy profile

[8] Such features can be implicitly modeled within the present framework by letting $0 \in I$ be an indifferent "fictitious" player. Then, θ_0 parameterizes residual uncertainty, about which no nonfictitious player has private information, and $\mathcal{A}_0(h)$ is the set of realizations of a chance move. Knowledge about the objective probabilities of chance moves can be modeled by restrictions on players' beliefs.

[9] Battigalli and Prestipino (2013) allow for imperfect monitoring, a dimension of common uncertainty, and chance moves, but they maintain that the game structure is finite.

$s = (s_i)_{i \in I}$ in $S = \prod_{i \in I} S_i$ determines a unique terminal history

$$s \mapsto \zeta(s) = (s(\varnothing), s(s(\varnothing)), s\left(s(\varnothing), s(s(\varnothing))\right), \ldots),$$

where $s(h) = (s_i(h))_{i \in I}$ is the action profile determined by s at h. The set of strategy profiles consistent with history h is

$$S(h) = \{s \in S : h \preceq \zeta(s)\}.$$

Note that[10]

$$S(h) = \{s \in S : \forall (h', a) \in H, (h', a) \preceq h \Rightarrow s(h') = a\} = \prod_{i \in I} \mathrm{proj}_{S_i} S(h).$$

Then, if h occurs, i is informed that the co-players are implementing strategies in subset $S_{-i}(h) = \mathrm{proj}_{S_{-i}} S(h)$. Since feasibility constraints do not depend on types, observed choices cannot reveal hard information about θ_{-i}. We can therefore represent i's information at h about the types of the co-players and the strategies they are implementing with the set $\Sigma_{-i}(h) = \Theta_{-i} \times S_{-i}(h)$. Note that if $a = (a_i, a_{-i})$ and $a' = (a_i', a_{-i})$ belong to $\mathcal{A}(h)$, then $\Sigma_{-i}(h, a) = \Sigma_{-i}(h, a')$. Thus, we implicitly assume that player i's beliefs about the strategies and types of others, conditional on a given history, are independent of his own actions. Let

$$\mathcal{C}_i = \{\Sigma_{-i}(h) : h \in H\}$$

denote the set of conditioning events for player i.

Lemma 4 *Fix a simple dynamic game structure* $\langle I, \bar{H}, (\Theta_i, A_i)_{i \in I}\rangle$ *and a player* i *in* I. *Then,* $(\Sigma_{-i}, \mathcal{C}_i)$ *is a simple conditional space.*

Adding to a game structure a map from terminal nodes to outcomes (or consequences), we obtain a game form, which is the mathematical representation of the rules of the game; adding to the game form players' type-dependent preferences over (lotteries of) outcomes, we obtain a game with incomplete information.[11]

Definition 6 A *simple dynamic game form* is a tuple

$$\langle I, \bar{H}, (\Theta_i, A_i)_{i \in I}, Y, g\rangle,$$

where $\langle I, \bar{H}, (\Theta_i, A_i)_{i \in I}\rangle$ is a simple dynamic game structure, Y is a compact metric *space of outcomes*, and $g : \Theta \times Z \to Y$ is a continuous *outcome function*. A *simple dynamic game* with incomplete information is a tuple

$$\Gamma = \langle I, \bar{H}, (\Theta_i, A_i)_{i \in I}, Y, g, (v_i)_{i \in I}\rangle,$$

[10] More generally, in all games with perfect recall, $S(h_i) = \mathrm{proj}_{S_i} S(h_i) \times \mathrm{proj}_{S_{-i}} S(h_i)$ for each information set h_i of each player i, which is all we really need.

[11] The intermediate step from game structure to game form could be skipped, defining utility/payoff functions over terminal nodes. We keep this step because it adds conceptual clarity at little cost.

where $\langle I, \bar{H}, (\Theta_i, A_i)_{i\in I}, Y, g\rangle$ is a simple dynamic game form and, for each player i, $v_i : \Theta_i \times Y \to \mathbb{R}$ is a continuous *utility function*; the composition

$$(\theta, z) \mapsto u_i(\theta, z) = v_i(\theta_i, g(\theta, z))$$

is called *payoff function* of player i.

Henceforth, we maintain the following:

Assumption 1 Γ is a simple dynamic game.

Comment: Structure Γ is also called "game with payoff uncertainty" (Battigalli and Siniscalchi 2007) or "basic game" (Bergemann and Morris 2016) to distinguish it from Bayesian games (Harsanyi 1967–1968), that is, richer structures implicitly describing players' infinite hierarchies of exogenous initial beliefs about θ. The epistemic structures considered in this paper implicitly describe players' infinite hierarchies of conditional beliefs based on $(\Theta \times S, H)$, from which one can derive infinite hierarchies of exogenous initial beliefs. But there is no need to refer to Bayesian games in our analysis.

Example 2 (*Disclosure, see* Battigalli 2006) Consider a buyer–seller signaling game. The seller (player 1) has private information about the quality $\theta \in \Theta_1 = \{1, \ldots, K\}$ of the object to sell. After some message $m \in A_1$ of the seller, the buyer (player 2) chooses a price $a_2 \in A_2 = [0, b]$ for the object, where the upper bound b is large: $b > K$. Partial and terminal histories are, respectively,

$$H = \{\varnothing\} \cup (A_1 \times \{w\}) \cong \{\varnothing\} \cup A_1, \quad Z = (A_1 \times \{w\}) \times (\{w\} \times A_2) \cong A_1 \times A_2,$$

where w denotes the pseudo-action of "waiting." To ease notation, let $Y = Z$ and $g = \mathrm{id}_Z$. The utility of the seller is strictly increasing in the action of the buyer, while the utility of the buyer is a quadratic loss function $u_2(\theta^*, m, a_2) = -(\theta^* - a_2)^2$. It is assumed for simplicity that the message of the sender has the form "quality is at least k," which we denote by $[\theta \geq k]$, for some $k \in \Theta_1$ (note that Θ_1 and A_1 are isomorphic in an obvious sense). Moreover, information is verifiable and the seller is harshly punished if his message conveys false information: for all θ^* in Θ_1, $[\theta \geq k]$ in A_1, and a_2 in A_2,

$$u_1(\theta^*, [\theta \geq k], a_2) = \begin{cases} a_2 & \text{if } \theta^* \geq k, \\ a_2 - P & \text{otherwise,} \end{cases}$$

where $P > b$. With this, it is dominant for the seller to tell the truth, that is, to send message $[\theta \geq k]$ given his true type θ^* only if $\theta^* \geq k$. The game has finite horizon, but it is infinite, like the set of actions of the buyer. Nonetheless, it is readily verified that it is a simple game. ▲

Example 3 (*Reputation, see* Schmidt 1993) Consider a finite, two-player stage game $G = \langle I = \{1, 2\}, A_1, A_2, v_1, v_2\rangle$ ($v_i : A_1 \times A_2 \to \mathbb{R}$, $i \in \{1, 2\}$). Assume for simplicity that each player has a unique best reply to each pure action of the co-player in G. Let G be infinitely repeated, with discount factors δ_1 and δ_2. Player 1 is more patient than player 2: $\delta_1 > \delta_2$. The partial histories are the finite sequences in $H = (A_1 \times A_2)^{<\mathbb{N}_0}$, and the set of terminal histories is $Z = (A_1 \times A_2)^{\mathbb{N}}$. The constraint correspondence is simply $\mathcal{A}_i(h) = A_i$, for all h in H and each i.

