Reasoning, games, action and rationality

August 11 - 15, ESSLLI 2008

Lecture 4

Eric Pacuit Olivier Roy

October 5, 2008

1 Lecture 4: Hard Knowledge and Prior Beliefs

1.1 Hard Knowledge and Nash Equilibrium

1.1.1 Mutual Knowledge of strategy

Iterated elimination of strictly dominated strategies is a very intuitive concept, but for many games it does not tell anything about what the players will or should choose. The coordination game of Table **??** is a typical example. Rationality and common belief of rationality does not exclude any strategy choice for Ann or for Bob. Whether choosing **A** is rational for Bob depends on what he expects Ann will choose, and believing that she is rational does not help him to narrow his expectations to **a** or **b**. But what if Bob knows (that is has correct beliefs about) Ann's strategy choice? Intuitively, this is quite clear what he should do, that is what is his best response: coordinate with her. If he knows that she plays **a**, for instance, then playing **A** is clearly the only rational thing for him to choose, and vice versa if he knows that she plays **b**. The situation is entirely symmetrical for Ann. If she is sure that Bob plays **A**, and she is right about it, then her only rational choice is to choose **a**, and vice versa if she knows that Bob plays **B**. In terms of type structure, the only states where Ann is rational and her type knows Bob's strategy choice and where Bob is also rational and his type knows Ann's strategy choices are states where they play either **aA** or **bB**. The same hold for epistemic plausibility models: they play either **aA** or **bB** in all states where Ann and Bob are rational and in all situation they considers possible they play the same strategy.

The strategy profiles **aA** and **bB** are *Nash equilibria* of that game. This solution concept is usually described as the profiles where no player has an incentive to unilaterally deviate from his strategy choice. In other words, a Nash equilibrium is a combination of strategies, one for each player, such they all play their best response given the strategy choices of the others. This is precisely what happens at the profiles **aA** and **bB** in our coordination example. At **aA**, for instance, **a** is the best response of Ann, given that Bob plays **A**, and similarly for Bob.

This key intuition being Nash equilibrium, the idea of "best response given the action of others", is what the mutual knowledge of strategy ensures in the example above. If Ann knows at a state that Bob chooses **A**, she is not only convinced of that, but she is also correct: Bob *is* playing **A** at that state. By playing rationally given her information, Ann is thus playing her best response given Bob's strategy choice. This is not only the case in our small example, but also in general.

Theorem 1.1 *[**?**] For any two-players strategic game and epistemic model for that game, if at state w both players are rational and know each other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.*

This epistemic characterization differs from the one of iterated elimination of dominated strategies in three ways. First, it requires mutual *knowledge* and not only beliefs. It is crucial that each player is right about the strategy choice of the other. Otherwise his response might simply fall off the mark. Second, it only requires *one* level of interactive knowledge. It is not necessary that both players know that they know each other's choices, nor is any higher-order level of information necessary. Finally, the result is stated only for games with *two* players. One can find the general result [**?**].

To add : paragraph on the various modal characterizations of Nash equilibrium.

1.1.2 Announcement of strategy

Nash equilibrium can also be characterized dynamically, much in the flavor of **?** 's approach to iterated elimination of strictly dominated strategies. Just like one only needs one level of interactive knowledge to ensure equilibrium choices, only one announcement is needed to narrow down a relational model to the equilibrium state: the announcement of each player's strategy.

Theorem 1.2 *For any two-players strategic game and relational model for that game, if at state w after announcing player i's strategy choice in the original model j is rational, and after announcing player j's strategy choice in the original model i is rational, then the profile they play at w is a Nash equilibrium.*

Again, there's an intuitive connection between the static and the dynamic characterizations of Nash equilibrium. Each announcement provides the players with all the information they need to make the best decision given the other's strategy choice. Observe, however, that the two announcements of strategy are not simultaneous nor successive. They are both made individually in the original model. A joint announcement of all player's strategy choice trivialize any posterior statement about rationality. It leaves only one possible strategy profile to consider possible for all players, which makes them trivially rational. What is required is rather a "test" of each player's rationality posterior to the announcement of his opponent's choice. This is indeed what captures the idea that each strategy in a Nash equilibrium is a best response, given the choices of others.

The two characterizations of Nash equilibrium rest on mutual knowledge of strategy choice, a strong condition if any. In the coordination game of Table **??**, for instance, it is not clear how Ann or Bob could ever acquire such hard information by looking only at the structure of the game—recall that knowledge is defined here as *true* belief. To be truthfully informed about each other's strategy choices of others they need a lot of additional information, which the announcement in Theorem **??** bluntly provides. More generally, cases where mutual knowledge of strategy choice obtains are exceptional, and characterizations of equilibrium play which use this condition are bound to apply to a limited number of scenarios. One way to get a more general understanding of equilibrium play is to refine the structure of information in game playing situations, by distinguishing *correlated prior beliefs* and *asymmetric information*, as we shall see in the next section.

