

Beyond Common Priors

Paulo Barelli*

February 2004

Abstract

One property (called action-consistency) that is implicit in the common prior assumption (CPA) is identified and shown to be the driving force of the use of the CPA in a class of well-known results. In particular, we show that Aumann (1987)'s Bayesian characterization of correlated equilibrium, Aumann and Brandenburger (1995)'s epistemic conditions for Nash equilibrium, and Milgrom and Stokey (1982)'s no-trade theorem are all valid without the CPA but with action-consistency. Moreover, since we show that action-consistency is much less restrictive than the CPA, the above results are more general than previously thought, and insulated from controversies around the CPA.

JEL Classification: C70, C72, D80, D82

1 Introduction

Harsanyi (1967, 1968) introduced the common prior assumption (CPA) as a useful condition to be used in the analysis of games of incomplete information. Ever since it has been widely used in game theory and information economics. Indeed, the CPA is associated with most applied work with differential information, and there are some very important results (e.g. no trade theorems) that depend crucially on the CPA. Nonetheless, it is a restrictive assumption that has generated some controversies in the literature, especially in terms of its interpretation. A vast literature emerged from such controversies, a literature that focuses mainly on characterizing the CPA.

*I am grateful to Max Amarante under whose supervision this research has been undertaken. I also thank Adam Brandenburger, Prajit Dutta, Larry Epstein, Emanuele Gerratana, Barton Lipman, Alberto Martin, Sangwon Park, Marco Scarsini, and Paolo Siconolfi for important comments and discussions.

One feature that has not been fully considered by this literature is how important the assumption is. That is, whether the CPA is crucial to results that are known to hold under the CPA or not. If some results are valid without the CPA, then not only we improve our understanding about the nature of the results themselves, but we may also hope to dissociate the results from the controversies around the CPA.

This paper addresses such issues. We identify a property (which we term action-consistency) that underlies a class of results that use the CPA, and show that such property is not only different from the CPA, but it is also much less restrictive. More precisely, consider the class of results that depend both on the CPA and on (common) knowledge of some events related to actions that agents might take in a strategic situation. We show that every such result is ultimately based on action-consistency, and not on the CPA. As an implication, we also show that not even the common knowledge assumptions are crucial for such results. It follows that action-consistency reveals that the results in the class above are more general than they were thought to be, and not necessarily liable to controversies around the CPA.

For instance, when we say (see Aumann (1987)) that under the CPA rational players play a correlated equilibrium, we are in fact restricting the number states of the world (strategic situations) where we can ensure that players play a correlated equilibrium. With our result, Aumann's proposition is shown to be valid in many more states, since it is valid under action-consistency and there are states compatible with action-consistency and not compatible with the CPA. The following simple example illustrates the idea. Consider a 2 player game with the following payoff matrix

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

and say that each player can take one of three types represented by the following table:

	L ₁	L ₂	R ₁
U ₁	$\frac{1}{4}, \frac{3}{8}$	$\frac{3}{4}, \frac{1}{8}$	0, 0
U ₂	$\frac{3}{4}, \frac{1}{8}$	$\frac{1}{4}, \frac{3}{8}$	0, 0
D ₁	$\frac{1}{4}, \frac{1}{2}$	$\frac{1}{4}, \frac{1}{2}$	$\frac{1}{2}, 1$

where the letter designated to the type indicates the associated action (U₁ plays U, and so on). The numbers in the table represent the beliefs of each type; for instance, if column player is of type L₂, then he believes that $(\frac{1}{8}, \frac{3}{8}, \frac{1}{2})$ are the probabilities of the row player being of types U₁, U₂, D₁, respectively. It is not hard to verify that the CPA is violated: there is no probability distribution that generates the types above as conditional distributions. Nevertheless, the types above satisfy the property that we call action-consistency: the marginal distributions over the set of actions can be generated as conditional distributions of a probability distribution. Indeed, the probability distribution whose marginal over the events UL, UR, DL, and DR is $(\frac{1}{3}, 0, \frac{1}{3}, \frac{1}{3})$ is one such distribution. Such distribution is a correlated equilibrium distribution of the game. The CPA is violated, but we do have rationality of the players (the action associated to each type is a best reply to the beliefs of the type), and action-consistency. And Aumann's result still hold, for we can say that players indeed take part in a correlated equilibrium.

As the example shows, there are instances where the full-force of the CPA is not needed. One immediate implication is that action-consistency is less restrictive than the CPA (since, clearly, CPA implies action-consistency), and one question is to determine how less restrictive it is. The answer is: a lot less. We show that the set of states of the world compatible with action-consistency is a dense subset of the state space, while the set of states compatible with the CPA is not a dense subset of the state space. Given that action-consistency reveals where the CPA bites, and it is also much less restrictive, we can say that action-consistency is a meaningful way of relaxing the CPA.

Since action-consistency is less restrictive than the CPA, it provides one way out of controversies around the CPA. For instance, the usual interpretation of the CPA is that differences in beliefs ought to come from differences in information alone, not differences in opinion. This interpretation is derived

from a philosophical position, namely, that probabilities should be based only on information. It is one of the sources of controversy over the CPA. Under action-consistency, all that is required is that players be consistent on their beliefs about events related to the actions they might take in the given strategic situation. Hence, players might have different opinions about many events (for instance, those that they are allowed to hold inconsistent beliefs), so that it allows beliefs (probabilities) to depend on more than only information. Moreover, the denseness result mentioned above shows that any other kind of controversy over the CPA need not carry over to action-consistency, for every state of the world is approximately action-consistent.

In order to illustrate that the crucial feature of the CPA is the action-consistency property that is implicit in it, we present three well-known theorems in games and information economics: Two of them refer to equilibria in games: Aumann (1987)'s Bayesian characterization of correlated equilibrium and Aumann and Brandenburger (1995)'s epistemic conditions for Nash equilibrium. The third one is Milgrom and Stokey (1982)'s no-trade theorem. All of them can be stated with action-consistency replacing the CPA. After presenting such theorems, we then argue that they belong to a well-defined class of results¹, and that all results in that class can be stated under action-consistency instead of the CPA.

1.1 The CPA

The crucial, and most insightful, concept developed by Harsanyi is that of type-agent.² The relevant attributes of a player (agent) are summarized by his type, which is known by himself but not by his opponents. The opponents are not sure about the type of the given player: each of them considers that the given player could have had many other attributes (that is, that he could have taken many other types). Such situation can be described by a space of type-agents: a set of points, where each point is a list of types of the players involved in a particular strategic situation. Given that what is uncertain for each given player is the types of his opponents, each type ought to be associated with one probability distribution over the types of the opponents (following, as Harsanyi did, the Bayesian tradition). In fact, one can identify a type with a probability distribution over (opponents's) types. Hence, each point in the space of type-agents is a list of probability distributions.

¹That of results that are stated with the CPA in conjunction to an assumption of common knowledge of events that are related to actions in the given strategic situation.

²Harsanyi actually followed an idea of Selten.

Here is where Harsanyi introduced the assumption of **consistency**: these probability distributions are consistent with each other if they can be generated as posteriors of a single prior (where the updating comes from each player knowing his own type). That is, each type is just the conditional probability distribution that can be computed from the prior distribution by Bayes's formula. This single prior has come to be known as the common prior, and frequently the assumption of consistency is viewed as the common prior assumption (CPA). The following example illustrates. There are two players, and each player can take two types. For instance, when player 1 is of type 2 (that is, t_{12}), he believes that player 2 will take types 1 and 2 (t_{21} and t_{22}) with probabilities γ and $1 - \gamma$, respectively.

	t_{21}	t_{22}
t_{11}	α, β	$1-\alpha, \eta$
t_{12}	$\gamma, 1-\beta$	$1-\gamma, 1-\eta$

These four types are consistent if they can be computed as conditional distributions of a single probability distribution, a statement that is equivalent to requiring that

$$\frac{\beta}{1-\beta} \frac{1-\alpha}{\alpha} = \frac{\eta}{1-\eta} \frac{1-\gamma}{\gamma},$$

and in which case the consistent distribution is given by $\Delta^{-1}(\alpha\beta\gamma, (1-\alpha)\beta\gamma, \alpha(1-\beta)\gamma, \alpha(1-\beta)(1-\gamma))$, where $\Delta = \alpha(1-\beta) + \beta\gamma$. It is clear that consistency is a property of a particular set of types, and that this property might easily fail. Indeed, in all likelihood four arbitrary numbers α, β, γ and η will not satisfy the equality above.

1.2 Interpretation

The CPA's usefulness is easy to understand: it is much simpler to specify one probability distribution than many of them, one for each type of each player. The controversy begins when it comes to interpret the concept. The usual interpretation is as follows: there is an ex ante stage where prior beliefs are formed, and an interim stage (when the actual strategic situation that we like to study takes place), where each player gets to know his type and updates his prior beliefs accordingly. The CPA is the requirement that

the prior beliefs, in the ex ante stage, coincide. If the interim beliefs (the types) differ, it must be because players received different information. The question now is to justify why the prior beliefs would coincide.

The usual justification is exactly that differences in beliefs ought to come from differences in information only, not from differences in opinion. Prior beliefs represent opinions, types represent beliefs, and information is represented by each player knowing his own type at the interim stage. There is no reason, the argument goes, to have differences in beliefs after having received the same information. One problem with such argument is that prior to any given strategic situation, the players have most likely been subject to other strategic situations, and have (conceivably) received different information. So the beliefs that they formed in those prior situations will not coincide in general. But these beliefs are nothing else than the prior beliefs in the given strategic situation, so consistency seems to be a rather unlikely state of affairs.

1.3 Why Relax the CPA

We just argued that the main problem with the CPA is its interpretation; there is a philosophical position to be taken, that of probabilities being generated by information alone. Not to mention the debate around the CPA mentioned above. Such considerations lead naturally to a quest for a better understanding of the CPA. Indeed, the literature characterizing the CPA is now vast (see, among others, Feinberg (2000), Heifetz (2001), Nehring (2001), Lipman (2003), and Samet (1998)). The perspective taken here can be seen as yet another attempt to get a better understanding of the CPA. By identifying action-consistency as the crucial feature of the CPA in many applications, we provide some answers to some problems associated with the CPA.