To ease notation, let $Y = Z$ and $g = \mathrm{id}_Z$. The game has private values, with Θ_2 a singleton and $\Theta_1 \cong \Theta = \{\theta^\circ, \theta^c, \ldots\}$. The payoff functions of types θ° and θ^c of player 1 and player 2 are given by

$$u_1(\theta^\circ, (a^t)_{t=1}^{\infty}) = (1-\delta_1)\sum_{t=1}^{\infty}\delta_1^{t-1}v_1(a^t), \qquad u_2((a^t)_{t=1}^{\infty}) = (1-\delta_2)\sum_{t=1}^{\infty}\delta_2^{t-1}v_2(a^t),$$
$$u_1(\theta^c, (a^t)_{t=1}^{\infty}) = (1-\delta_1)\sum_{t=1}^{\infty}\delta_1^{t-1}v_1^c(a^t),$$

where $v_1^c(a) = 1$ if $a = (a_1^*, a_2)$ for any $a_2 \in A_2$, and $v_1^c(a) = 0$ otherwise, with

$$a_1^* = \arg\max_{a_1 \in A_1} v_1\left(a_1, \arg\max_{a_2 \in A_2} v_2(a_1, a_2)\right).$$

In other words, player 1 can be either a normal (θ°), or a "crazy" type (θ^c); in the latter case, it is always dominant for him to choose the *Stackelberg action* a_1^*. The strategy s_1^* that plays a_1^* at each h is called *Stackelberg strategy.* Even if there are other types besides θ° and θ^c, they do not matter for our analysis.

The game just described is simple, as is every infinite repetition of a finite stage game. ▲

4.2 Epistemic characterization results

Following Battigalli and Siniscalchi (1999), one can construct a canonical epistemic type structure based on a simple dynamic game. Each epistemic type t_i is an infinite hierarchy of CPSs $t_i = (\mu_i^m)_{m\in\mathbb{N}}$ where μ_i^1 is the first-order CPS of type t_i, an element of $\Delta^{\mathcal{C}_i}(\Sigma_{-i})$, μ_i^2 is the second-order CPS of type t_i, an element of $\Delta^{\mathcal{C}_i}(\prod_{j\neq i}(\Sigma_j \times \Delta^{\mathcal{C}_j}(\Sigma_{-j})))$, so that $\mu_i^1 = \mathrm{mrg}_{\Sigma_{-i}}\,\mu_i^2$, and so on. Denoting by T_i the set of infinite hierarchies satisfying coherence and common full belief in coherence[12] (a compact metrizable space), there is a canonical homeomorphism $\beta_i : T_i \to \Delta^{\mathcal{C}_i}(\Sigma_{-i} \times T_{-i})$ that determines, for each epistemic type (hierarchy) t_i, a CPS $\beta_i(t_i)$ over information types, strategies, and epistemic types (hierarchies) of the other players.[13]

[12] Roughly, this means that everybody's lower-order beliefs are the marginals of higher-order beliefs and there is common belief of this conditional on each history.

[13] Furthermore, Battigalli and Siniscalchi (1999) prove that the canonical type structure is also universal, i.e., that every other type structure can be universally (hence uniquely up to isomorphism) mapped into it preserving beliefs hierarchies. On such belief-preserving morphisms, see Heifetz and Samet (1998) and Friedenberg and Meier (2011).

4.2.1 Complete type structures

Although this construction of the canonical structure is the conceptual backdrop of our epistemic analysis, we do not need to make it explicit for our formal results. In fact, it is enough to work with any type structure with belief maps that satisfy the continuity and surjectivity properties of the canonical structure.

Definition 7 Fix a simple dynamic game Γ. A *complete epistemic type structure* based on Γ is a tuple $\langle I, (T_i, \beta_i)_{i \in I} \rangle$ where, for each player i, T_i is a compact metrizable space of *epistemic types* and $\beta_i : T_i \to \Delta^{\mathcal{C}_i}(\Sigma_{-i} \times T_{-i})$ is a continuous surjective *belief map*.[14] A profile $(\theta_i, s_i, t_i)_{i \in I}$ in $\prod_{i \in I} \Sigma_i \times T_i$ is a *state of the world*, (θ_i, s_i, t_i) in $\Sigma_i \times T_i$ is a *personal state* of player i.

The consideration of incomplete type structures entails restrictions on players beliefs that are implicitly assumed to be transparent to the players, but not explicit in the theoretical analysis.[15] We therefore maintain the following:

Assumption 2 The given type structure $\langle I, (T_i, \beta_i)_{i \in I} \rangle$ is complete.

To ease notation, we let $\beta_{i,h}$ denote the component of β_i that corresponds to conditioning event $\Sigma_{-i}(h)$. Similarly, for each $\mu_i \in \Delta^{\mathcal{C}_i}(\Sigma_{-i} \times T_{-i})$ (respectively, $\mu_i \in \Delta^{\mathcal{C}_i}(\Sigma_{-i})$), we let $\mu_{i,h}$ denote the belief of i conditional on $\Sigma_{-i}(h)$. The topological assumptions of the definition above, together with Lemma 1, imply that $\Delta^{\mathcal{C}_i}(\Sigma_{-i} \times T_{-i})$ is a Polish space; therefore, it makes sense to assume that all the belief maps β_i (and thus all their components $\beta_{i,h}$) are continuous—a convenient technical property satisfied by the canonical structure. Completeness requires that for each CPS μ_i about the information types, strategies, and epistemic types of the co-players there is an epistemic type t_i with $\beta_i(t_i) = \mu_i$, that is, each conceivable system of beliefs is represented in the type structure. In particular, this implies that, for any history $h \in H$, **if** there is a profile of information types, strategies, and beliefs of the co-players $(\theta_{-i}, s_{-i}, \mu_{-i})$ such that the strategies s_{-i} are consistent with h and are best replies[16] to the beliefs μ_{-i} given θ_{-i}, **then** it is possible for i, upon observing h, to believe that the co-players are rational, because there is a profile of epistemic types t_{-i} such that $(\theta_{-i}, s_{-i}, t_{-i})$ satisfies rationality and is consistent with h. This is a key feature of the analysis of forward-induction reasoning of Battigalli and Siniscalchi (2002, 2007), which we extend here.[17]

4.2.2 Epistemic assumptions

Fix a simple dynamic game Γ and a complete type structure $\langle I, (T_i, \beta_i)_{i \in I} \rangle$ based on Γ, which we interpret as the canonical type structure for Γ. Our basic assumption is

---

[14] Each type t_i is also a type in the sense of Harsanyi (1967–1968), because it determines an initial belief about the information and epistemic types of the co-players: $t_i \mapsto \mathrm{mrg}_{\Theta_{-i} \times T_{-i}} \beta_{i,\varnothing}(t_i)$. Therefore, Γ and $\langle I, (T_i, \beta_i)_{i \in I} \rangle$ together yield a Bayesian game.

[15] See Battigalli and Friedenberg (2012), and Battigalli and Prestipino (2013).

[16] The formal definition of best reply is given below.

[17] Battigalli and Friedenberg (2012) analyze forward-induction reasoning when contextual assumptions rule out completeness.

that players are rational in the following sense. Let $S_i(h) = \text{proj}_{S_i} S(h)$ denote the set of strategies of i that do not prevent history h. Also, let

$$\begin{array}{llll} U_i: & \Theta \times S & \rightarrow & \mathbb{R} \\ & (\theta, s) & \mapsto & u_i(\theta, \zeta(s)) \end{array}$$

denote the strategic-form payoff function of i.

Definition 8 Strategy s_i^* is a *sequential best reply* to CPS μ_i in $\Delta^{\mathcal{C}_i}(\Sigma_{-i})$ given θ_i if, for every h in H such that $s_i^* \in S_i(h)$,

$$s_i^* \in \arg\max_{s_i \in S_i(h)} \int U_i(\theta_i, s_i, \sigma_{-i}) \mu_i(\mathrm{d}\sigma_{-i} | \Sigma_{-i}(h)).$$

Player i is *rational* at personal state (θ_i, s_i, t_i) if s_i is a sequential best reply to $\text{mrg}_{\Sigma_{-i}} \beta_i(t_i)$ given θ_i.