1.2 Prior beliefs, correlation and equilibrium

In this section we will see how common *prior* beliefs, under common knowledge of rationality, lead to equilibrium play. To introduce the notion of prior belief we will go back to the various stages of information disclosure that we introduced in Section **??**. This will lead to discuss the *common prior assumption* and the related "Harsanyi doctrine", according to which different beliefs should only be accountable in terms of different information. We will then look at correlated beliefs and, after a short digression on mixed strategies we will look at equilibrium play.

1.2.1 Prior and posterior beliefs uncertainty

One can distinguish three stages of information disclosure in games. *Ex ante*, the players have not yet chosen their strategy and have not received any information the choices of others. *Ex interim*, they have made their decisions, but they only have partial information about the decision of others. *Ex post*, all relevant information—strategy choices, outcome, payoffs—is revealed to all players.

So far we focused on the *ex interim* stage, where each player is at least "informed" about his own strategy choice. This is particularly explicit in epistemic plausibility models, where the players' beliefs are *conditional* on a their hard information at a state. The situation is similar type structure, where there is no uncertainty about one's strategy choice and information.

At the *ex ante* stage, on the other hand, *no* information has been yet disclosed. The agents have not made their decision yet, and they have little to no information about the decision and information of others. As we will see shortly, many authors see the *ex ante* stage as a stage where only exogenous uncertainty (see Section **??**) can make some game playing situation more plausible than other for some player.

The distinction between *ex ante* and *ex interim* can be captured via the notions of *prior* and *posterior* probabilities, the familiar to the Bayesian. In type structure, a *prior* probability distribution p_i for player i is simply a probability distribution on the set of states, i.e. on the product $S \times T$ of the sets of strategy and type profiles. Let t'_i be the set of states (σ, t) such that $t_i = t'_i$, and $p_i(t'_i)$ be the sum of the probabilities of all states in t'_i. For simplicity we will assume here that p_i assigns non-zero probability to all type of player i, that is $p_i(t'_i) > 0$ for all types $t'_i \in T_i$. Whenever this condition does not hold one can use the various devices introduced in Section **3!!**. We will say that a type structure $\mathbb{T}$ is *generated* by the set of priors $\{p_i\}_{i \in I}$ whenever the following holds:

$$\lambda_i(t_i)(E) = \frac{p_i(E \cap t_i)}{p_i(t_i)}$$

where $E \cap t_i$ is, slightly abusing our notation, the set of pairs $(s_{-i}, t_{-i}) \in E$ such that there is a $s_i \in S_i$ for which $((s_i, s_{-i}), (t_i, t_{-i})) \in t_i$. In other worlds, a type structure $\mathbb{T}$ is generated by a prior probability whenever the beliefs of each agents at a type t_i are is posterior beliefs conditional on the fact that he is of that type.

This idea can be captured in epistemic plausibility models as well, by looking only at the relations $\leq_i$, without restricting them to the equivalence classes $[w]_i$ of each state. We mentioned in Section **??** that on epistemic plausibility models where the leq_i are locally connected the relations $\sim_i$ become superfluous. To capture the idea of prior beliefs it will be useful, however, to take the relations $\leq_i$ to be *fully* connected, that is to assume that for all w, w' either $w \leq_i w'$ or $w' \leq_i w$, and to take the relation $\sim_i$ as a separate relation such that $w \sim_i w'$ implies $\sigma_i(w) = \sigma_i(w')$. The relation $\leq_i$ then represents the *prior* beliefs of i, and its restrictions to a certain *information cell* $[w]_i$ represents i's *posterior* beliefs, that is his beliefs conditional on having received the information associated with the state w.

The distinction between *ex ante* and *ex interim* stages highlights the role of the private or asymmetric information in assessing the rationality of a strategy choice. This becomes particularly clear when the players starts with *common* prior beliefs, an important assumption in the game-theoretic literature that we now briefly look at.

1.2.2 Common priors

Many authors have argued that differences in posterior beliefs should always be thought as resulting from *different* or *asymmetric information* received by each players during a given game playing situation, and not by differences in prior beliefs [**? ? ?**]. This has led to the idea that the players should have the same *prior beliefs*, in absence of any substantial information about a particular game playing situation. If players end up with different beliefs at the *ex interim* stage, it should be because they have received different information about the actual game playing situation.

This idea, known as the *common prior assumption*, can be easily captured in epistemic plausibility models and in type spaces. In the later it simply boils down to set take the same prior belief p_i for all players i in I. Similarly, in epistemic plausibility models it means taking the same plausibility ordering $\leq_i$ for all players.

It is important to see that common prior beliefs do not implies identical posterior beliefs. Indeed, different private information can result in different posterior beliefs even in contexts where the priors *were* common. The beliefs at each type depicted in Tables **??** and **??**, for instance, can be seen as coming from the common prior probability distribution p over $S \times T$ such that $p((t_{Ann}, t_{Bob}), aA) = p((t_{Ann}, u_{Bob}), aB) = 1/2$ and $p(s) = 0$ for all other states. Despite this common ground, Ann and Bob have different beliefs about Bob's type at state $((t_{Ann}, t_{Bob}), aA)$. Bob is indeed certain of his own information,

$$\lambda_{Bob}(t_{Bob})(t_{Bob}) = \frac{p(t_{Bob})}{p(t_{Bob})} = 1$$

but Ann is not certain about Bob's:

$$\lambda_{Ann}(t_{Ann})(t_{Bob}) = \frac{p(t_{Bob} \cap t_{Ann})}{p(t_{Ann})} = 1/2$$

The same can be said about the epistemic plausibility model of Figures **??** and **??**. One can just take the union of $\leq_{Ann}$ and $\leq_{Bob}$ as inducing Ann and Bob's (common) prior beliefs, and recover their posterior beliefs by taking the respective restriction of this relation to the partitions depicted in Figure **??**.