First, as mentioned above, action-consistency does not require that prior beliefs be equal, and hence is compatible with probabilities being generated by opinions as well as information. Second, any debate around the CPA carries over to the results that depend on the CPA. Given that some of those results are widely used theorems in game theory and information economics, the theorems themselves might be subject to the same kind of criticism that is targeted at the CPA.³ Hence, by dissociating those results from the CPA,

³Indeed, Gul (1998)'s critique to the CPA was targeted at Aumann (1987)'s Bayesian characterization of correlated equilibrium.

we insulate them from controversies around the CPA. As a by-product, we get a better grasp of the very nature of those well-known results as well. Third, the fact that the set of states compatible with action-consistency is a dense subset of the state space provides an answer to any controversy around the CPA: given that any state of the world can be approximated by an action-consistent state, in applications we can use the assumption regardless of any debate around its plausibility. We cannot do the same with the CPA: we show below that the set of consistent states is a negligible subset of the state space

1.4 Action-Consistency

Harsanyi's consistency is the requirement that the types of the players can be computed as conditional probabilities of a single prior. This means that, for a given player, his beliefs over every event in the state space can be computed as the conditional probability of a single prior. Action-consistency reduces the class of events that have to satisfy the above property; that is, under action-consistency, only the beliefs over events related to actions in the game are required to be computable from a single prior. In other words, only the marginal distributions of the types over actions are assumed to be consistent. Informally, the CPA (or consistency) requires too high a degree of consistency among the players; they have to have consistent beliefs for every event in the state space, including those that are not directly related to the actions of the players in the game. Since what matters in most applications is the joint action ultimately taken by the group of players, it is clear that the CPA might be too strong a requirement for some applications. Action-consistency allows for inconsistencies on intrinsically meaningless (for the particular application) events.

1.5 What We Learn from Action-Consistency

Given that we can restate well-known results using action-consistency instead of the CPA, action-consistency casts new light on these results. For instance, we are used to think (after Aumann and Brandenburger (1995)'s classic paper) that the CPA is a required condition for players to play a Nash equilibrium. Well, not necessarily: they only have to be consistent in events related to actions in the game. The same logic applies to all other theorems considered here. It follows that action-consistency reveals features of those theorems that were hidden by the CPA. In particular, it reveals the theorems hold more generally than their original formulations. Indeed, the

denseness result mentioned before shows that the theorems hold to a considerably higher degree of generality. In addition, we point out one important finding arising from the use of action-consistency: *ex ante* knowledge.

1.5.1 Ex Ante Knowledge

One by-product of stating well-known results with action-consistency instead of the CPA is that the common knowledge requirements in each of these results can also be relaxed. That is, each of the results considered here has the following form: under the CPA (and some other conditions), if some event is common knowledge, then the result follows. For instance, under rationality and knowledge of the game, and under the CPA, if the conjectures of each player about what action will be taken by his opponents are common knowledge, then players play a Nash equilibrium. Once we use action-consistency in the place of the CPA, then we are also led to use some other requirement in the place of common knowledge of the event. Formally, the requirement is that the action-consistent measure puts probability one to the event that was supposed to be common knowledge. Such condition will be called *ex ante* knowledge of the event. Using this terminology, theorems that required the CPA together with common knowledge of some event can be restated with action-consistency and *ex ante* knowledge of the given event. Hence, there is yet another gain in generality: the theorems do not necessarily require common knowledge assumptions.

It must be emphasized that *ex ante* knowledge does not mean knowledge in the usual sense, which is an interim concept. *Ex ante* knowledge is to be interpreted as players having agreed on what is *ex ante* perceived as possible. If we challenge the *ex ante* perspective (as in Gul (1998), and in Gul and Dekel (1997)), then such concept has to be equally scrutinized.

1.6 Epistemic Conditions for Nash Equilibrium and No Trade

Aumann and Brandenburger (1995) show that a set of conditions, including the CPA, are sufficient conditions for a Nash equilibrium to be played. Such result is still valid with action-consistency instead of the CPA. What is needed for Nash equilibrium is only knowledge of the game and of rationality of the players, and *ex ante* knowledge of the conjectures. Given that Aumann and Brandenburger claimed that their conditions were tight (i.e., could not be relaxed), the above result can be viewed as an addition to the epistemic literature. Milgrom and Stokey (1982) show that, in an exchange economy,

under the CPA, if the initial endowments are ex ante Pareto optimal, and it is common knowledge that a particular trade is feasible and acceptable, then agents are indifferent between this trade and the initial endowment. Again, action-consistency and ex ante knowledge of acceptability are enough.

The no-trade result above suggests that action-consistency can also be used in agreement theorems. Indeed, we can state Aumann (1976)'s agreeing to disagree and Sebenius and Geanakoplos (1983)'s no bets theorems using action-consistency and ex ante knowledge. But differently from the previous results, we can only do so by restricting the class of events that agents are supposed to agree (or not to bet on). Under action-consistency we can only hope to get agreement on events related to actions in the game (and this is in fact achieved if the beliefs over the events are ex ante knowledge).

The results above (and the one presented next) show the scope of action-consistency: there is a class of results that can be stated in their original formulation, but with action-consistency replacing the CPA. This is the class of results that are stated with the CPA in conjunction to an assumption of common knowledge of events that are related to actions in the given strategic situation. Results not in this class can be stated with action-consistency but have to be adapted to action-related events (as it is the case for agreeing to disagree and no bets).

1.7 Action-Consistency, Rationality, Correlated Equilibrium and the Common Prior Debate

Aumann (1987) shows that Bayesian rationality and correlated equilibrium are equivalent under the CPA. Such result is still valid with action-consistency instead of the CPA. As with the other results considered here, we improve our understanding of the nature of Aumann's result. But here there is one additional consideration: Aumann's result is the subject of the common prior debate: Gul (1998) and Gul and Dekel (1997) criticize Aumann's result on the grounds that the common prior cannot be interpreted as an ex ante belief. The fact that Aumann's result can be stated without the CPA partially insulates the result from the critique. The equivalence of rationality and correlated equilibrium holds beyond the CPA.

At this point one might argue that the insulation provided by action-consistency is not enough: after all, we would still be using the action-consistent measure as the correlating measure, so we would still be interpreting the correlating measure as an ex ante belief. The answer to such

argument is twofold. First, first-order beliefs are bona-fide Bayesian beliefs. Since action-consistency is related to consistency of first-order beliefs, it is not subject to the same critique targeted at the CPA. Second, the denseness result shows that action-consistency is a property of almost any given state of the world. It follows that we do not have to assume that players agreed in the ex ante stage: in all likelihood we can view players *as if* they had agreed.

1.8 Restrictiveness of Action-Consistency

Action-consistency is less restrictive than the CPA. It is one property that is hidden by the CPA and that might hold without the full force of the CPA. We show that the set of states of the world that are compatible with action-consistency is a dense subset of the beliefs space. In addition, we show that if the set of actions available to the players (agents) is finite, then every state of the world is compatible with action consistency. Such results are to be contrasted with the fact that the set of states compatible with the CPA is a meager subset of the beliefs space.

Although the results above provide an answer for how less restrictive action-consistency is, there is one detail that must be emphasized. A state is compatible with action-consistency (CPA) if it belongs to the support of one action-consistent measure. In the common prior case, such condition is enough to ensure that players agree (in the interim stage) on the set of possible states of the world. This is not the case with action-consistency. In other words, the fact that a state is action-consistent does not mean that the Bayesian game is necessarily action-consistent. The interim game, determined only by the state of the world, is action-consistent. But the Bayesian game, which is the game viewed from an ex ante perspective, is not necessarily action-consistent.

In what follows we present the material of Sections 1.3 through 1.8 above in detail. In order to do that, we first present the framework and some of the discussions around the CPA in Section 2. Section 4 introduces action-consistency and discusses some of its properties. Section 5 contains a first set of results. There, we start with the equivalence between rationality and correlated equilibrium, and then present the epistemic conditions for Nash equilibrium, no-trade and agreement theorems. Section 3 shows action-consistency is not a restrictive assumption, whereas the CPA is shown to be a restrictive assumption. Some examples are shown in Section 7. An Appendix constructing a belief space under rationality is also provided.

2 Preliminaries

This Section presents the concepts used in the paper. It begins with the beliefs space, which was introduced by Harsanyi (1967, 1968) and formalized by the now classic Mertens and Zamir (1985). It is our basic object, used throughout the paper. Of great importance for our concern is one particular subset, that of consistent states. This is the set of states that are in the support of a consistent probability measure. Because of this, we recall the concept of a consistent probability measure and present some ways of characterizing it. In particular, we argue that a consistent probability measure can be viewed as a fixed point of a class of stochastic operators. This approach was taken by Samet (1998) in the context of finite belief spaces, and here it is applied to a general beliefs space. Such an approach proves to be quite useful in Section 3, when we deal with the size of the subset of consistent states.

Having dealt with consistency, we go on to present the common prior debate mentioned above, which illustrates that there is no consensus in the literature when it comes to interpreting the CPA. We anticipated above that consistency is less restrictive than CPA. Accordingly, here we build on that and argue that what consistency really entails is a property which we term common conditional prior. With respect to the common prior debate, common conditional prior falls in a middle ground: it allows for differences in opinions as well as differences in information, while keeping with the notion of consistency.

2.1 The Beliefs Space

Consider a game of incomplete information composed of a finite set I of players, who have private information about the realization of some aspect of the game. One of Harsanyi's key findings is that each player thinks about what other players think. That is, in order to describe the beliefs of the players in such a game, it is not enough to consider only beliefs over the uncertain aspect of the game. One must also consider what each player believes about the beliefs of his opponents. But then we have another round: we must consider what each player believes about the beliefs each of his opponents about the beliefs of the opponents of each of them. That is, we must consider hierarchies of beliefs about beliefs. This is the problem solved by Mertens and Zamir (1985): they presented a (universal) way of dealing with such hierarchies.

More precisely, they showed that for any given set K representing the uncertainty associated with a given game, one can construct the space of hierarchies of beliefs about beliefs associated with K (see Appendix A). That is, each player is assumed to have a belief about the uncertainty he faces, and this is represented by a regular probability measure over K . As every player will then have such beliefs, each such player will face another round of uncertain events, which is what his opponents believed in the first place. This is again represented by a regular probability measure over the set of probability measures found previously. And so on. The resulting space can be shown to be homeomorphic to $\Omega = K \times \prod_{i \in I} \Theta^i$, with $\Theta^i = \Delta(K \times \prod_{j \neq i} \Theta^j)$, where $\Delta(X)$ denotes the space of regular probability measures on X endowed with the weak-* topology.