Comment: By standard arguments, s_i is a sequential best reply to μ_i given θ_i if and only if it is realization-equivalent to a strategy with the one-shot-deviation property given (μ_i, θ_i); that is, if and only if there exists some s_i^* such that

$$\forall h \in H, \quad s_i^*(h) \in \arg\max_{a_i \in \mathcal{A}_i(h)} \int U_i(\theta_i, s_i^*|_h a_i, \sigma_{-i}) \mu_i(\mathrm{d}\sigma_{-i} | \Sigma_{-i}(h)),$$

and

$$\forall s_{-i} \in S_{-i}, \quad \zeta(s_i^*, s_{-i}) = \zeta(s_i, s_{-i}),$$

where $s_i^*|_h a_i$ is the strategy consistent with h that chooses a_i at h and coincides with s_i^* at each history that does not weakly precede h (hence, at all histories following h). In simple dynamic games with finite horizon, this is equivalent to the standard "folding-back" rationality of dynamic programming. If players' preferences were dynamically inconsistent, then the conceptually appropriate best reply condition should be the one-shot deviation property.

To ease notation, we let $r_i(\mu_i, \theta_i)$ denote the set of sequential best replies to μ_i given θ_i. Then, we can define the event "player i is rational" as

$$R_i = \{(\theta_i, s_i, t_i) \in \Sigma_i \times T_i : s_i \in r_i(\text{mrg}_{\Sigma_{-i}} \beta_i(t), \theta_i)\}.$$

Lemma 5 *For each $i \in I$, the sequential best reply correspondence*

$$\begin{array}{llll} r_i: & \Delta^{\mathcal{C}_i}(\Sigma_{-i}) \times \Theta_i & \rightrightarrows & S_i \\ & (\mu_i, \theta_i) & \mapsto & r_i(\mu_i, \theta_i) \end{array}$$

is upper hemicontinuous with nonempty values, and R_i is nonempty and closed.

As argued by Battigalli (2003) and Battigalli and Siniscalchi (2003), it is often plausible in applications to assume that players' first-order beliefs satisfy some restrictions and that this is transparent. Such restrictions may depend on the information type. For example, we may want to assume that beliefs satisfy independence across co-players or that beliefs about θ_{-i} given θ_i are derived from a common prior on Θ. Battigalli (2003, 2006) and Battigalli and Siniscalchi (2003) provide many other examples and applications.

Definition 9 A system of first-order *belief restrictions* is a profile of closed-graph correspondences $\Delta = (\Delta_i)_{i\in I}$, where $\Delta_i : \Theta_i \rightrightarrows \Delta^{\mathcal{C}_i}(\Sigma_{-i})$ and Δ_i has nonempty values for each i in I.

We let $\Delta_{i,\theta_i} \subseteq \Delta^{\mathcal{C}_i}(\Sigma_{-i})$ denote the resulting nonempty closed set of CPSs for information type θ_i, and we let $[\Delta_i]$ denote the event in $\Sigma_i \times T_i$ that i satisfies the restrictions Δ_i, that is,

$$[\Delta_i] = \{(\theta_i, s_i, t_i) : \mathrm{mrg}_{\Sigma_{-i}}\, \beta_i(t_i) \in \Delta_{i,\theta_i}\}.$$

Remark 3 Since the graph of Δ_i is closed and the map

$$(\theta_i, s_i, t_i) \mapsto \left(\theta_i, s_i, \mathrm{mrg}_{\Sigma_{-i}}\, \beta_i(t_i)\right)$$

is continuous, $[\Delta_i]$ is closed.

Battigalli and Siniscalchi (2007) consider two systems of epistemic assumptions: common initial belief in rationality and in the belief restrictions, and common strong belief in rationality and in the belief restrictions. Here, we focus on the latter, which models a form of forward-induction reasoning. The analysis of the former is simpler and can be carried out adapting the techniques of this paper (see Sect. 5). For each event $E_{-i} \subseteq \Sigma_{-i} \times T_{-i}$, define $\mathrm{SB}_i(E_{-i}) \subseteq \Sigma_i \times T_i$ as the event that i *strongly believes* E_{-i}:

$$\mathrm{SB}_i(E_{-i}) = \Sigma_i \times \{t_i \in T_i : \forall h \in H, (E_{-i} \cap (\Sigma_{-i}(h) \times T_{-i}) \neq \emptyset) \Rightarrow \beta_{i,h}(t_i)(E_{-i}) = 1\}.$$

In words, i strongly believes E_{-i} if he is certain of E_{-i} conditional on every history h that does not contradict E_{-i}. Since $\Sigma_{-i}(\varnothing) = \Sigma_{-i}$, strong belief implies belief at the beginning of the game.

Remark 4 If E_{-i} is closed, the set

$$M_i(E_{-i}) = \bigcap_{h\in H:(\Sigma_{-i}(h)\times T_{-i})\neq\emptyset} \left\{\mu_i \in \Delta^{\mathcal{C}_i}(\Sigma_{-i} \times T_{-i}) : \mu_{i,h}(E_{-i}) = 1\right\}$$

of CPSs μ_i that strongly believe E_{-i} is closed; since β_i is continuous,

$$\mathrm{SB}_i(E_{-i}) = \Sigma_i \times \beta_i^{-1}(M_i(E_{-i}))$$

is closed as well.

Consider the following sequence of events: for each player i,

- $R_i^{\Delta,1} = R_i \cap [\Delta_i]$;
- for each positive integer m, $R_i^{\Delta,m+1} = R_i^{\Delta,m} \cap \mathrm{SB}_i(\prod_{j\neq i} R_j^{\Delta,m})$.

Note that the sequence is well defined, because—by an easy induction argument—Lemma 5 and Remarks 3–4 imply that each set $R_i^{\Delta,m}$ is closed. To ease notation, let $R_i^{\Delta,\infty} = \bigcap_{m\in\mathbb{N}} R_i^{\Delta,m}$.

The sequence of events $\left(\prod_{i\in I} R_i^{\Delta,m}\right)_{m\in\mathbb{N}}$ embodies the "best rationalization principle": If $(\theta_i, s_i, t_i) \in R_i^{\Delta,\infty}$, then player i is rational at (θ_i, s_i, t_i) and $\beta_i(t_i)$ strongly believes each event in the decreasing sequence $\left(R_{-i}^{\Delta,m}\right)_{m\in\mathbb{N}\cup\{\infty\}}$; this means that player i always believes in the highest level $m \in \mathbb{N}\cup\{\infty\}$ of "strategic sophistication" consistent with what he observes.

Definition 10 Player i satisfies Δ-rationality and mutual strong belief in Δ-rationality of order $m-1$ at personal state (θ_i, s_i, t_i) if $(\theta_i, s_i, t_i) \in R_i^{\Delta,m}$; player i satisfies Δ-rationality and common strong belief in Δ-rationality at personal state (θ_i, s_i, t_i) if $(\theta_i, s_i, t_i) \in R_i^{\Delta,\infty}$.

4.2.3 Characterization theorem

We are interested in the behavioral implications of the foregoing epistemic assumptions. In particular, for each information type θ_i, the section at θ_i of $\mathrm{proj}_{\Sigma_i} R_i^{\Delta,m}$ is the set of strategies consistent with Δ-rationality and mutual strong belief in Δ-rationality of order $m-1$ given θ_i; similarly, the section at θ_i of $\mathrm{proj}_{\Sigma_i} R_i^{\Delta,\infty}$ is the set of strategies consistent with Δ-rationality and common strong belief in Δ-rationality given θ_i. The main result of this paper yields a characterization of these behavioral implications by means of a solution concept, called "strong Δ-rationalizability," related to Pearce's 1984 notion of rationalizability for finite games in extensive form.

