Lesson 5: Common knowledge and agreeing to disagree

5.1 Common knowledge: an example

Aumann, in '87, qualifies correlated equilibria as expression of bayesian rationality.
To understand and discuss the content of Aumann's assertion, we need some further technical tools.
We need a richer language. I will begin with an example:
You are a DM

$$\text{What you get depends on: } \begin{cases} a \in A & \leftarrow & \text{your action} \\ \omega \in \Omega & \leftarrow & \text{true state of nature} \end{cases}$$

Typical problem for decision under incertanty.
You must choose a before knowing ω

Consider an example. You are offered to bet on the result of a throw of couple of dice.
You will gain G if the sum of dice is 8, and you will pay L otherwise.
Here a reasonable (not the unique which is possible or reasonable) representation of Ω is $\Omega = \{1, \ldots, 6\}^2$.
With $p(i, j) = \frac{1}{36} \quad \forall (i, j) \in \Omega$

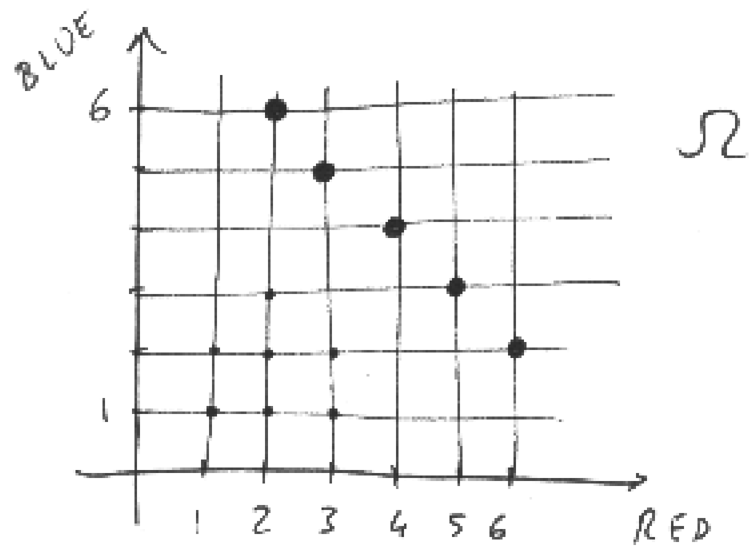

Assume that you are simply an expected money maximizer. That is, your vN-M utility function is (linear with) money.

TO BET: $\frac{5}{36}G - \frac{31}{36}L$
NOT TO BET: 0

So, you will bet if $\frac{5}{36}G \geq 31L\ldots$
Notice that your choice is *NOT CONDITIONED* upon ω. Obviously. You don't know ω. We are obeying to some minimal realism assumption.
Would be different if you knew ω.
Of course, if $\omega \in \{(2,6),(3,5),\ldots,(6,2)\}$ then you would "bet", getting G.
Otherwise you would not bet, getting 0.
Not serious.

But, there are interesting "intermediate" cases.
For example, you could be allowed to see the result of die 1 just before betting.
This means that you have partial information.
Or that you have an information partition
$P = \{\{(1,2),\ldots,(1,6)\},\{(2,1),\ldots,(2,6)\},\ldots,\{(6,1),\ldots,(6,6)\}\} = \{P_1,\ldots,P_6\}$.
This (info partition) is a standard tool.
The interpretation is obvious. If ω is the true state of nature, the DM knows only $P(\omega)$, the element of the partition to which ω belongs.
So, the action of DM can be *contingent* on $P(\omega)$.
Of course, to decide, the DM will re-compute the probability distribution based on his partial information.

2 cases:
$\omega = (1,j)$, i.e. we are in P_1.
The probability that 8 obtains is zero, so:

TO BET: $-L$
NOT TO BET: 0.

$\omega = (k,j)$ with $k \neq 1$; i.e. we are in P_k, $k = 2,\ldots,6$. The probability that 8 obtains is $\frac{1}{6}$. So:

TO BET: $\frac{1}{6}G - \frac{5}{6}L$
NOT TO BET: 0.

All of this with just one DM.

If the DM are two (or more)? Clearly, the key issue here is that they *may have different (partial) information.*
For example, DM_2 could know the result of the "second" die.
So, he has a different information partition.

$P_2 = \{\{(1,1),\dots,(6,1)\},\{(1,2),\dots,(6,2)\},\dots,\{(1,6),\dots,(6,6)\}\}$
For example, if the true ω is $(1,3)$, w.r.t. to the bet:
1 assignes $prob = 0$ to the event E that the sum of dice is 8.
2 assignes $prob = \frac{1}{6}$ to the same event E.
So, if we have that $\frac{1}{6}G - \frac{5}{6}L > 0, DM_2$ will bet, while 1 not.