The basic domain of uncertainty, K , includes the exogenous uncertainty (states of nature, K_0) and the set of possible payoff functions $U = \prod_{i \in I} U^i$, where $U^i = \{u^i : A \rightarrow \mathbb{R}\}$ and $A = \prod_{i \in I} A^i$ is the joint action set. That is, $K = K_0 \times U$. Here, since we are interested in applications, we will consider one additional ingredient: the joint action set, A , can also be viewed as part of the basic domain of uncertainty. That is, we will use $K = K_0 \times U \times A$, and obtain a description of the game that not only describes the beliefs held by the players, but also the joint action taken by them.⁴

The space Ω is viewed as the space of states of the world, where each state is a tuple consisting of one element of the basic domain of uncertainty and a list of regular probability measures $(t_\omega^i)_{i \in I}$,

$$\omega = (k_\omega, t_\omega^1, \dots, t_\omega^I).$$

Each $t_\omega^i \in \Delta(K \times \prod_{j \neq i} \Theta^j)$ is viewed as the type of player i at ω (a belief over K and the types of the opponents). One defines then the type partition \mathcal{H}^i , where each $H^i \in \mathcal{H}^i$ is of the form $H_\omega^i = \{\hat{\omega} \in \Omega : t_\omega^i = t_\omega^i\}$. Given the Borel σ -field Σ induced by K on Ω (the topological properties of K are inherited by Ω) one computes the sub- σ -field generated by \mathcal{H}^i for each player, which will also be denoted by \mathcal{H}^i . Players have the obvious private information over U and A : each player knows his action $a^i \in A^i$ and his payoff function $u^i \in U^i$. This means that the projections a_ω^i and u_ω^i are measurable with respect to player i 's type partition \mathcal{H}^i .⁵

⁴We could have used $K = K_0 \times U$ together with a mapping $f : \Omega \rightarrow A$ describing the state-contingent joint action.

⁵It follows that we are considering the subspace of the universal belief space where the above kind of private information is respected. See Mertens and Zamir (1985), page 14.

In other words, Ω formalizes Harsanyi's space of type-agents. The interpretation of Ω is fairly intuitive: each state $\omega \in \Omega$ is one realization of the relevant uncertainty associated with the game: one type for each player is realized, and one point in K is realized as well. Since the types represent all beliefs about beliefs, we have a representation of all the relevant uncertainty. In addition Ω has a product structure that defines naturally the type partitions that represent the information structure: given ω , player i knows his type, so considers that any state in H_ω^i might have taken place.

Following Mertens and Zamir, K (and a fortiori Ω) is assumed compact and Hausdorff.⁶

What about consistency (or the CPA)? It is a property that may or may not be satisfied by a given state in Ω . That is, given $\omega \in \Omega$, the list of probability distributions $(t_\omega^1, \dots, t_\omega^I)$ is consistent if each one of these measures can be computed as the conditional probability distribution of a single probability distribution, where the conditioning is on each player's type partition. Clearly this need not be the case. For instance, if every state that one player considers possible is not considered possible by some other player, then these two players cannot be consistent with each other. In fact, there is no common prior in a situation like that.

In formal terms:⁷

Definition 1 *A consistent probability measure over Ω is given by $P \in \Delta(\Omega)$ such that*

$$P(E) = \int t_\omega^i(E)P(d\omega) \tag{1}$$

for every $i \in I$ and every event $E \in \Sigma$.

Definition 2 *A state ω is called **consistent** if it belongs to the support of one consistent probability measure.*

In other words, a consistent P is a particular average of the player's types. A state ω is consistent with the CPA if it belongs to the support of a common prior.

⁶Brandenburger and Dekel (1993), Heifetz (1989) and Mertens, Sorin and Zamir (1994) show that the construction carries through without the compactness assumption. The first assumes that K is Polish, while the other two only assume that K is Hausdorff.

⁷See Mertens and Zamir (1985), Section 4.

We conclude this presentation of the beliefs space with the concept of a beliefs subspace. To understand it, consider for example a two-player game. For a given state of the world ω , player 1 considers every state in $Supp(t_\omega^1)$ possible. But every state therein is associated with player 2 being of a particular type; so player 1 also considers that every state in $Supp(t_\omega^2)$, for every $\omega \in Supp(t_\omega^1)$, is possible. The same reasoning can be applied for player 2. And so on. The end of this construction is a subset of Ω where both players agree, in the sense that both consider the same set of states to be possible. In other words, the subset just created has the property that if a state belongs to subset, then both $Supp(t_\omega^1)$ and $Supp(t_\omega^2)$ belong to the subset.

More formally:

Definition 3 *A beliefs subspace is a non-empty subset Y of Ω such that $\omega \in Y \Rightarrow Supp(t_\omega^i) \subset Y$ for every $i \in I$.*

That is, a beliefs subspace is a simultaneously self-supporting set. It is clear that the intersection of beliefs subspaces is again a belief subspace, so one can construct minimal subspaces satisfying a given property by computing the intersection of all beliefs subspaces satisfying that property. For any given state ω , let Y_ω denote the minimal beliefs subspace generated by ω . Such subspace is common knowledge among the players.

That is, the crucial use of beliefs subspaces is related to knowledge, where knowledge is defined as follows. An event $E \in \Sigma$ is known by player i at state ω if $t_\omega^i(E) = 1$. The set of states where i knows E is $K^i(E) = \{\omega : t_\omega^i(E) = 1\}$. The intersection of these sets is the set of states where E is mutual knowledge: $K(E) = \bigcap_{i \in I} K^i(E)$. The set of states where E is common knowledge is given by $CK(E) = K(E) \cap K(K(E)) \cap K(K(K(E))) \cap \dots$. One can show that $CK(Y_\omega) = Y_\omega$ for each minimal beliefs subspace Y_ω .

The stochastic operators perspective offers another way of understanding minimal beliefs subspaces. For any given stochastic operator, say T , an ergodic class is defined as the support of a minimal invariant measure (a measure P such that $P = TP$, and such that there is no other invariant measure whose support is included in $Supp(P)$). Once we consider the tuple $(T^i)_{i \in I}$ of operators, we are led to the concept of simultaneously ergodic classes: those classes that are simultaneously ergodic for each T^i . It is clear that simultaneously ergodic classes are minimal beliefs subspaces.

And it follows that if ω is consistent, then $Supp(P)$ is common knowledge among the players. Also, given an event $E \in \Sigma$, if there is a common prior P such that $P(E) = 1$, then E is common knowledge. Such perspective will be used to define ex ante knowledge in Section 4.

2.2 The Common Prior Debate

The relation between the CPA and consistency has generated a heated debate in the literature. The plausibility of having individuals have the same ex ante likelihood assessment is not easy to defend. But it is also not easy to come up with a case against it. Recently, a more specific critique has been raised about the plausibility of interpreting the common prior assumption in a beliefs space setting. Gul (1998) and Dekel and Gul (1997) are the ones that first came forward with that critique. Given that Gul advanced this view in a discussion with Aumann (see Gul (1998) and Aumann (1998)), let's oversimplify and say that there are two extreme positions: that of Aumann, and that of Gul.

Although Aumann does provide compelling arguments in favor of the CPA, what is crucial in his perspective is the use he makes of the CPA. Under the CPA, he shows that there is a strong connection between rationality and equilibria. In particular, in Aumann (1987) he shows that an objective equilibrium, called correlated equilibrium, is an expression of rationality under the CPA. The common prior is viewed as the correlating measure (the probability distribution used by players to correlate their actions). One issue is whether such view is justified or not. Gul (1998) doesn't think so.

The idea is as follows. The beliefs space is constructed as the solution of the Bayesian circle of beliefs about beliefs, which by definition takes into account every prior that the players might have had. A consistent probability measure on the beliefs space is just that - one probability measure that can be used to compute the types of the players as its conditional measures. It does not generate the types, it just characterizes a particular set of them. And it is not a prior in the Bayesian sense, since it is not a representation of a preference ordering over a set of acts (non negative random variables) over the beliefs space. Hence it cannot be interpreted as a prior, and its very denomination - common prior - is bound to be misleading. Hence, Aumann's use of the common prior as a correlating measure is not justified. It is supposed to be a measure that the players agreed upon before playing the game. Since they didn't agree on that measure, and it is not even a belief, how can we interpret it as a correlating measure?

An easy answer is that even though there was no agreement, the consistent measure does lend itself to the interpretation that it was as if players did agree on it. More than that, it is formally a correlating measure, regardless of where it came from.

A more involved answer is to advocate for the ex ante view of the game. True, a consistent measure is not a prior. But it is after all a probability measure over a state space. We are used to call such entities beliefs whenever they represent preference orderings. And in those cases we do so because it is as if that measure was used as a subjective probability by the decision maker. Well, a consistent measure can also be viewed as representing a self-referential belief, a belief about how the world functions and how the player fits that scheme. It is an *as if* argument all the same.

In short, neither view can be easily dismissed. In most parts of this paper we use the CPA as a mathematical property characterizing a given state (set of states). Nonetheless, whichever side one takes, the debate itself reveals how multifaceted the issue is.

3 Common Prior is Restrictive

The debate above shows one particular critique to the CPA. In this Section another critique will be presented: we will show that the set of states consistent with the CPA (i.e., the set of consistent states) is a negligible subset of Ω . It follows that the CPA eliminates almost all strategic situations, which makes it a very restrictive assumption.

Theorem 1 *The set of consistent states is a meager subset of Ω .*

Proof. It is plain that the set of consistent states in Ω has empty interior. All we have to do is to show that it is closed.

Let M be the set of continuous functions defined on Ω that are simultaneously invariant. That is, $M = \{f \in C(\Omega) : f(\omega) = \int f dt_\omega^i \text{ for every } i \in I\}$. Let \mathcal{D} be the collection $\bigcap_{f \in M} f^{-1}$ of subsets of Ω . Since \mathcal{D} is generated by the level sets of the continuous mapping $\varphi : \Omega \rightarrow \prod_{f \in M} f(\Omega)$, defined by $\varphi(\omega) = (f(\omega))_{f \in M}$, \mathcal{D} is an upper semicontinuous partition of the compact and Hausdorff space Ω .

A subset $F \subset \Omega$ is simultaneously self supporting if $\omega \in F$ implies that $Supp(t_\omega^i) \subset F$ for every $i \in I$. It is standard to show that each such F

contains a minimal simultaneously self supporting set, and that each such minimal F is the support of a simultaneously invariant probability measure P . Such P is a common prior. It follows that each minimal F is contained in exactly one $D \in \mathcal{D}$. The sets D that support common priors are called consistent sets, and will be denoted by E . Let C denote the collection of such sets, that is, C is the set of consistent states.

Now pick a sequence (ω_n) in C such that $\omega_n \rightarrow \omega$. Since the quotient mapping $\phi : \Omega \rightarrow \Omega/\mathcal{D}$ is continuous, $\phi(\omega_n) \rightarrow \phi(\omega)$. Let U be any open set containing $E = \phi(\omega)$. Since \mathcal{D} is an usc partition, there exists an open set $V \subset U$ such that $E \subset V$ and every element D of \mathcal{D} that intersects V is contained in U . This means that the sequence of consistent sets $E_n = \phi(\omega_n)$ is eventually inside U , so we can focus on this sequence of consistent sets. Now let P_n be any common prior on E_n . Since $\Delta(\Omega)$ is compact, and the set of common priors is closed, there exists a convergent subsequence $P_{n_k} \rightarrow P$, where P is a common prior. The support of P is included in U , which is arbitrary. Therefore $Supp(P) \subset E$, E is a consistent set, and ω is a consistent state. It follows that C is closed. ■

In itself, this result is to be expected. In fact, several authors have mentioned that consistency is not the norm.⁸ Nyarko (1991) argued that most games violate Harsanyi's Doctrine, which is to say that most games are inconsistent. More formally, in the context of a finite type space, he shows that the set of consistent states is a subset of measure zero. The result above confirms Nyarko's findings, but in the context of a general beliefs space.