Definition 11 (cf. Battigalli 2003) Fix a system of belief restrictions Δ. For each player i, let

$$\Sigma_i^{\Delta,1} = \{(\theta_i, s_i) : \exists \mu_i \in \Delta_{i,\theta_i}, s_i \in r_i(\mu_i, \theta_i)\};$$

then recursively define, for each i in I and for each positive integer m,

$$\Sigma_i^{\Delta,m+1} = \Big\{(\theta_i, s_i) : \exists \mu_i \in \Delta_{i,\theta_i}, s_i \in r_i(\mu_i, \theta_i), \forall k \in \{1,\dots,m\}, \forall h \in H$$
$$\Sigma_{-i}^{\Delta,k} \cap \Sigma_{-i}(h) \neq \emptyset \Rightarrow \mu_i\left(\Sigma_{-i}^{\Delta,k} | \Sigma_{-i}(h)\right) = 1\Big\},$$

where $\Sigma_{-i}^{\Delta,k} = \prod_{j\neq i} \Sigma_j^{\Delta,k}$; finally, let $\Sigma_i^{\Delta,\infty} = \bigcap_{m\in\mathbb{N}} \Sigma_i^{\Delta,m}$. Strategy s_i is *m-strongly Δ-rationalizable* for information type θ_i if $(\theta_i, s_i) \in \Sigma_i^{\Delta,m}$; strategy s_i is *strongly Δ-rationalizable* for θ_i if $(\theta_i, s_i) \in \Sigma_i^{\Delta,\infty}$.

In words, a strategy is strongly Δ-rationalizable for information type θ_i if it can be justified as a sequential best reply to a first-order CPS μ_i that strongly believes each set $\Sigma_{-i}^{\Delta,k}$ of k-strongly Δ-rationalizable profiles of information types and strategies of the co-players. On top of the applications mentioned—for example—by Battigalli (2003), Battigalli (2006), and Battigalli and Siniscalchi (2003), this solution concept is used in the analysis of self-enforcing agreements[18] and of robust implementation/mechanism design.[19]

Comment: Battigalli and Prestipino (2013) show that, if the belief restrictions Δ are "closed under compositions," the solution algorithm can be simplified, requiring at each step $m+1$ that μ_i strongly believes $\Sigma_{-i}^{\Delta,m}$ (only). In particular, this simplification holds when the restrictions concern only independence, or initial beliefs (e.g., exogenous beliefs), or there are no restrictions at all.[20]

The following theorem and corollary are the main game theoretic results of this paper. We provide a proof in the main text because its simplicity illustrates the power of Lemma 3.

Theorem 1 *For every player i, pair (θ_i, s_i) in Σ_i, and positive integer m, strategy s_i is m-strongly Δ-rationalizable for information type θ_i if and only if there exists an epistemic type t_i in T_i such that i satisfies Δ-rationality and mutual strong belief in Δ-rationality of order $m-1$ at personal state (θ_i, s_i, t_i); that is,*

$$\forall i \in I, \forall m \in \mathbb{N},\ \Sigma_i^{\Delta,m} = \mathrm{proj}_{\Sigma_i} R_i^{\Delta,m}.$$

Proof Let $R_i^{\Delta,0} = \Sigma_i \times T_{-i}$, so that the basis step is trivial. Suppose by way of induction that $\Sigma_i^{\Delta,m} = \mathrm{proj}_{\Sigma_i} R_i^{\Delta,m}$ for all i and $m \le n$. Let $(\theta_i, s_i, t_i) \in R_i^{\Delta,n+1}$; then, the inductive hypothesis implies that $\mu_i^1 = \mathrm{mrg}_{\Sigma_{-i}} \beta_i(t_i)$ is a CPS in Δ_{i,θ_i} that strongly believes $\Sigma_{-i}^{\Delta,m}$ for each $m \le n$ and $s_i \in r_i(\mu_i^1, \theta_i)$; hence, $(\theta_i, s_i) \in \Sigma_i^{\Delta,n+1}$. For the other direction, let $(\theta_i, s_i) \in \Sigma_i^{\Delta,n+1}$. By definition of $\Sigma_i^{\Delta,n+1}$, there exists some $\nu_i \in \Delta_{i,\theta_i}$ such that ν_i strongly believes $\Sigma_{-i}^{\Delta,m} = \mathrm{proj}_{\Sigma_{-i}} R_{-i}^{\Delta,m}$ for each $m \le n$ and $s_i \in r_i(\nu_i, \theta_i)$. By Lemma 3 and the inductive hypothesis, there is a CPS μ_i on $(\Sigma_{-i} \times T_{-i}, \mathcal{C}_i)$ such that μ_i strongly believes $(R_{-i}^{\Delta,m})_{m=0}^n$. Since β_i is surjective, there exists some t_i in T_i so that $\beta_i(t_i) = \mu_i$; therefore, $(\theta_i, s_i, t_i) \in R_i^{\Delta,n+1}$. $\square$

Corollary 1 *For every player i and pair (θ_i, s_i) in Σ_i, strategy s_i is strongly Δ-rationalizable for information type θ_i if and only if there exists an epistemic type t_i in T_i such that i satisfies Δ-rationality and common strong belief in Δ-rationality at personal state (θ_i, s_i, t_i); that is,*

$$\forall i \in I,\ \Sigma_i^{\Delta,\infty} = \mathrm{proj}_{\Sigma_i} R_i^{\Delta,\infty}.$$

[18] Catonini (2017b), and Harrington (2017). See Sect. 5.

[19] Artemov et al. (2013) analyze robust virtual implementation with respect to Δ-rationalizability in static mechanisms. Mueller (2016) applies strong rationalizability to dynamic mechanisms. Bergemann and Morris (2016 and references therein) apply to static mechanisms the solution concept without belief restrictions and call it "belief-free rationalizability."

[20] Battigalli and Siniscalchi (2002, 2003) correctly rely on such simplified algorithm.

Proof One inclusion is immediate (and does not use compactness):

$$\Sigma_i^{\Delta,\infty} = \bigcap_{m\in\mathbb{N}} \Sigma_i^{\Delta,m} = \bigcap_{m\in\mathbb{N}} \mathrm{proj}_{\Sigma_i} R_i^{\Delta,m} \supseteq \mathrm{proj}_{\Sigma_i}\left(\bigcap_{m\in\mathbb{N}} R_i^{\Delta,m}\right) = \mathrm{proj}_{\Sigma_i} R_i^{\Delta,\infty},$$

where the first and last equalities hold by definition, the second equality follows from Theorem 1, and the inclusion is obvious. To show the converse, let $\sigma_i \in \Sigma_i^{\Delta,\infty} = \bigcap_{m\in\mathbb{N}} \mathrm{proj}_{\Sigma_i} R_i^{\Delta,m}$, so that for every $m \in \mathbb{N}$ there exists t_i^m in T_i such that $(\sigma_i, t_i^m) \in R_i^{\Delta,m}$. This implies that, for every $m \in \mathbb{N}$, the section $T_{\sigma_i}^m$ of $R_i^{\Delta,m}$ at σ_i is nonempty. Like $(R_i^{\Delta,m})_{m\in\mathbb{N}}$, also $(T_{\sigma_i}^m)_{m\in\mathbb{N}}$ is a decreasing sequence of closed nonempty subsets of a compact space. By the finite intersection property of compact spaces, $\bigcap_{m\in\mathbb{N}} T_{\sigma_i}^m$ is nonempty. If t_i is an element of the intersection, we have $(\sigma_i, t_i) \in \bigcap_{m\in\mathbb{N}} R_i^{\Delta,m}$. Hence, $\sigma_i \in \mathrm{proj}_{\Sigma_i} \bigcap_{m\in\mathbb{N}} R_i^{\Delta,m}$, and the claim follows. □

Example 4 Consider Example 2. Strategies of the seller are identified with his actions, while the set of strategies of the buyer is $S_2 = A_2^{A_1\times\{w\}} \cong A_2^{A_1}$. Therefore, the set of beliefs of the seller (whose only observable event about the buyer corresponds to the empty history) is identified with $\Delta(S_2)$, and we do not impose exogenous restrictions, that is, $\Delta_{1,\theta} = \Delta(S_2)$ for each θ. The conditioning events for the buyer correspond to the empty history and the possible messages by the seller:

$$\mathcal{C}_2 = \{\Theta_1 \times A_1, \Theta_1 \times \{[\theta \geq 1]\}, \ldots, \Theta_1 \times \{[\theta \geq K]\}\}.$$

The belief system of the buyer is given by a CPS in $\Delta^{\mathcal{C}_2}(\Theta_1 \times A_1)$. We assume that the buyer is "mildly skeptical," in the sense that, after observing a message, he assigns probability greater than some small positive value to the worst state consistent with it. Formally,

$$\Delta_2 = \{\mu \in \Delta^{\mathcal{C}_2}(\Theta_1 \times A_1) : \forall k \in \Theta_1, \mu\,(k|[\theta \geq k]) \geq \varepsilon\},$$

where $\mu\,(\cdot|[\theta \geq k])$ is the marginal on Θ_1 of the component of the CPS conditional on $\Theta_1 \times \{[\theta \geq k]\}$, and ε is a fixed positive real number (we assume ε small and, in particular, less than 1).