The common prior assumption is, as we mentioned, strongly associated with the idea that difference in posterior beliefs should be understood in terms of difference in information. This thesis originated from the work of Harsanyi on games with *incomplete* information, that is games where the players are uncertain about each other's payoffs, and has become known as the *Harsanyi doctrine*. Harsanyi indeed showed in [**?**] that, provided some consistency constraints on a set of (posterior) beliefs, one can see them as if they were coming from common priors.

The common prior assumption and the associated Harsanyi doctrine are not uncontroversial in the game-theoretic literature, especially when it concerns games of incomplete information (see [**?** **?**]). Morris [**?**] offers an extensive survey of the arguments in its favor, the most convincing being probably that if the common prior assumption is dropped, almost every strategy choice can be seen as rational. This idea is central in Aumann's justification of Bayesian rationality in games, to which we turn after a short digression on correlated beliefs and mixed strategies.

See [**?**] for more on the common prior assumption, from a logical point of view. See also [**?**] for an interesting note on such logical characterizations.

1.2.3 Correlated and non-correlated prior beliefs

The definition of common prior that we presented above does not exclude cases where the expectations of one player are correlated to his own strategy choice. That is, can happen that the players consider that the choices of the other players are not independent of his own strategy choice.

Suppose, for instance, that the following are Ann and Bob's (common) prior beliefs for the coordination game in Table **??**, and that Ann and Bob have only one type, t_{Ann} and t_{Bob}, respectively. By generating the type struc-

	A	B
a	1/2	0
b	0	1/2

Table 1: Correlated prior beliefs

ture from this common prior, one gets that Ann knows that whenever she chooses **a** Bob chooses **A**, and similarly when she chooses **b**. In other words, conditional on herself choosing **a**, Ann *knows* that Bob will choose **A**, and conditional on herself choosing **b** she knows that Bob will choose **B**. Ann and Bob's beliefs and strategy choices are *correlated* in this case.

In general, a probabilistic common prior distribution p is correlated when it cannot be factorized in a set of independent probability distributions p_i, one for each player, on the set of strategy choices and possible types of others.

Definition 1.3 (Correlated common prior - the probabilistic case) A common prior probability distribution p is *correlated* when there is no set of probability distributions $\{p_i\}_{i \in I}$ such that $p_i = \Delta(S_{-i}, T_{-i})$ and for all state $s \in S \times T$, $p(s) = \Pi_{i \in I} p_i(s_{-i}, t_{-i})$. $\triangleleft$

Correlated beliefs have been extensively discussed in the literature, especially questions regarding the source of correlations [**? ?**]. They have also been applied to a wide range of game-theoretic "puzzles", from coordination games [**?**] the Prisoner's Dilemma [**?**] and, as we will see shortly, they allow for a natural generalization of the idea of equilibrium in games.

All of this concerns the probabilistic version of correlated beliefs. How to define them in epistemic plausibility models is still an open question to us (although [**?** , p.99] might lead the way here).

1.2.4 Pure and mixed strategies

So far we only worked with so-called *pure* strategies, that is elements of the strategy sets S_i for each agent i, leaving aside *mixed strategies*. A mixed strategy ρ for player i is a probability distribution $\Delta(S_i)$ on his set of strategies, and the *support* of a mixed strategy ρ is the set of pure strategies which get assigned a non-zero probability.

Mixed strategies are not only crucial in Aumann's view on Bayesian rationality in games, which we will look at shortly, but are also key ingredients of mainstream game theory. Nash's celebrated result that every game in strategic form as an equilibrium [**?**], for instance, rests on mixed strategies. In many games, such as *Matching Pennies* (see [**?** , p.17]), there is no Nash equilibrium in pure strategy. All equilibrium of that games involve randomizations on the strategy sets S_i.

There is no consensus, however, on how to interpret mixed strategies. Some authors propose to take them as object of choices, meaning that to play a mixed strategy literally means randomizing over one's own strategy set. Others, see for instance [**?**], propose to look at a mixed strategy ρ_i of player i as the beliefs or the expectations of the other players $j \neq i$ about what i's choices. See [**?** , 37-44] for a review of the various interpretations of mixed strategies.

Furthermore, while mixed strategy can be naturally included in type structures, it is not clear how to render them in epistemic plausibility models. In particular, to include them one would have to revise the definition of rationality that we introduced on page **??**, and it is still open to us how to do it meaningfully.

For the remaining of this section, we will thus focus on type structures, which also provide a natural environment for the interpretation of mixed strategy as beliefs of others.

1.2.5 "Rational" expectations and correlated equilibrium