4 Action-Consistency

In the previous Section we presented the framework of analysis, examined some of the controversies associated with the CPA, and provided yet another critique to the CPA. It is now time to introduce the concept of action-consistency. Action-consistency is weaker than the CPA as it is (immediately) implied by the latter. In Section 5 we show that action-consistency is what is really needed when the CPA is used to produce a certain class of results. This is the ultimate motivation for the concept: that is, action-consistency is the minimal consistency requirements for the validity of a class of results. In what follows we will present the formal definition of action-consistency, providing an alternative motivation for it.

⁸Mertens and Zamir (1985) argue informally, via examples, that most beliefs subspaces are inconsistent. The same informal argument is found in Heifetz (2001).

Given that the beliefs space Ω can be identified with $K \times \prod_{i \in I} \Theta^i$, which has a cartesian product structure, there is a natural way to relax consistency: assume that it holds only for certain components of $K \times \prod_{i \in I} \Theta^i = (A \times U \times K_0) \times \prod_{i \in I} \Theta^i$, in a sense to be defined shortly. One natural candidate is K , given that it represents the basic domain of uncertainty, and some form of consistency on K would be interpreted as consistency of beliefs over the given uncertainty. Another candidate is A , the exogenous joint strategy set, since consistency therein would be interpreted as consistency of beliefs over the joint action of the group of players. The results in Section 5 suggest that the choice of A is meaningful for some applications: it reveals the crucial property that is implicit in the CPA. This choice leads to the concept of action-consistency. On the other hand, there might be situations where we need consistency over $A \times U$.

To assume consistency on a component of Ω means to assume that the marginal distributions of types over that component satisfy some form of consistency. More precisely:

Definition 4 *An **action-consistent probability measure** over Ω is given by $\mu \in \Delta(\Omega)$ such that*

$$\mu([E]) = \int t_\omega^i([E])\mu(d\omega), \quad (2)$$

for all events $E \in \Sigma_A$ and all $i \in I$, where Σ_A is the σ -algebra over A induced by Σ and $[E] = \{\omega \in \Omega : a_\omega \in E\}$.

Clearly, if μ is a consistent measure, then it is also an action-consistent measure since $[E] \in \Sigma$. Below we will also consider a case where consistency over $A \times U$ is required. This comes from simply having $[E] = \{\omega \in \Omega : (a_\omega, u_\omega) \in E\}$ in (2) above. In any case, action-consistency refers to (2) over events defined by (at most) K . For most cases below, we will use only events defined by the joint action set A , and this justifies the term “action-consistency”.⁹

As it was the case with consistency, the fact that $t_\omega^i([E])\chi_H = t_\omega^i([E] \cap H)$

⁹One could use the term K -consistency to denote consistency on K or on any of its components in general.

for every $H \in \mathcal{H}^i$ together with (2) implies that

$$\begin{aligned} \int_H t_\omega^i([E])\mu(d\omega) &= \int t_\omega^i([E])\chi_H\mu(d\omega) \\ &= \int t_\omega^i([E] \cap H)\mu(d\omega) = \mu([E] \cap H) \\ &= \int_H \chi_{[E]}\mu(d\omega), \end{aligned}$$

so, by Radon-Nikodým,

$$\mu([E]|\mathcal{H}^i)(\omega) = t_\omega^i([E]), \mu\text{-a.e.}, \quad (3)$$

for all events $E \in \Sigma_A$ and all $i \in I$. Hence, action-consistency means that the marginal (over A) of the type of player i can be computed as the marginal (over A) of the conditional probability of μ given \mathcal{H}^i , for every player $i \in I$. In that sense, what is required is much less than consistency, since the types of the players might be inconsistent while their marginal distributions satisfy (2).¹⁰

Another similarity with consistency is that action-consistency is also characterized by the intersection of closed convex hulls, but now of the marginal distributions of the types. The argument is analogous: the marginal of t_ω^i is an invariant measure of the stochastic operator T^i restricted to measures defined on the events $[E]$ for $E \in \Sigma_A$.

$$t_\omega^i([E]) = \int t_v^i([E])t_\omega^i(dv),$$

for every $E \in \Sigma_A$ and every player. Therefore, an action-consistent μ is a convex combination of the marginal distributions of $(t_\omega^i)_{\omega \in \Omega}$. Hence, calling those marginal distributions m_ω^i , we have that an action-consistent measure must belong to

$$\bigcap_{i \in I} G_A^i,$$

where $G_A^i = \overline{\text{co}}(m_\omega^i)_{\omega \in \Omega}$. It is immediate that G_A^i is the projection of G^i (defined above) over A : while consistency required that the sets G^i themselves

¹⁰Note that action-consistency is *not* the same as having an information partition on A (or K) and a “common prior” there, and then considering the set of distributions over Ω with the given “common prior” as the marginal of those distributions over A (or K).

had non empty intersection, action-consistency only requires that their projections over A intersect.¹¹

The above characterization also reveals that action-consistency is related only to first-order beliefs. At each state $\omega \in \Omega$, the marginal distribution of t_ω^i over K is by construction player i 's first-order beliefs. The marginal m_ω^i over A is just the marginal distribution of the first order beliefs over the joint action set. It follows that action-consistency only restricts the first element of each hierarchy of beliefs about beliefs. As mentioned above, it does not mean that higher order beliefs are not relevant, for they do constrain first order beliefs.

The intuition behind action-consistency is simple: players might disagree badly on events that are not directly related to the actions in the game, as long as they agree on the likelihood of the actions. One player might have an absurd theory from the viewpoint of his opponent, but as long as this theory delivers the same result in terms of $\Delta(A)$, then these two players are indeed consistent to a certain extent. In a sense, it is consistency at “face-value”, since one does not require players to go deep into the roots of the conjectures over A .

Going back to the example of the two investors considered earlier, say that the game they play after the announcement by the central banker is that of taking a position in an asset, and that it is an equilibrium to have one going short and the other long on the asset. It is clear that it does not matter why one investor thought that the other would do the opposite of what he did: one of them might use advanced financial models to help his decision, while the other just looks at the stars and somehow decides what the other must have thought. As long as they are right in their conjectures about the opponent's behavior, we have an equilibrium, which reflects the fact that the investors did entertain some form of consistency.

Another example is the game played by New Yorkers every day when they walk on Manhattan's sidewalks. It is a convention in the city that everyone walks on the right side (at least in Downtown Manhattan). One has to know virtually nothing about the (type of the) opponents, just that they are New Yorkers, to conclude that they will take the action of walking on the right side. In formal terms, New Yorkers are action-consistent in that particular game. This example shows that action-consistency might be particularly plausible for large games (i.e., a large number of players), where consistency on every event (the CPA) is a rather strong requirement.

¹¹This gives the necessary geometric intuition for Proposition 6 below.

4.1 Ex Ante Knowledge

Since action-consistency is not related to all orders of beliefs about beliefs, its relation with common knowledge is not as tight as the CPA's. That is, we saw above that if P is a common prior, then the event $Supp(P)$ is common knowledge. It follows that if E is an event such that $P(E) = 1$, then E is common knowledge. For action-consistency, the analogous concept is that of ex ante knowledge:

Definition 5 *Given an action-consistent μ we say that an event $E \subset \Omega$ is **ex ante knowledge** for μ if $\mu(E) = 1$.*

The term ex ante knowledge refers to the idea that it is as if players had agreed in ex ante terms that the event E would take place. It is not the usual notion of knowledge, which is an interim concept. But it is the notion of knowledge that can be characterized with action-consistency. It is clear that ex ante knowledge is not common knowledge, for action-consistency only deals with first-order beliefs, and common knowledge is a property of the whole hierarchy of beliefs. It is also clear that it is a less restrictive requirement: when the action-consistent μ is a full-fledged common prior, and E is common knowledge, then $\mu(E) = 1$, so that E is also ex ante knowledge.

Given that in Section 5 we will present a class of results that can be stated with action-consistency and ex ante knowledge in the place of common prior and common knowledge, the above remark shows that there is an added generality stemming from action-consistency: the well-known results considered here do not even require common knowledge assumptions.

Apart from not being the usual notion of knowledge, the concept of ex ante knowledge raises two additional issues: that of the relation of action-consistency and Bayesian games and that of the ex ante perspective. The latter was discussed before: the ex ante stage is not modelled, so to argue that players ex ante agreed on some event might not be a sensible statement. Yet again, one can view it as an *as if* statement, or argue that the fact that the ex ante stage is not modelled does not preclude us from making statements about what might be viewed as happening before the interim stage.

4.1.1 Bayesian Game \times Interim Game

The relation between action-consistency and Bayesian games is related to the relation between common knowledge and ex ante knowledge. A Bayesian game is given by $\langle I, A, T, p, u \rangle$, where I and A are as before, $T = \prod_{i \in I} T^i$ is the set of types, $p = (p^i)_{i \in I}$ is a vector of prior distributions and $u = (u^i)_{i \in I}$ is a vector of payoff functions. Hence, in the context of a beliefs space Ω , a Bayesian game is defined by a beliefs subspace, which defines the set of types T . Under the CPA, the fact that $Supp(P)$ defines a beliefs subspace, implies that $Supp(P)$ also defines a Bayesian game (indeed, a consistent Bayesian game). This is not the case with action-consistency, for $Supp(\mu)$ need not be a beliefs subspace. The ex ante knowledge event associated with μ is not a Bayesian game, whereas it would be if used a common prior (and the common knowledge event associated with it).

In Section 6 below we argue that (almost) every state is action-consistent. This means that the game played at the interim stage is action-consistent, so we can say that most (interim) games are action-consistent. But the discussion above shows that this need not be the case for the associated Bayesian game. It is not necessarily the case that a particular beliefs subspace will be contained in the support of an action-consistent measure. Whenever this is the case, the Bayesian game is action-consistent.

5 Using Action-Consistency instead of CPA

In this Section we show that action-consistency is what is needed to establish a class of well-known results. We present two theorems in games and one in general equilibrium that can be stated with action-consistency replacing the CPA. We also present two agreement theorems that make explicit the scope of action-consistency.