Battigalli (2006) shows that the only Δ-rationalizable message of a seller of type k is $a_1 = [\theta \geq k]$, and the Δ-rationalizable response of the buyer to such message is $a_2 = k$. The result follows from forward-induction reasoning: First note that strong belief in the rationality of the seller implies that the buyer believes that he tells the truth. With this, weak skepticism implies that the strategy of the buyer satisfies $s_2\,([\theta \geq K]) > s_2([\theta \geq K-1])$. Given that the seller believes in $R_2 \cap [\Delta_2] \cap \mathrm{SB}_2(R_1)$, if $\theta = K$ he will never send message $[\theta \geq K-1]$, because he believes he can make the buyer increase a_2 by sending message $[\theta \geq K]$. Thus, $\mathrm{SB}_2(R_1) \cap \mathrm{SB}_2\,(R_1 \cap \mathrm{SB}_1\,(R_2 \cap [\Delta_2] \cap \mathrm{SB}_2(R_1)))$ implies that if the rational buyer receives message $[\theta \geq K-1]$, he is certain that $\theta = K-1$ and chooses $s_2([\theta \geq K-1]) = K-1$. An induction argument shows that $s_2([\theta \geq k]) = k$ for each k. Formally, the following "full disclosure" result holds:

$$\Sigma_1^{\Delta,\infty} = \{(\theta^*, [\theta \geq k]) \in \Theta_1 \times A_1 : \theta^* = k\},$$
$$\Sigma_2^{\Delta,\infty} = \{s_2 \in S_2 : \forall k \in \{1, \ldots, K\}, s_2([\theta \geq k]) = k\}.$$

Hence, by the characterization result in Corollary 1, these are the sets of type–strategy pairs consistent with rationality, weak skepticism, and common strong belief in rationality and weak skepticism. ▲

Example 5 Consider now the infinitely repeated game described in Example 3. There are no restrictions on the beliefs of player 1: that is,

$$\forall \theta_1 \in \Theta_1, \Delta_{1,\theta_1} = \Delta^{\mathcal{C}_1}(S_2),$$

while player 2 initially assigns probability at least ε to the crazy type θ^c:

$$\Delta_2 = \left\{\mu_2 \in \Delta^{\mathcal{C}_2}(\Theta_1 \times S_1) : \mu_{2,\varnothing}(\{\theta^c\} \times S_1) \geq \varepsilon\right\},$$

where $\varepsilon \in (0, 1)$ is given. Call this restriction ε-possibility of θ^c. Adapting a proof by Fudenberg and Levine (1989) one can show that if the Stackelberg action a_1^* is played long enough, and player 2 strongly believes that player 1 is rational,[21] then player 2 eventually assigns conditional probability close to 1 to a_1^*.

Suppose that the stage game has "conflicting interests," that is, the Stackelberg action holds player 2 to his maxmin payoff:

$$v_2\left(a_1^*, \arg\max_{a_2 \in A_2} v_2\left(a_1^*, a_2\right)\right) = \min_{a_1 \in A_1} \max_{a_2 \in A_2} v_2\left(a_1, a_2\right).$$

Then, adapting arguments from Schmidt (1993) and Watson (1993), one can show that the Stackelberg payoff for player 1 is an approximate lower bound on his strongly Δ-rationalizable expected payoff as he becomes infinitely patient; the reason is that he can eventually "teach" player 2 to play the stage-game best reply to the Stackelberg action:[22]

$$\forall \mu_1 \in \Delta\left(S_2^{\infty,\Delta}\right), \lim_{\delta_1 \to 1} \max_{s_1 \in S_1} \mathbb{E}_{\mu_1}\left(U_1\left(s_1, \cdot\right)\right)$$
$$\geq \lim_{\delta_1 \to 1} \mathbb{E}_{\mu_1}\left(U_1\left(s_1^*, \cdot\right)\right) \geq v_1\left(a_1^*, \arg\max_{a_2 \in A_2} v_2\left(a_1^*, a_2\right)\right).$$

21 In this case, 2's belief in the rationality of 1 after any Stackelberg history $h^* = \left(\left(a_1^*, a_2^1\right), \left(a_1^*, a_2^2\right), \ldots\right)$ is implied by ε-possibility of θ^c.

22 If θ^c is a "commitment type" whose only feasible action is the Stackelberg actions, then the result also holds for "weak," or "initial" Δ-rationalizability. See also Battigalli and Watson (1997), who consider a countable sequence a short-run players in the role of player 2.

By Corollary 1, under common strong belief in rationality and ε-possibility of θ^c, the expected payoff of a rational and very patient player 1 is approximately bounded below by the Stackelberg payoff. ▲

5 Discussion

In this section, we discuss some extensions and generalizations of our analysis.

5.1 More general information structures

For the sake of simplicity, we restricted our attention to simple incomplete-information dynamic games with observable actions, distributed knowledge of θ (i.e., pooling players' private information, the true state of nature θ is identified), and without chance moves. All these simplifications can be removed without affecting the main characterization result.[23]

5.2 Common initial belief in rationality

We focused on epistemic assumptions that capture a form of forward-induction reasoning because this makes the analysis more challenging compared to other assumptions such as common initial belief in rationality, or common belief in rationality conditional on a fixed family of observable events. It is quite straightforward to extend to simple infinite dynamic games the characterization results concerning such assumptions obtained for finite games. Adapting the statement and proof of Lemma 3, we obtain the following:

Lemma 6 *Let X and T be compact metrizable spaces and let $\mathcal{C}$ be a countable subcollection of $\mathcal{B}(X)$. Fix some family $\mathcal{F} \subseteq \mathcal{C}$ and a closed event $E \subseteq X \times T$. For each CPS ν on $(X, \mathcal{C})$ such that $\nu_C(\text{proj}_X E) = 1$ for every $C \in \mathcal{F}$, there is a CPS μ on $(X \times T, \mathcal{C})$ such that $\text{mrg}_X \mu = \nu$ and $\mu_C(E) = 1$ for every $C \in \mathcal{F}$.*

Letting $\mathcal{F} = \{X\} \subseteq \mathcal{C}$ in this modified lemma, we can prove that the behavioral implications of rationality (and Δ-restrictions) and common initial belief in rationality (and in the Δ-restrictions) are characterized by "weak" or "initial" (Δ-)rationalizability (see Battigalli 2003, Battigalli and Siniscalchi 1999, 2007). A similar result holds for rationality and common belief in rationality conditional on a fixed family $\mathcal{F}$ of nonterminal histories (see Battigalli and Siniscalchi 1999).