5.1 Bayesian Characterization of Correlated Equilibrium

Aumann (1987) provided a strong connection between rationality and equilibria. In that paper, Aumann showed that, under the CPA, Bayesian rationality and correlated equilibrium are equivalent, which is a much stronger result than what was known at that time, that is, that Bayesian rationality implied only a subjective form of equilibrium. Aumann's result states that equilibrium behavior is an unavoidable consequence of the assumption of rationality. Later such statement was criticized by Gul (1998) (see Section

2.2 above) Here we show that the driving force leading to Aumann's result is action-consistency. The CPA is not needed. At once, this provides an answer to Gul's critique.

To begin with, let us define the concept of correlated equilibrium.

Definition 6 *Given a correlating space F (taken to be an exogenous measure space) with information partitions \mathcal{F}^i for each player $i \in I$, and a tuple of probability distributions $(P^i)_{i \in I}$ on $\Delta(F)$, called subjective priors, a joint action $f : F \rightarrow A$ is a **subjective correlated equilibrium** if, for all ω ,*

$$\int (u^i(f(v)) - u^i(a^i, f^{-i}(v)))P^i(dv|\mathcal{F}^i)(\omega) \geq 0 \quad (4)$$

for every $i \in I$ and every $a^i \in A^i$, where $u^i : A \rightarrow \mathbb{R}$ is player i 's payoff function, and $f^i : F \rightarrow A^i$ is \mathcal{F}^i -measurable.

If $P^i = P$ for every i then we have an **objective correlated equilibrium**.

As in Aumann (1987), the payoff functions are known (that is, U is a singleton, and the payoff functions are u^1, \dots, u^I). The following result is an immediate consequence of the definition of action-consistency (a random variable $f : \Omega \rightarrow R$ is A -measurable if it is constant on each $[E]$, for each $E \in \Sigma_A$).

Lemma 1 *If a random variable $f : \Omega \rightarrow \mathbb{R}$ is A -measurable and μ is action-consistent then*

$$\int f(v)t_\omega^i(dv) = \int f(v)\mu(dv|\mathcal{H}^i)(\omega)$$

for all $i \in I$ and $\omega \in \Omega$.

Proof. The result is true for $f = \chi_{[E]}$, any $E \in \Sigma_A$ by (3) above. Hence it is also true for all simple functions, and the result follows from the Monotone Convergence Theorem. ■

In the Appendix we present the implications of the assumption of Bayesian rationality. In particular, we derive the event Ω^R where all players are rational, and show that for every such state therein, $\omega \in \Omega^R$,

$$\int (u^i(a_v) - u^i(a^i, a_v^{-i}))t_\omega^i(dv) \geq 0, \quad (5)$$

for all $i \in I$ and $a^i \in A^i$.

We can now state Aumann's proposition in its full generality:

Proposition 1 *Given an action-consistent $\mu \in \Delta(\Omega)$, consider the event Ω^R where all players are rational, and assume that Ω^R is ex ante knowledge for μ . Then the projections $a_\omega : \Omega \rightarrow A$ form a correlated equilibrium under the correlating space (Ω^R, μ) .*

Proof. Put $f^{a^i}(\omega) = u^i(a_\omega) - u^i(a^i, a_\omega^{-i})$ for each $a^i \in A^i$, and note that f^{a^i} is A -measurable. From Lemma 1 and (5), for all $\omega \in \Omega^R$,

$$\int f^{a^i}(v) \mu(dv | \mathcal{H}^i)(\omega) = \int f^{a^i}(v) t_\omega^i(dv) \geq 0.$$

for every $i \in I$ and $a^i \in A^i$, which shows that μ acts as a correlating measure over Ω^R (cf. (4)). ■

Hence it is as if players agreed to play according to μ before knowing their types and then decided to stick to it after the realization of ω , in the sense that no player wanted to deviate to any available action a^i . The critique to such a result was mentioned before: μ is not a Bayesian belief, and so one is not allowed to treat it as such and impose ex ante agreement. In any case, it is plain that μ acts as a correlating measure, and that we are entitled to state that it is as if players reached an agreement before playing, even though we know they didn't necessarily reach an agreement. In addition, one can still invoke the self-referential (ex ante) view and call μ a belief, and use the logic of the common prior (with the added generality that now the common prior refers only to events on A).

What is relevant is that, given the results in Section 3, we can go one step further and say that Bayesian rationality is enough for an objective equilibrium. Aumann provided the link between rationality and equilibrium having to use one additional assumption (the CPA). Here it is shown that such additional assumption is not needed in general, and rationality is in fact almost enough for a form of objective equilibrium. As we show in Section 6, if the action set is finite then rationality is indeed enough for an objective equilibrium, for every state of the world is in the support of an action-consistent μ .

5.2 Epistemic Conditions for Nash Equilibrium

In a now classic paper, Aumann and Brandenburger (1995) provided sufficient epistemic conditions for a group of players to play a Nash equilibrium. They showed that, in a finite game, if the state of the world satisfies common knowledge of conjectures, knowledge of rationality and of the game, and consistency then players are necessarily part of a Nash equilibrium. They went further and, by means of several examples, showed that their conditions could not be improved upon. In particular, they used their Example 5.3 to argue that the CPA was indispensable. It turns out, though, that it is dispensable to a certain extent. We show below that action-consistency is enough for their argument.

Following Aumann and Brandenburger, we assume that A is finite. As before, we use $[\]$ to denote an event, so that $[a]$ denotes the event $\{\omega \in \Omega : a_\omega = a\}$ ($[a^{-i}]$ and $[a^i]$ are defined analogously). At a given state $\omega \in \Omega$, player i 's conjecture about the joint action taken by his opponents is the marginal of t_ω^i over A^{-i} , and it will be denoted by ϕ_ω^i (that is, $\phi_\omega^i(a^{-i}) = t_\omega^i([a^{-i}])$). The crucial step leading to the claimed generalization of Aumann and Brandenburger's result is stated in the following Lemma.

Lemma 2 *Let $\mu([E]) = \int t_\omega^i([E])\mu(d\omega)$ for all $E \subset A$ and all $i \in I$. If $\phi_\omega^i = \phi^i$ for all $\omega \in \text{Supp}(\mu)$, then $\mu([a]) = \prod_{i \in I} \mu([a^i])$ for all $a \in A$.*

Proof. Putting $E = [a^{-i}]$,

$$\mu([a^{-i}]) = \int t_\omega^i([a^{-i}])\mu(d\omega) = \int \phi^i(a^{-i})\mu(d\omega) = \phi^i(a^{-i}).$$

Given that a_ω^i is \mathcal{H}^i -measurable, it follows that the full support of t_ω^i is either in $[a^i]$ or disjoint from $[a^i]$, so that $t_\omega^i([a]) = t_\omega^i([a^{-i}])\chi_{[a^i]}(\omega)$. Hence

$$\mu([a]) = \int t_\omega^i([a])\mu(d\omega) = \int \phi^i(a^{-i})\chi_{[a^i]}(\omega)\mu(d\omega) = \phi^i(a^{-i})\mu([a^i]).$$

Summing up, $\mu([a]) = \mu([a^{-i}])\mu([a^i])$. The result now follows by induction on $i \in I$. ■

We are now in position to show that one can use action-consistency instead of the CPA to get sufficient conditions for Nash equilibrium.

Proposition 2 *Suppose that at the given state of the world there is an action-consistent μ such that the event $\{\omega : \phi_\omega^i = \phi^i \text{ for every } i \in I\}$ is ex ante knowledge for μ . Suppose also that the game u and rationality are mutually known at the given state. Then for each player j , all the conjectures ϕ^i of players i other than j induce the same conjecture σ^j about j , and $(\sigma^1, \dots, \sigma^I)$ is a Nash Equilibrium of u .*

Proof. By assumption $\phi_\omega^i = \phi^i$ for all $\omega \in \text{Supp}(\mu)$, so Lemma 2 shows that the stated conditions imply that $\mu([a]) = \prod_{i \in I} \mu([a^i])$ for all $a \in A$. Putting $\sigma^i(a^i) = \mu([a^i])$, we get that $\phi^i(a^j) = \sigma^j(a^j)$ for all $i \neq j$, and $\phi^i(a^{-i}) = \prod_{j \neq i} \sigma^j(a^j)$. Given mutual knowledge of the payoffs and of rationality, it follows that each action a^j with $\phi^i(a^j) > 0$ for some $i \neq j$ maximizes expected value of u^j with respect to ϕ^j , and a fortiori also with respect to $\prod_{k \neq j} \sigma^k(a^k)$. Hence $(\sigma^1, \dots, \sigma^I)$ is a Nash Equilibrium of u . ■

The novelty of Proposition 2 above is, of course, the use of action-consistency instead of the CPA, which can be viewed as an addition to the epistemic literature. Another addition is the use of ex ante knowledge instead of common knowledge of the conjectures, since the former is less restrictive than the latter.

Another important issue comes when we use the results of Section 6 below. Since A is finite, Corollary 2 below shows that every state of the world is action-consistent. Hence, if one accepts the assumption of rationality and of knowledge of the payoff functions, then the only important condition for Nash equilibrium is that the conjectures be ex ante knowledge.

5.3 No Trade

Now consider an economy with I agents and L goods. Each agent has an utility function $u^i : \mathbb{R}^L \rightarrow \mathbb{R}$, and an endowment vector $e_\omega^i \in \mathbb{R}_+^L$ for each ω . A trading plan $z = (z^1, \dots, z^I)$, where $z^i : \Omega \rightarrow \mathbb{R}^L$ is feasible if $e_\omega^i + z_\omega^i \geq 0$ for every i and ω and $\sum_{i \in I} z_\omega^i \leq 0$ for every ω . The trading plan is acceptable at ω if

$$\int (u^i(e_v^i + z_v^i) - u^i(e_v^i)) t_\omega^i(dv) \geq 0.$$

Let A be the set of pairs of feasible trading plans and endowments: that is, A is the set of (z, e) where z is feasible and $e = (e^1, \dots, e^I)$. Given an action-consistent μ , we say that the initial endowment is ex ante Pareto optimal

if

$$\int (u^i(e_v^i + z_v^i) - u^i(e_v^i))\mu(dv) \leq 0,$$

for every i and every feasible trading plan z . Milgrom and Stokey (1982) showed that if it is common knowledge that there is an acceptable trading plan for all agents, and there is a common prior, then each agent is indifferent between the trading plan and the zero trade, if the initial endowment is ex ante Pareto optimal (hence they do not trade). Such result needs only action-consistency and ex ante knowledge of acceptability. To prove this claim, we need the following result:

Lemma 3 *Let f be an A -measurable random variable, and let μ be an action-consistent probability distribution. Then*

$$\mu(f) = \mu(t^i(f))$$

where $\mu(f) = \int f(\omega)\mu(d\omega)$, and $t_\omega^i(f) = \int f(v)t_\omega^i(dv)$, so that $t^i(f)$ is the random variable that takes the value $t_\omega^i(f)$ at ω , and $\mu(t^i(f))$ is just the integral of such random variable with respect to μ .