5.3 Transparent belief restrictions

Battigalli and Prestipino (2013) prove, for finite games, that strong Δ-rationalizability also characterizes the behavioral implications of a stronger set of epistemic assumptions than those considered here, that is, stronger than rationality, $[\Delta]$ and common

[23] For an epistemic analysis of finite, but otherwise general dynamic games of incomplete information, see Battigalli and Prestipino (2013).

strong belief in rationality and $[\Delta]$. Specifically, define the mutual *full belief* operator $\mathrm{B}(\cdot)$ as follows: for every event $E=\prod_{i\in I}E_i$,

$$\mathrm{B}(E)=\prod_{i\in I}\mathrm{B}_i(E_{-i})=\prod_{i\in I}\left\{(\theta_i,s_i,t_i):\forall h\in H,\beta_{i,h}(t_i)(E_{-i})=1\right\}.$$

An event is *transparent* at a state (θ,s,t) if E is true and there is common full belief of E at (θ,s,t), that is,

$$(\theta,s,t)\in\bigcap_{n=0}^{\infty}\mathrm{B}^n(E)=:\mathrm{B}^*(E),$$

where $\mathrm{B}^0(E)=E$. Also, as in Battigalli and Siniscalchi (2002), let

$$\mathrm{CSB}(E)=\prod_{i\in I}E_i\cap\mathrm{SB}_i(E_{-i})$$

denote the mutual correct strong belief operator. It is routine to check that

$$\prod_{i\in I}R_i^{\Delta,m}=\mathrm{CSB}^{m-1}(R\cap[\Delta])$$

for each m. Battigalli and Prestipino (2013) prove that

$$\forall m\in\mathbb{N},\ \mathrm{proj}_{\Theta\times S}\mathrm{CSB}^{m-1}(R\cap[\Delta])=\prod_{i\in I}\Sigma_i^{\Delta,m}=\mathrm{proj}_{\Theta\times S}\mathrm{CSB}^{m-1}\left(R\cap\mathrm{B}^*([\Delta])\right)$$

and

$$\mathrm{proj}_{\Theta\times S}\mathrm{CSB}^{\infty}(R\cap[\Delta])=\prod_{i\in I}\Sigma_i^{\Delta,\infty}=\mathrm{proj}_{\Theta\times S}\mathrm{CSB}^{\infty}\left(R\cap\mathrm{B}^*([\Delta])\right),$$

that is, one can replace $[\Delta]$ with transparency of $[\Delta]$ in the epistemic assumptions considered here. This is equivalent to saying that Δ-rationalizability characterizes the behavioral implications of rationality and common strong belief in rationality in the substructure of the canonical structure obtained by assuming transparency of the restrictions Δ.[24] By inspection of the proof of Battigalli and Prestipino, it is clear that it can be applied *verbatim* to the more general case of simple games considered here.

5.4 Self-enforcing agreements and forward induction

Let Δ represent the restrictions implied by belief in the compliance with a preplay nonbinding agreement. Harrington (2017) applies strong Δ-rationalizability to analyze incomplete collusive agreements in infinite oligopoly models where firms set prices in

[24] See the discussion in Battigalli and Friedenberg (2012) and Battigalli and Prestipino (2013).

a finite grid. Such games are simple, and our result provides an epistemic foundation of Harrington's analysis. Catonini (2017a, b) develops a theory of self-enforcing (possibly incomplete) agreements in *finite* games based on forward-induction reasoning. He argues that strong Δ-rationalizability can capture aspects of such a theory, although a somewhat different solution concept, selective rationalizability, is more appropriate. Catonini's argument refers to the result by Battigalli and Prestipino (2013) mentioned above, which shows that Δ-rationalizability implicitly gives "epistemic priority" to belief in the agreement over belief in rationality when they cannot be reconciled. Catonini (2017a) shows that selective rationalizability characterizes the behavioral implications of $\mathrm{CSB}^{\infty}([\Delta] \cap \mathrm{CSB}^{\infty}(R))$. Lemma 3 allows to extend Catonini's theory to simple infinite games.

5.5 Evidence games

Example 2 is a special case of the class of *evidence games* where an agent (player 1) has the ability to exhibit verifiable information about a discrete state of nature and a principal (player 2) chooses a reward. There is a rich literature about mechanisms and evidence games,[25] which are therefore an important subset of the class of infinite dynamic games with incomplete asymmetric information analyzed in this paper. We conjecture that the forward-induction analysis of Example 4 and—more generally—the epistemic analysis allowed by our results can be fruitfully applied to shed light on this literature.

5.6 Nonsimple games

The main reason to restrict our analysis to simple dynamic games is that this allows us to prove that the topological properties of the primitive uncertainty conditional space $(X, \mathcal{C})$ are inherited by $\left(X \times \Delta^{\mathcal{C}}(X), \mathcal{C}\right)$: if $(X, \mathcal{C})$ is a simple conditional space, so is $\left(X \times \Delta^{\mathcal{C}}(X), \mathcal{C}\right)$ (see Battigalli and Siniscalchi (1999) and Lemmas 1, 4). However, the requirement that each conditioning event C in $\mathcal{C}$ is both closed and open (clopen) is restrictive in some applications. For example, suppose that Θ_{-i} is uncountable and that i gets signals about θ_{-i}; then, even if signals are discrete, the set of profiles $\left(\theta_j, s_j\right)_{j \neq i}$ consistent with an information set h_i, $\Sigma_{-i}(h_i)$ may be not clopen. Similarly, facial expressions and other observable features of co-players in face-to-face interaction may be signals of, say, co-players' first-order beliefs; then, the set of profiles $\left(\theta_j, s_j, \mu_j^1\right)_{j \neq i}$ consistent with an information set h_i may be not clopen. We can prove a version of Lemma 3 that does not assume the clopeness of conditioning events in $\mathcal{C}$ and perhaps can be applied to more general constructions of hierarchical belief spaces.[26]

25 See, for example, Hart et al. (2017) and the references therein.

26 We thank Gabriele Beneduci for proving this generalization. See the working paper version: http://www.igier.unibocconi.it/folder.php?vedi=6337&tbn=albero&id_folder=4878.

Appendix: Proofs

Proof of Lemma 3

Fix some $\nu \in \Delta^{\mathcal{C}}(X)$ that strongly believes $(\text{proj}_X E_1, \ldots, \text{proj}_X E_n)$. Define recursively a partition $\{\mathcal{C}_0, \mathcal{C}_1, \ldots\}$ of $\mathcal{C}$ as follows: let $\mathcal{C}_0 = \{C \in \mathcal{C} : \nu_X(C) > 0\}$. Suppose we have defined $\mathcal{C}_0, \ldots, \mathcal{C}_m$, and let

$$\bar{\mathcal{C}}_{m+1} = \{\bar{C} \in \mathcal{C} \setminus (\mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m) : \forall C \in \mathcal{C},\ C \supset \bar{C} \Rightarrow C \in \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m\}$$

be the collection of conditioning events $\bar{C}$ not in $\mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m$ whose *strict* predecessors C are all in $\mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m$; then, let

$$\mathcal{C}_{m+1} = \{D \in \mathcal{C} \setminus (\mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m) : \exists \bar{C} \in \bar{\mathcal{C}}_{m+1}, \bar{C} \supseteq D, \nu_{\bar{C}}(D) > 0\}.$$

For convenience, let $\bar{\mathcal{C}}_0 = \{X\}$. Note that, by definition of $\mathcal{C}_m$ and since ν is a CPS, we have $\bar{\mathcal{C}}_m \subseteq \mathcal{C}_m$ for every $m = 0, 1, \ldots$ (in particular, $\bar{\mathcal{C}}_0 = \{X\} \subseteq \mathcal{C}_0$).

Claim 1 *Every element of $\mathcal{C}_m$ has a unique predecessor in $\bar{\mathcal{C}}_m$.*

First note that it is true by definition that each element of $\mathcal{C}_m$ has at least one predecessor in $\bar{\mathcal{C}}_m$; we only have to prove uniqueness. The claim is obvious for $m = 0$. For $m = 1, 2, \ldots$, let $C \in \mathcal{C}_m$ and consider any two *distinct* predecessors $C', C'' \supset C$ (if they exist); then—by the tree-like property of $\mathcal{C}$—either $C'' \supset C'$ or $C' \supset C''$. Thus, by definition of $\bar{\mathcal{C}}_m$, at most one of C' and C'' can belong to $\bar{\mathcal{C}}_m$. □

Claim 2 $\{\mathcal{C}_0, \mathcal{C}_1, \ldots\}$ *is a partition of $\mathcal{C}$.*

By the construction above, it is clear that the collections $\mathcal{C}_m$ are pairwise disjoint. Let us show that $\mathcal{C} = \bigcup_{n \geq 0} \mathcal{C}_n$. Pick any C in $\mathcal{C}$. By definition of simple conditional space, the collection of predecessors (supersets) of C in $\mathcal{C}$, denoted by $\mathcal{C}^{\supseteq}(C)$, is finite and totally ordered by $\supseteq$. Hence, $\mathcal{C}^{\supseteq}(C)$ can be written as $\{C_0, \ldots, C_m\}$ where all C_ks are distinct ($k = 0, \ldots, m$) and $X = C_0 \supset \cdots \supset C_{m-1} \supset C_m = C$. Remark that, for every k less than m, $\{C_0, \ldots, C_k\}$ is the set of strict predecessors of C_{k+1}. We show inductively that

$$C_k \in \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_k$$

for all $k = 0, \ldots, m$, so that, in particular, $C = C_m \in \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_m \subseteq \bigcup_{n \geq 0} \mathcal{C}_n$. Indeed, $C_0 = X \in \mathcal{C}_0$ is trivially true. Suppose that $C_k \in \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_k$ for all $k = 0, \ldots, \ell \leq m - 1$: if $C_{\ell+1} \in \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_\ell$ we are done; if not, every strict predecessor of $C_{\ell+1}$ belongs to $\mathcal{C}_0 \cup \cdots \cup \mathcal{C}_\ell$ by the inductive hypothesis, which implies—by definition of $\bar{\mathcal{C}}_{\ell+1}$—that $C_{\ell+1} \in \bar{\mathcal{C}}_{\ell+1} \subseteq \mathcal{C}_{\ell+1}$.

By Claims 1–2, we can define a map $e : \mathcal{C} \to \mathcal{C}$ that associates each $C \in \mathcal{C}$ with the *earliest* predecessor $\bar{C}$ of C such that $\nu_{\bar{C}}(C) > 0$: let $\mathcal{C}_m$ be the cell of partition $\{\mathcal{C}_0, \mathcal{C}_1, \ldots\}$ that contains C, then $\bar{C} = e(C)$ is the unique predecessor of C in $\bar{\mathcal{C}}_m$.

Note that $C \cap \text{proj}_X E_m \neq \emptyset$ if and only if $E_m \cap (C \times T) \neq \emptyset$, for all $m = 0, \ldots, n$ and C in $\mathcal{C}$. Since ν strongly believes each $\text{proj}_X E_m$, by Lemma 2 we can find an array of probability measures $\bar{\mu} = (\bar{\mu}_C)_{C \in \mathcal{C}}$ in $[\Delta(X \times T)]^{\mathcal{C}}$ (possibly not a CPS) such that $\text{mrg}_X \bar{\mu} = \nu$ and $\bar{\mu}$ strongly believes $\mathcal{E}$. More explicitly: for all C in $\mathcal{C}$ and for all $m = 0, \ldots, n$,

$$\text{mrg}_X \bar{\mu}_C = \nu_C,$$

and

$$E_m \cap (C \times T) \neq \emptyset \Rightarrow \bar{\mu}_C(E_m \cap (C \times T)) = 1.$$

We use some of these measures to construct the desired CPS μ:

For all $m = 0, 1, \ldots$, for all C in $\mathcal{C}_m$, and for all E in $\mathcal{B}(X \times T)$, let

$$\mu_C(E) = \frac{\bar{\mu}_{e(C)}(E \cap (C \times T))}{\bar{\mu}_{e(C)}(C \times T)},$$

where the denominator is positive because $\bar{\mu}_{e(C)}(C \times T) = \text{mrg}_X \bar{\mu}_{e(C)}(C) = \nu_{e(C)}(C) > 0$ by definition of $e(C)$. Since $\{\mathcal{C}_0, \mathcal{C}_1, \ldots\}$ is a partition of $\mathcal{C}$, we obtain an array $(\mu_C)_{C \in \mathcal{C}}$ in $[\Delta(X \times T)]^{\mathcal{C}}$. Clearly, $\mu_C(C \times T) = 1$ for all C in $\mathcal{C}$.

Claim 3 *ν is the marginal of μ.*

For each C in $\mathcal{C}$ and each event of the form $E \times T$ for E in $\mathcal{B}(X)$,

$$\begin{aligned}\mu_C(E \times T) &= \frac{\bar{\mu}_{e(C)}((E \cap C) \times T))}{\bar{\mu}_{e(C)}(C \times T)} = \frac{\nu_{e(C)}(E \cap C)}{\nu_{e(C)}(C)} \\ &= \frac{\nu_C(E \cap C)\, \nu_{e(C)}(C)}{\nu_{e(C)}(C)} = \nu_C(E),\end{aligned}$$

where we used the definition of $e(C)$ and the fact that ν satisfies the chain rule (1). □

Claim 4 *μ strongly believes $\mathcal{E}$.*

Let $m \in \{1, \ldots, n\}$ and $C \in \mathcal{C}$. We must show that $E_m \cap (C \times T) \neq \emptyset$ implies $\mu_C(E_m) = 1$. If $E_m \cap (C \times T) \neq \emptyset$, then $C \cap \text{proj}_X E_m \neq \emptyset$, and $e(C) \cap \text{proj}_X E_m \neq \emptyset$ because $C \subseteq e(C)$ by definition. Since $\bar{\mu}$ strongly believes $\mathcal{E}$, $\bar{\mu}_{e(C)}(E_m) = 1$, which implies $\bar{\mu}_{e(C)}(E_m \cap (C \times T)) = \bar{\mu}_{e(C)}(C \times T)$ and thus

$$\mu_C(E_m) = \frac{\bar{\mu}_{e(C)}(E_m \cap (C \times T))}{\bar{\mu}_{e(C)}(C \times T)} = \frac{\bar{\mu}_{e(C)}(C \times T)}{\bar{\mu}_{e(C)}(C \times T)} = 1.$$

□

Claim 5 *μ satisfies the chain rule* (1).

Fix an event E and two conditioning events C and D in $\mathcal{C}$ with $E \subseteq D \times T \subseteq C \times T$. We must show that $\mu_C(E) = \mu_D(E)\, \mu_C(D \times T)$. There are two cases: (i) If

$e(C) = e(D)$, then

$$\begin{aligned}\mu_C(E) = \frac{\bar{\mu}_{e(C)}(E)}{\bar{\mu}_{e(C)}(C \times T)} &= \frac{\bar{\mu}_{e(D)}(E)}{\bar{\mu}_{e(D)}(C \times T)} \\ &= \frac{\bar{\mu}_{e(D)}(E)}{\bar{\mu}_{e(D)}(D \times T)} \frac{\bar{\mu}_{e(D)}(D \times T)}{\bar{\mu}_{e(D)}(C \times T)} \\ &= \mu_D(E)\, \mu_C(D \times T).\end{aligned}$$

(ii) If $e(C) \neq e(D)$, then $D \subseteq e(D) \subset C \subseteq e(C)$ and $\nu_{e(C)}(e(D)) = 0$. Therefore $\bar{\mu}_{e(C)}(E) = \bar{\mu}_{e(C)}(D \times T) = 0$, which implies that (1) holds:

$$\begin{aligned}\mu_C(E) &= \frac{\bar{\mu}_{e(C)}(E)}{\bar{\mu}_{e(C)}(C \times T)} \\ &= 0 \\ &= \frac{\bar{\mu}_{e(D)}(E)}{\bar{\mu}_{e(D)}(D \times T)} \frac{\bar{\mu}_{e(C)}(D \times T)}{\bar{\mu}_{e(C)}(C \times T)} \\ &= \mu_D(E)\, \mu_C(D \times T).\end{aligned}$$

□

Since ν is the marginal of μ, which strongly believes $\mathcal{E}$ and satisfies all the properties of a CPS, the theorem is proved. □

Other proofs

Proof of Lemma 4

Let $H_n = H \cap A^n$ be the set of nonterminal histories of length n. By assumption, H_n is finite for each $n = 0, 1, \ldots$; therefore, $H = \bigcup_{n \in \mathbb{N}_0} H_n$ is countable. Each set $\mathcal{A}_j(h)$ is a compact metrizable space; hence, the countable Cartesian product $S_j = \prod_{h \in H} \mathcal{A}_j(h)$ is compact and metrizable as well. The set Θ_j is compact metrizable and $I \setminus \{i\}$ is countable, hence $\Sigma_{-i} = \prod_{j \neq i}(\Theta_j \times S_j)$ is a compact metrizable space.