Proof. The result is true if $f = \chi_{[E]}$ for $E \subset A$ by definition of action consistency. But then it is also true for simple functions, and the result follows from the Monotone Convergence Theorem. ■

Note that, using A as above, $u^i(e^i + z^i) - u^i(e^i)$ is an A -measurable random variable.

Proposition 3 *Let μ be an action-consistent probability distribution and suppose that $e = (e^1, \dots, e^I)$ is ex ante Pareto optimal with respect to μ . Suppose it is ex ante knowledge for μ at ω that some feasible trading plan z is acceptable to all agents. Then each agent is indifferent between z and the zero trade.*

Proof. Let $f^i = u^i(e^i + z^i) - u^i(e^i)$. Since z is ex ante Pareto optimal,

$$\mu(f^i) \leq 0,$$

and it is also acceptable in the support of μ :

$$\mu(t^i(f^i)) \geq 0.$$

From Lemma 3,

$$0 \geq \mu(f^i) = \mu(t^i(f^i)) \geq 0$$

Hence $\mu(t^i(f^i)) = 0$ for every agent, which means that $t_i(f^i) = 0$ in the support of μ . But this is simply

$$\int u^i(e_v^i + z_v^i)t_\omega^i(dv) = \int u^i(e_v^i)t_\omega^i(dv),$$

which means that z is indifferent to the zero trade. ■

5.4 Agreement on Actions - How Far We Can Go With Action-Consistency

Above we were able to prove three well-known results with the assumption of action-consistency in the place of the CPA. This shows that action-consistency is sufficient to establish a certain class of results. Now we wish to identify such a class. To do so, we present a form of Aumann (1976)'s agreeing to disagree result, which is now stated in terms of agreement on events related to actions, and also the no-bets generalization, again stated in terms of bets on actions. Such results characterize action-consistency (much in the same way that their original formulations characterize the CPA), and also show the limits of action-consistency: the class of results that do not require the CPA in its full force are the results that are based on knowledge of events related to the actions that the agents can take. Since the agreement theorems are originally related to any event in the state space, they cannot be fully restated in their original formulations with action-consistency. Yet, if for a particular problem what is at issue is agreement on events related to actions, then the versions of the agreement theorems that can be stated with action-consistency are enough to prevent disagreement.

5.4.1 Agreeing to Disagree

Aumann (1976) proved that it cannot be common knowledge that agents disagree if they have a common prior. Hence agents cannot agree to disagree. When we restrict to action-related events, replacing the CPA by action-consistency brings about a subtle distinction in the formulation of the agreement theorem. This is related to our previous distinction between knowledge and ex ante knowledge. In fact, the reformulation of the theorem for action-related events is that agents cannot agree to disagree on these

events. But now the first “agree” refers to ex ante knowledge instead of common knowledge. The second “agree” is as in Aumann: it is considered in terms of expected values of some event, so that agents i and j agree (at ω) on their opinions with respect to an event E if $t_\omega^i(E) = t_\omega^j(E)$.

Proposition 4 *If the expected values of an event $[E]$, $E \subset A$, are ex ante knowledge for an action-consistent μ (i.e., if $t_\omega^i([E]) = p_i$, for every agent i , is ex ante knowledge for μ), then $p_i = p$.*

Proof. Since μ is action-consistent,

$$\mu([E]) = \int t_\omega^i([E])\mu(d\omega) = p_i,$$

so that $p_i = p = \mu([E])$. ■

Hence, agents cannot agree to disagree on the likelihood of events related to actions in the game, if they are action-consistent, and if ‘to agree’ means to be ex ante knowledge.

5.4.2 No Bets

The agreement idea can be extended to bets, which are random variables defined on the state space. An agent i is said to accept the bet $f : \Omega \rightarrow \mathbb{R}$ at ω if $t_\omega^i(f) \geq 0$. Sebenius and Geanakoplos (1983) show that if it is common knowledge that agents are willing to enter a zero-sum bet then each agent expects to get zero out of the bet (i.e., each agent cannot expect to get a strictly positive return from the bet). We say that a bet f is on A if it is A -measurable: it is a random variable $f : \Omega \rightarrow \mathbb{R}$ which is constant on each event $[a] \subset \Omega$, $a \in A$. A group of agents is willing to bet on A if there are I random variables $f^i : \Omega \rightarrow \mathbb{R}$, all of them A -measurable, such that $\sum_{i \in I} f^i = 0$ and $\int f(v)t_\omega^i(dv) \geq 0$ for every $i \in I$.

Proposition 5 *If it is ex ante knowledge for an action-consistent μ that agents are willing to bet on A , then all of them expect to get zero out of the bets (i.e., if $\sum_{i \in I} f^i = 0$ and $t_\omega^i(f^i) \geq 0$, $\forall i$, are ex ante knowledge for μ , then $t_\omega^i(f^i) = 0$, $\forall i$).*

Proof. For a given i , using Lemma 3, we have that

$$\mu\left(\sum_{i \in I} t^i(f^i)\right) = \sum_{i \in I} \mu(t^i(f^i)) = \sum_{i \in I} \mu(f^i) = \mu\left(\sum_{i \in I} f^i\right) = 0,$$

which means that $\sum_{i \in I} t^i(f^i) = 0$ in the support of μ . But each $t^i(f^i)$ is a non-negative random variable, so that $t^i_\omega(f^i) = 0$ for every $\omega \in \text{Supp}(\mu)$. ■

5.4.3 No Trade Revisited

Proposition 3 above was stated under the assumption that U is a singleton. That is, the payoff functions were assumed known to every agent. In Milgrom and Stokey's formulation, the payoff functions might depend on the state of the world as well. Following the ideas above, we can restate their theorem by assuming more than action-consistency (but still less than consistency). That is, assume that (2) holds for every event $E \in \Sigma_A \otimes \Sigma_U$, where A is as before and U is the set of payoff functions. Then $u^i(e^i + z^i) - u^i(e^i)$ is an $A \times U$ -measurable random variable, and Milgrom and Stokey's theorem can be stated in their original form, but with the above notion of consistency instead of the CPA, and also only with ex ante knowledge of acceptability.

5.5 Summing Up

The five results presented in Propositions 1 through 5 above make it clear that there is a class of results that do not depend on the CPA in its full generality. It is the class of results where some action-related event is assumed common knowledge. For such results, what the CPA does is to ensure consistency of ex ante beliefs in terms of actions, but this is precisely what action-consistency means. It follows that for such a class of results we can use action-consistency instead of the CPA. For the class of results where events related to payoff functions are also needed, we can use consistency on $A \times U$ as we did above.

6 Action-Consistent States

Above we introduced the concept of action-consistency as the relevant feature of the CPA, at least for a certain class of results. We also argued that it is a plausible condition, in the sense that it does not impose equality of opinions, just consistency of beliefs over actions. Here we show that action

consistency is a weak assumption. The set of states compatible with action-consistency is a dense subset of Ω , in general. And if A is finite, then every state in Ω is compatible with action-consistency. In view of the result in Section 3 above, action-consistency is much weaker than the CPA.

First, recall that a state is compatible with the CPA if it is consistent, i.e., if it belongs to the support of a consistent probability measure. Analogously,

Definition 7 *A state $\omega \in \Omega$ is **action-consistent** if $\omega \in \text{Supp}(\mu)$ for some action-consistent probability distribution μ .*

The definition above comes straight from the concept of a consistent state. Nevertheless, while a consistent state is contained in a minimal beliefs subspace that is common knowledge among the players, an action-consistent state does not share this property. Indeed, this is the reason why we used ex ante knowledge instead of common knowledge. Below we discuss this issue in more detail.

We can now determine whether action-consistency is a restrictive assumption or not. The idea is that, since the set of action-consistent states is a subset of Ω , the restrictiveness of the assumption can be determined by the size of this subset. If it is a “large” subset, then most strategic situations will be action-consistent, so that the assumption can be regarded as not restrictive, or weak. Likewise, it will be a strong assumption if the subset is “small”. The following results show that the set of action-consistent states can be viewed as a large subset: it is in general a dense subset of Ω , and if A is finite, in particular, it is Ω itself.

Proposition 6 *Let Y be a finite beliefs subspace. Then every $\omega \in Y$ is action-consistent.*

Proof. Let $\{a_1, \dots, a_N\}$ be the finite set of joint actions associated with Y . Let $Y = \{\omega_1, \dots, \omega_M\}$, and construct an $IN \times M$ matrix \mathbf{H} , with entry $((i, n), m)$ given by $t_{\omega_m}^i([a_n]) - \chi_{[a_n]}(\omega_m)$. We will show that we can find an action-consistent probability distribution that puts positive mass on any given $\omega \in Y$, say, on ω_1 . Consider an $(IN + 1)$ -dimensional vector b of zeros in all entries except for the last one, and append to the matrix \mathbf{H} an $(IN + 1)$ th row with one in the first entry and zeros elsewhere. It is then clear that we are looking for an M -dimensional vector μ such that

$\mathbf{H}\mu = b$ and $\mu \geq 0$. This means that we want to make sure that b belongs to the cone generated by the columns of \mathbf{H} . But since we can increase the number of columns, M , by considering other states (with the same finite set of associated joint actions, $\{a_1, \dots, a_N\}$) as much as we please, we can always choose the new states in a way that ensures that b belongs to the cone generated by the columns of H . ■

With the above result in hands, the following two corollaries are immediate.

Corollary 1 *The set of action-consistent states is dense in Ω .*

Proof. This follows immediately from Mertens and Zamir's Theorem 3.1, which asserts that, for each $\omega \in \Omega$, and any neighborhood U of ω , there exists a $\bar{\omega}$ in U that belongs to a finite beliefs subspace. Since this $\bar{\omega}$ is action-consistent, we are done. ■

In addition, the following is also a consequence of Proposition 6:

Corollary 2 *If A is finite, then every $\omega \in \Omega$ is action-consistent.*

Proof. Just repeat the proof of Proposition 6 using $A = \{a_1, \dots, a_N\}$. ■

The three results above show clearly that action-consistency is a weak assumption. In an infinite belief space, almost every state, even the ones on inconsistent subspaces, can be viewed as belonging to the support of a suitable action-consistent measure. This does not mean that this measure contains a given subspace in its support. Only the given state is in the support. Examples usually deal with finite belief subspaces, and it is quite simple to come up with such examples where the given subspace violates action-consistency. But this does not mean that one cannot find other subsets of the associated infinite belief space which can be used to construct a suitable action-consistent measure containing any given state in its support.

In terms of applications, all we need is action-consistency coupled with ex ante knowledge of some action related event. Such procedure works, as we showed above, but it also shows that the connection between common prior and common knowledge is lost when we use action-consistency. We can always ensure that the game, at the interim stage, is action-consistent. But we cannot ensure that the Bayesian game (from the ex ante perspective)

is action-consistent. Whenever the connection between action-consistency and common knowledge exists, the Bayesian game is action-consistent.