We have to show that $\mathcal{C}_i$ is a countable collection of clopen subsets of Σ_{-i} with a tree-like structure. The collection $\mathcal{C}_i$ is countable because H is countable, and it inherits its tree-like structure from the countable tree H: indeed, $\Sigma_{-i} = \Sigma_{-i}(\varnothing) \in \mathcal{C}_i$, and the set of strict predecessors of any $\Sigma_{-i}(h)$ in $\mathcal{C}_i$ is the finite collection $\mathcal{C}_i^{\prec}(h) = \{\Sigma_{-i}(h') : h' \prec h\}$; pick any pair of distinct predecessors $\Sigma_i(h')$ and $\Sigma_i(h'')$ in $\mathcal{C}_i^{\prec}(h)$. Since H is a tree, either $h' \prec h''$ and $\Sigma_{-i}(h'') \subset \Sigma_{-i}(h')$ or $h'' \prec h'$ and $\Sigma_{-i}(h') \subset \Sigma_{-i}(h'')$.

Obviously, $\Sigma_{-i}(\varnothing) = \Sigma_{-i}$ is closed and open in Σ_{-i}. To see that $\Sigma_{-i}(h)$ is closed for each h in $H \setminus \{\varnothing\}$, let $h = (a^1, \ldots, a^n)$ and note that

$$\Sigma_{-i}(h) = \prod_{j \neq i} \left(\Theta_j \times \{a_j^1\} \times \cdots \times \{a_j^n\} \times \prod_{h' \in H : h' \not\prec h} \mathcal{A}_j(h') \right)$$

where each set in the product is closed. To see that $\Sigma_{-i}(h)$ is also open in Σ_{-i}, we first show that $\Sigma(h)$ is open in Σ. Note that $\{\Sigma(h') : h' \in H_n\}$ is the finite partition of subsets of Σ obtained from the preimages through map $(\theta, s) \mapsto \zeta(s)$ of the elements of the finite partition $\{Z(h') : h' \in H_n\}$ of Z. Each $\Sigma(h')$ is a product of closed subsets (see above) and it is closed as well; therefore, the finite union $\bigcup_{h' \in H_n \setminus \{h\}} \Sigma(h')$ is closed and $\Sigma(h) = \Sigma \setminus \bigcup_{h' \in H_n \setminus \{h\}} \Sigma(h')$ is open. Since a Cartesian product $C = \prod_{j \in I} C_j \subset \Sigma$ is open if and only if each C_j is open, it follows that each $\Sigma_j(h)$ is open and $\Sigma_{-i}(h) = \prod_{j \neq i} \Sigma_j(h)$ is open. □

Proof of Lemma 5

The result follows from standard compactness-continuity arguments and is therefore omitted (see Battigalli 2003). □

References

Aliprantis, C.R., Border, K.C.: Infinite Dimensional Analysis. Springer, Berlin (2006)
Arieli, I.: Rationalizability in continuous games. J. Math. Econ. **46**, 912–924 (2010)
Artemov, G., Kunimoto, T., Serrano, R.: Robust virtual implementation: toward a reinterpretation of the Wilson doctrine. J. Econ. Theory **148**, 424–447 (2013)
Battigalli, P.: Rationalizability in infinite, dynamic games of incomplete information. Res. Econ. **57**, 1–38 (2003)
Battigalli, P.: Rationalization in signaling games: theory and applications. Int. Game Theory Rev. **8**, 1–27 (2006)
Battigalli, P., Friedenberg, A.: Forward induction reasoning revisited. Theor. Econ. **7**, 57–98 (2012)
Battigalli, P., Prestipino, A.: Transparent restrictions on beliefs and forward induction reasoning in games with asymmetric information. B.E. J Theor. Econ. **13**(1), 1–53 (2013)
Battigalli, P., Siniscalchi, M.: Hierarchies of conditional beliefs and interactive epistemology in dynamic games. J. Econ. Theory **88**, 188–230 (1999)
Battigalli, P., Siniscalchi, M.: Strong belief and forward induction reasoning. J. Econ. Theory **106**, 356–391 (2002)
Battigalli, P., Siniscalchi, M.: Rationalization and incomplete information. Adv. Theor. Econ. **3**(1), Article 3 (2003)
Battigalli, P., Siniscalchi, M.: Interactive epistemology in games with payoff uncertainty. Res. Econ. **61**, 165–184 (2007)
Battigalli, P., Watson, J.: On ‘reputation’ refinements with heterogeneous beliefs. Econometrica **65**, 369–374 (1997)
Ben Porath, E.: Rationality, nash equilibrium, and backward induction in perfect information games. Rev. Econ. Stud. **64**, 23–46 (1997)
Bergemann, D., Morris, S.: Belief-free rationalizability and informational robustness. In: Cowles Foundation Discussion Paper 2066, Yale University (2016)
Catonini, E.: Rationalizability and Epistemic Priority Orderings. Mimeo. Higher School of Economics, Moscow (2017a)
Catonini, E.: Self-enforcing Agreements and Forward-Induction Reasoning. Mimeo. Higher School of Economics, Moscow (2017b)
Chen, Y., Luo, X., Qu, C.: Rationalizability in general situations. Econ. Theory **61**(1), 147–167 (2016). https://doi.org/10.1007/s00199-015-0882-y
Dekel, E., Siniscalchi, M.: Epistemic game theory. In: Young, P., Zamir, S. (eds.) Handbook of Game Theory with Economic Applications, vol. 4, pp. 619–702. North-Holland, Amsterdam (2015)
Friedenberg, A., Meier, M.: On the relationship between hierarchy and type morphisms. Econ. Theory **46**(3), 377–399 (2011). https://doi.org/10.1007/s00199-010-0517-2

Fudenberg, D., Levine, D.K.: Reputation and equilibrium selection in games with a patient player. Econometrica **57**, 759–778 (1989)
Harrington, J.: A theory of collusion with partial mutual understanding. Res. Econ. **71**, 140–158 (2017)
Harsanyi, J.: Games of incomplete information played by bayesian players. Parts I, II, III. Manag. Sci. **14**, 159–182, 320–334, 486–502 (1967–1968)
Hart, S., Kremer, I., Perry, M.: Evidence games: truth and commitment. Am. Econ. Rev. **107**, 690–713 (2017)
Heifetz, A., Samet, D.: Topology-free typology of beliefs. J. Econ. Theory **82**(2), 324–341 (1998)
Lubin, A.: Extensions of measures and the von Neumann selection theorem. Proc. Am. Math. Soc. **43**, 118–122 (1974)
Mueller, C.: Robust virtual implementation under common strong belief in rationality. J. Econ. Theory **162**, 407–450 (2016)
Pearce, D.: Rationalizable strategic behavior and the problem of perfection. Econometrica **52**, 1029–1050 (1984)
Schmidt, K.: Reputation and equilibrium characterization in repeated games with conflicting interests. Econometrica **61**, 325–351 (1993)
Shimoji, M., Watson, J.: Conditional dominance, rationalizability, and game forms. J. Econ. Theory **83**, 161–195 (1998)
Tan, T., Werlang, S.: The Bayesian foundation of solution concepts of games. J. Econ. Theory **45**, 370–391 (1988)
Watson, J.: A ‘reputation’ refinement without equilibrium. Econometrica **61**, 199–205 (1993)
Varadarajan, V.S.: Groups of automorphisms of Borel spaces. Trans. Am. Math. Soc. **109**, 191–220 (1963)