The CPA is used virtually everywhere in information economics and games of incomplete information, if not for other reason, because of its usefulness. But not only it is an assumption that is subject to philosophical criticisms, it is also a restrictive assumption. Hence, to stick to the CPA is, to a large extent, to choose usefulness over plausibility. Action-consistency, in a precise sense, eliminates the need for such a choice: it is not a restrictive assumption, and it is (almost) as useful as the CPA. Although one can still criticize the ex ante aspects of action-consistency, the view provided by Corollaries 1 and 2 allows us to use it without having to defend its plausibility. This is the main contribution of this paper.

7 Examples

Example 1 Consider the a game with two players, where each of them can be of 4 types as in table below

	t_{21}	t_{22}	t_{23}	t_{24}
t_{11}	$1, \frac{1}{2}$	$0, 0$	$0, 0$	$0, 0$
t_{12}	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, 1$	$0, 0$	$0, 0$
t_{13}	$0, 0$	$0, 0$	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{4}$
t_{14}	$0, 0$	$0, 0$	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{3}{4}$

Each row corresponds to a type of player 1 and each column to a type of player 2. Hence, at state (t_{11}, t_{21}) , player 1 believes that only (t_{11}, t_{21}) might take place while player 2 believes that there is a 50/50 chance of states (t_{11}, t_{21}) and (t_{12}, t_{21}) taking place. And so on. Clearly there are two belief subspaces, $C_1 = ((t_{11}, t_{21}), (t_{12}, t_{21}), (t_{12}, t_{22}))$ and $C_2 = ((t_{13}, t_{23}), (t_{13}, t_{24}), (t_{14}, t_{23}), (t_{14}, t_{24}))$. For any state in C_1 (C_2), both players construct the same minimal beliefs subspaces C_1 (C_2). If the state lies outside of both C_1 and C_2 then one player will construct C_1 and the other C_2 as the minimal beliefs subspaces. This is clearly an inconsistent situation. Moreover, C_2 is an inconsistent subspace since there is no measure that generates the types at C_2 as conditional measures. Or, the convex hulls of the types defined at

C_2 are disjoint. Hence, there is only one consistent subspace, C_1 , and the unique consistent measure puts probability $\frac{1}{3}$ to states (t_{11}, t_{21}) , (t_{12}, t_{21}) , and (t_{12}, t_{22}) .

Example 2 Now consider the 2×2 game presented above, with the following payoff matrix

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

and assume that if player 1 is of type t_{11} then he plays U (and he plays D for the other 3 types), and if player 2 is of type t_{21} then she plays L (and she plays R for the other 3 types). For the types described in table T_1 , if the state of the world is on C_1 (the consistent subspace), then the unique consistent measure defines a correlated equilibrium, the one where each of the three joint actions $((UL), (DL), (DR))$ is played with probability $1/3$. This is readily verified to satisfy the required inequalities. Alternatively, it also follows from Proposition 1 because the assignment of actions to types above satisfied Bayesian rationality. Moreover, if the state of the world is on the inconsistent subspace C_2 , then although there is no consistent distribution over C_2 , the marginal distributions of the types over A do agree (types at C_2 play DR with probability one), so we have action-consistency. Since Bayesian rationality is also satisfied at C_2 , we have a correlated equilibrium distribution. In fact, at C_2 players play the Nash equilibrium DR , and this is because the conjectures are constant in the support of an action-consistent distribution: both players know that DR will be played, and the fact that they disagree on the likelihood of states of the world is not relevant for the game since they agree on what is payoff-relevant for the game. That's one example of the revised version of Aumann and Brandenburger's theorem.

Example 3 Another example of the validity of Aumann and Brandenburger's theorem under action-consistency is given by the following 3 player game. There are three players, the row, the column and the matrix player, and the payoff matrices are given by

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

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

The unique Nash equilibrium is given by $(\sigma_{Row}, \sigma_{Column}, \sigma_{Matrix}) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$.
The types are represented by the following four matrices,

	L ₁	R ₁	L ₂	R ₂
U ₁	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{8}, \frac{1}{4}$	0, 0, 0	0, 0, 0
D ₁	$\frac{1}{8}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	0, 0, 0	0, 0, 0
U ₂	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
D ₂	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
	W ₁			

	L ₁	R ₁	L ₂	R ₂
U ₁	$\frac{1}{4}, \frac{1}{4}, \frac{1}{8}$	$\frac{1}{4}, \frac{3}{8}, \frac{3}{8}$	0, 0, 0	0, 0, 0
D ₁	$\frac{3}{8}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	0, 0, 0	0, 0, 0
U ₂	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
D ₂	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
	E ₁			

	L ₁	R ₁	L ₂	R ₂
U ₁	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
D ₁	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
U ₂	0, 0, 0	0, 0, 0	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$
D ₂	0, 0, 0	0, 0, 0	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$
	W ₂			

	L ₁	R ₁	L ₂	R ₂
U ₁	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
D ₁	0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
U ₂	0, 0, 0	0, 0, 0	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$
D ₂	0, 0, 0	0, 0, 0	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{1}{4}, \frac{1}{4}$
	E ₂			

where again the action assigned to each type is represented by the letter associated with the type (type W_1 means that the matrix player plays W , and so on). It is clear that for the types indexed by 1 the theories are not consistent (there is no common prior with any such state in the support), while the types indexed by 2 are consistent (there is a common prior that puts probability $\frac{1}{8}$ to every state in the support of those types). The Nash equilibrium $(\sigma_{Row}, \sigma_{Column}, \sigma_{Matrix}) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is played when the state of the world is (U_1, L_1, W_1) , but since there is no common prior at such state, Aumann and Brandenburger's conditions would not identify such equilibrium.

The fact that there is no common prior with (U_1, L_1, W_1) in the support does not mean that there is no action-consistent distribution with (U_1, L_1, W_1) in the support. In fact, a probability distribution that puts probability $\frac{1}{8}$ for the states $((U_1, L_1, W_1), (U_2, R_2, W_2), (U_2, L_2, E_2), (U_2, R_2, E_2), (D_2, L_2, W_2), (D_2, L_2, E_2), (D_2, R_2, W_2),$ and (D_2, R_2, E_2) , and zero for all other states is easily seen to be action-consistent. And conjectures are ex ante knowledge, for each player thinks that the opponents will play each joint action with probability $\frac{1}{4}$. Given that we have (common) knowledge of the game and of rationality, the conditions in Proposition 2 above are satisfied, and using them one can conclude that players necessarily play a Nash equilibrium.

Example 4 Aumann and Brandenburger (1995) use their example 5.3 to illustrate a situation where the CPA fails. Here we show how the failure of the CPA does not imply failure of action-consistency, by embedding their example into a larger subset of the beliefs space. There are three players, and each has two actions available: the row player can take actions H and T , the column player can take actions h and t , and the matrix player can take actions W and E . Aumann and Brandenburger use the following types:

	h_1	t_1
H_1	$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$
T_1	$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$
	W_1	

where the subscript 1 refers to the type that play that action (i.e., H_1 refers to the type of the row player that takes action H). Clearly, there is no common prior, and no action-consistent measure defined only in that subspace. But this subspace is a subset of Ω . For instance, let's consider that each player can be of two types, indexed by 1 and 2 (i.e., H_1 and H_2 refer to the types of the row player that take action H), and that the following tables represent those types:

		h_1	t_1	h_2	t_2			h_1	t_1	h_2	t_2
H_1		$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$	0, 0, 0	0, 0, 0	H_1		0, 0, 0	0, 0, $\frac{1}{2}$	0, 0, 0	0, 0, 0
T_1		$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$	0, 0, 0	0, 0, 0	T_1		0, 0, 0	0, 0, $\frac{1}{2}$	0, 0, 0	0, 0, 0
H_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0	H_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
T_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0	T_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
W_1						E_1					
		h_1	t_1	h_2	t_2			h_1	t_1	h_2	t_2
H_1		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0	H_1		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
T_1		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0	T_1		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, 0
H_2		0, 0, 0	0, 0, 0	$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$	H_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, $\frac{1}{2}$
T_2		0, 0, 0	0, 0, 0	$\frac{1}{2}, \frac{1}{2}, \frac{3}{8}$	$\frac{1}{2}, \frac{1}{2}, \frac{1}{8}$	T_2		0, 0, 0	0, 0, 0	0, 0, 0	0, 0, $\frac{1}{2}$
W_2						E_2					

There are 64 states, and states (H_1, h_1, W_1) , (H_1, t_1, W_1) , (T_1, h_1, W_1) , and (T_1, t_1, W_1) represent Aumann and Brandenburger's subspace, now embedded in a larger subspace. None of the 64 states is consistent. But the distribution that puts probability $\frac{1}{4}$ to each of the following states, (H_1, h_1, W_1) , (T_1, t_1, W_1) , (T_1, h_1, W_2) , and (T_1, t_2, W_2) is an action-consistent probability distribution.

A The Belief Space under Rationality

Here we illustrate the workings of Mertens and Zamir's construction of the belief space for the particular case that $K = A$, and under the assumption of rationality. The resulting space can be viewed as the relevant state space when Bayesian rationality is the maintained hypothesis.

More precisely, the uncertainty that player i faces is given by A^{-i} , and hence player i has a prior on that set. For each such measure $t_1^i \in \Delta(A^{-i})$, rationality states that player i associate the best reply $a^i(t_1^i)$ given by

$$a^i(t_1^i) \in \arg \max_{\hat{a}^i} \int_{A^{-i}} u^i(\hat{a}^i, a^{-i}) dt_1^i(a^{-i}). \quad (6)$$

As player i performs the computation above for every possible $t_1^i \in \Delta(A^{-i})$, she forms a collection $(a^i(t_1^i), t_1^i)_{t_1^i \in \Delta(A^{-i})}$, which is a subset of $A^i \times \Delta(A^{-i})$.

For ease of notation, consider $\Delta(A^{-i})$ as a subset of $\Delta(A)$ with the given marginal distributions over A^{-i} . For the group of players, hence, we formed a subset of $A \times \Delta(A)^I$.

Now, each player also reasons about what the opponents might have believed in the first round. For player i , this is represented by a probability measure t_2^i on the relevant subset of $A^{-i} \times \Delta(A)^{I-1}$ from the first round of the opponents. Notice that defining $a^{-i} : \Delta(A)^{I-1} \rightarrow A^{-i}$ as the vector $(a^j(t_1^j))_{j \neq i}$ where $a^j(t_1^j)$ is given by (6), we have

$$\pi_{A^{-i}}^{-1} = \pi_{\Delta(A)^{I-1}}^{-1} \circ (a^{-i})^{-1}, \quad (7)$$

where $\pi_{A^{-i}}$ and $\pi_{\Delta(A)^{I-1}}$ are the projections of $A^{-i} \times \Delta(A)^{I-1}$ on its factors.

Once we consider the second round of beliefs of all players we form a subset of $A \times \Delta(A)^I \times (A \times \Delta(A)^I)^I$. The hierarchy of mutual beliefs goes on. It can be stated concisely as follows:

$$\begin{aligned} X_0 &= A \\ T_k &= \Delta(X_{k-1}) \\ X_k &= X_{k-1} \times (T_k)^I \text{ for all } k \geq 1, \end{aligned}$$

where it is understood that at each round of the hierarchy only the relevant subset is considered.

A hierarchy of beliefs is said to be coherent if the following two conditions are satisfied for all $i \in I$:

$$t^i(x_k) \circ \pi_{k,k-2}^{-1} = t^i(x_{k-1}) \quad (8)$$

$$t^i(x_k) \circ \pi_{k,\Delta(k-2)}^{-1} = \delta_{t^i(x_{k-1})}, \quad (9)$$

where $t^i(x_k)$ is player i 's likelihood assessment at $x_k \in X_k$ (an element of T_k), $\pi_{k,k-2}$ is the projection of X_k onto X_{k-2} , $\pi_{k,\Delta(k-2)}$ is the projection of X_k onto the I -th copy of T_{k-1} and δ is the Dirac measure. In words, a coherent hierarchy is one where each player knows her previous beliefs.

Remark 1 From (7) and (8), it follows that in a coherent hierarchy of beliefs the first order belief t_1^i must respect the marginal of t_2^i on $\Delta(A)^{I-1}$, in the sense that the induced distribution on A^{-i} , given by $t_2^i \circ \pi_{\Delta(A)^{I-1}}^{-1} \circ (a^{-i})^{-1}$, must be t_1^i .

The sequence of sets $(X_k)_{k \geq 1}$ satisfying the coherence conditions, together with the projections $\pi_{k,l}$ of X_k onto X_l , $l < k$, defines a projective system of sets. Given that A is finite, and using the weak-* topology on sets of measures¹², the sets in the system are all compact, and hence the projective limit $\Omega^R = \varprojlim X_k$ is non empty. Each thread $\omega \in \Omega$ is of the form

$$\omega = (a^1, \dots, a^I, t_1^1, \dots, t_1^I, t_2^1, \dots, t_2^I, \dots, t_k^1, \dots, t_k^I, \dots),$$

where t_k^i stands for $t^i(x_k)$. The collection $(t_k^i)_{k=1}^\infty$ can be represented by a probability measure on Ω^R , called t_ω^i , such that

$$t_k^i = t^i \circ \pi_{\omega,k}^{-1},$$

where $\pi_{\omega,k}$ is the projection of Ω^R onto X_k .¹³ t_ω^i is player i 's type at ω .

Hence, a belief space associated with a game is given by Ω^R . The way Ω^R was constructed was to respect rationality at each round, in the sense that no player would consider a situation where some player(s) would behave irrationally. From Remark 1, the coherence conditions imply that higher order beliefs constrain lower order beliefs. In particular, (iterated) elimination of dominated strategies is built in the construction.

At each $\omega \in \Omega^R$ each player has a prescription of which action to take, namely the projection on A^i , or the i -th component of ω . This action is optimal by construction since a^i is the best reply to the beliefs. Moreover, the projection $a_\omega \in A^{\Omega^R}$ of Ω^R onto A is also the best joint action taken when Ω^R is viewed as the state space. Let a_ω^i and a_ω^{-i} be the projections onto A^i and A^{-i} , and notice that a_ω^i is \mathcal{H}^i -measurable by construction.

Proposition 7 *Let $a_\omega \in A^{\Omega^R}$ be the projection of Ω^R onto A . Then, for each $\omega \in \Omega^R$ and each $i \in I$,*

$$\int u^i(f(v)) t_\omega^i(dv)$$

is maximized when $f = a_\omega$.

¹²All measures are assumed to be regular Borel measures, i.e., Radon.

¹³This result was proved in Mertens and Zamir (1985) using the Riesz-Markov theorem. It can also be viewed as a special case of Bochner's generalization of Kolmogorov's existence theorem (and for that we need the regularity of the measures).

Proof. Using Monotone Convergence Theorem again, we have

$$\int_{\Omega} u^i(a_v) t_{\omega}^i(dv) = \int_A u^i(a) m_{\omega}^i(da),$$

where $m_{\omega}^i \in \Delta(A)$ is the marginal of t_{ω}^i over A , i.e., it's given by $m_{\omega}^i(E) = t_{\omega}^i(a^{-1}(E))$ for every event $E \subset A$. But m_{ω}^i is t_1^1 at ω , so that the right hand side is maximized by construction (since the associated a^i is the best reply $a^i(t_1^i)$) ■

In particular, a^i is also the best reply to $t_{\omega}^i([a^{-i}])$, i.e., to player i 's conjecture about the opponents's behavior.

A.1 Bayesian Rationality in Ω

The construction above shows that the set of states in Ω where each player is rational is given by the subset Ω^R . That is, in the present setting with K_0 and U as singletons, a state of the world $\omega = (a_{\omega}, t_{\omega}^1, \dots, t_{\omega}^I)$ specifies a joint action $a = a_{\omega}$ for the tuple of players $t_{\omega}^1, \dots, t_{\omega}^I$. The i th component of a_{ω} does not have to be the best reply to t_{ω}^i . The beliefs subspace where a_{ω}^i is indeed the action chosen rationally by type t_{ω}^i for every i is the subspace where it is common knowledge that players are rational, which is Ω^R . It is, hence, just the minimal beliefs subspace generated by those very events. From Proposition 7 we conclude that for every $\omega \in \Omega^R$,

$$\int (u^i(a_v) - u^i(a^i, a_v^{-i})) t_{\omega}^i(dv) \geq 0,$$

for all $i \in I$ and $a^i \in A^i$, and $u^i : A \rightarrow \mathbb{R}$ is player i 's payoff function. Under common knowledge of Bayesian rationality, no player wants to deviate to some action other than the one given by the corresponding component of ω .

It is important to note the descriptive character of the beliefs space. One constructs a space that represents every possible situation, including situations where the players are irrational (in the sense of choosing actions that are not the best choices given their beliefs). Choice here is not what will be done given the space Ω : it is determined already by each state ω , by the very descriptive nature of the beliefs space. This almost mechanical flavor is not to be mistakenly thought of as lack of purposeful action. At any given interim game, players are choosing purposefully, and this situation is represented as a single state ω .

References

- [1] **Aumann, Robert 1976.** “Agreeing to Disagree”, *Annals of Statistics*, **4**: 1236-1239.
- [2] **Aumann, Robert 1987.** “Correlated Equilibrium as an Expression of Bayesian Rationality”, *Econometrica*, **55**: 1-18.
- [3] **Aumann, Robert 1998.** “Common Prior: A Reply to Gul”, *Econometrica*, **66**: 929-938.
- [4] **Aumann, Robert and Adam Brandenburger 1995.** “Epistemic Conditions for Nash Equilibrium”, *Econometrica*, **63**: 1161-1180.
- [5] **Barelli, Paulo 2003.** “An Asset Market Associated with a Game”, mimeo, Columbia University.
- [6] **Brandenburger, Adam and Eddie Dekel 1987.** “Rationalizability and Correlated Equilibrium”, *Econometrica*, **55**: 1391-1402.
- [7] **Brandenburger, Adam and Eddie Dekel 1993.** “Hierarchies of Beliefs and Common Knowledge”, *Journal of Economic Theory*, **59**: 189-198.
- [8] **Brandenburger, Adam, Eddie Dekel and John Geanakoplos 1992.** “Correlated Equilibrium with Generalized Information Structures”, *Games and Economic Behavior*, **4**: 182-201.
- [9] **Dekel, Eddie and Faruk Gul 1997.** “Rationality and Knowledge in Game Theory”, in Kreps and Wallis (eds.), *Advances in Economic Theory, Seventh World Congress*, Cambridge: Cambridge University Press.
- [10] **Feinberg, Yossi 2000.** “Characterizing Common Priors in the Form of Posteriors”, *Journal of Economic Theory*, **91**: 127-179.
- [11] **Gul, Faruk 1998.** “A Comment on Aumann’s Bayesian View”, *Econometrica*, **66**: 923-927.
- [12] **Harsanyi, John 1967-1968.** “Games of Incomplete Information Played by Bayesian Players”, Parts I-III, *Management Science*, **8**: 159-182, 320-334, 486-502.

- [13] **Heifetz, Aviad 1993.** “The Bayesian Formulation of Incomplete Information: The Non-Compact Case”, *International Journal of Game Theory*, **21**: 329-338.
- [14] **Heifetz, Aviad 2001.** “The Positive Foundations of the Common Prior Assumption”, Mimeo, The California Institute of Technology.
- [15] **Lipman, Barton 2003.** “Finite Order Implications of Common Priors”, *Econometrica*, **71**: 1255-1267.
- [16] **Mertens, Jean-François and Shmuel Zamir 1985.** “Formulation of Bayesian Analysis for Games with Incomplete Information”, *International Journal of Game Theory*, **14**: 1-29.
- [17] **Mertens, Jean-François, Sylvain Sorin and Shmuel Zamir 1994.** “Repeated Games. Part A. Background Material”, CORE Discussion Papers 9420.
- [18] **Milgrom, Paul and Nancy Stokey 1982.** “Information, Trade and Common Knowledge”, *Journal of Economic Theory*, **26**: 17-27.
- [19] **Morris, Stephen 2002.** “Typical Types”, mimeo, Cowles Foundation.
- [20] **Nehring, Klaus 2001.** “Common Priors under Incomplete Information: a Unification”, *Economic Theory*, **18**: 535-553.
- [21] **Nyarko, Yaw 1991.** “Most Games Violate the Harsanyi Doctrine”, New York University Working Paper.
- [22] **Polak, Ben 1999.** “Epistemic Conditions for Nash Equilibrium and Common Knowledge of Rationality”, *Econometrica*, **67**: 673-676.
- [23] **Samet, Dov 1998.** “Iterated Expectations and Common Priors”, *Games and Economic Behavior*, **24**: 131-141.
- [24] **Sebenius and Geanakoplos 1983.** “Don’t Bet on It: Contingent Agreements with Asymmetric Information”, *Journal of the American Statistical Association*, **18**: 424-426.
- [25] **Sine, Robert 1968.** “Geometric Theory of a Single Markov Operator”, *Pacific Journal of Mathematics*, **27**:155-168.
- [26] **Tan, Tommy and Sergio Werlang 1988.** “The Bayesian Foundations of Solution Concepts of Games”, *Journal of Economic Theory*, **45**: 370-391.