Nothing strange…

Notice the following.
If 2 KNOWS that 1 assignes $prob = 0$ to the event E, than 2 will revise his probability assessments! He understands that event E is actually impossible, analyzing the information received from the probability assessment of I.
NOTICE that for this to happen, it is essential that player 2 *KNOWS* P_1, the info partition of 1 (or that, at least, has some info about that).
So, the fact that 1 and 2 have different believes about E, *cannot* be a shared, a common information.
This is the key point of Aumann's "agreeing to disagree".

Please, notice that this was just a simple example. In particular, it was enough for player 2 to know the probability assigned by 1. One can construct more sophisticated examples, with more elaborate knowledge interactions.
I will turn now to a very sketchy introduction to the formalism of CK, just to have the minimal instruments for understanding both "agreeing to disagree" and "correlated equilibria as expression of bayesian rationality".

5.2 Connections with subjective equilibria

The simple example that we have seen shows how two DMs may have different probability assessments about an event, just because they have different (partial) information.
Notice that this was exactly the assumption that we needed for the subjective equilibria.
At the same moment, the example points to a possible weakness: exchanging information just on their probability assessments induces (possibly) a revision. So, it seems to be difficult to reconcile subjective equilibria with the "core" assumptions we made: in particular, about the common knowledge of the parameters of the game.
As we shall see, ther is really a problem, here.

5.3 Common knowledge

Common Knowledge

I will follow chapter 5 of Osborne and Rubinstein.
We have Ω (finite, always, to simplify techniques) and $\mathcal{P}_1, \mathcal{P}_2$ two (information) partitions of Ω.
Notice that a partition $\mathcal{P}_i$ $(i = 1, 2)$ identifies an information function P_i, in a obvious way:

$P_i : \Omega \to 2^{\Omega} \setminus \{\emptyset\}$. ($2^{\Omega}$ denotes the set of all subsets of Ω).

$P_i(\omega)$ is just the set of $\mathcal{P}_i$ who contains ω.
We shall say that an event $F \subseteq \Omega$ is *SELF EVIDENT* between 1 and 2 if *FOR ALL* $\omega \in F$ we have that $P_i(\omega) \subseteq F, i = 1, 2$.
An event $E \subseteq \Omega$ is CK between 1 and 2 *IN THE STATE* $\omega \in \Omega$ if there is a self-evident event F s.t.: $\omega \in F \subseteq E$.
Notice that the following result holds.

Theorem Given $\Omega, \mathcal{P}_1, \mathcal{P}_2$ and an event E, the following are equivalent:

1. - E is self-evident between 1 and 2.

2. - E is union of members of the partitions $\mathcal{P}_i, i = 1, 2$

Proof 1) $\Rightarrow$ 2). Because $\forall \omega \in E, P_i(\omega) \subseteq E$, for $i = 1, 2$ we have that $E = \cup_{\omega \in E} P_i(\omega)$,for $i = 1, 2$.
Notice that $P_i(\omega)$ is an element of the partition $\mathcal{P}_i$, due to the way in which we defined P_i.
2) $\Rightarrow$ 1) Since
$E = \cup_{\alpha \in A} P_{1,\alpha}$ with $P_{1,\alpha} \in \mathcal{P}_1 \quad \forall \alpha \in A$
$E = \cup_{\beta \in B} P_{2,\beta}$ with $P_{2,\beta} \in \mathcal{P}_2 \quad \forall \beta \in B$.
Clearly, every $\omega \in E$ will be in some $P_{1,\alpha}$ (with $P_{1,\alpha} \subseteq E$).
So, E is self-evident (between 1 and 2).

We shall come back to the connection between information partitions and information functions in the next section. For more details, see once more chapter 5 of Osborne and Rubinstein.

5.4 Agreeing to disagree

We have Ω (finite), and p, a probability distribution on Ω (to be interpreted later as the "common prior").

We remind that a function $P : \Omega \to 2^{\Omega}\backslash\{\emptyset\}$ is said to be an *information function*. We shall assume that P satisfies the following conditons:

$\omega \in P(\omega) \quad \forall \omega \in \Omega$
if $\omega' \in P(\omega)$, then $P(\omega) = P(\omega')$

It can be shown that P is "partitional" (i.e., there is a partition such that for all $\omega \in \Omega$, $P(\omega)$ is just the element of the partition containing ω) if and only if P satisfies the two conditions above.

Let P be an information function and let E be an event.
Given $\omega \in \Omega$, at ω the DM will assign to E the probability

$$p(E|P(\omega))$$

(i.e. the probability of E, conditional on $P(\omega)$).
In our example, E was the event: sum of dice =8.
And, for example, at $\omega = (1, 3)$ we had $p(E|P_1(\omega)) = 0 \quad p(E|P_2(\omega)) = \frac{1}{6}$.
Remark: the event that "DM i assigns the probability p_i to E is:

$$\{\omega \in \Omega : p(E|P_i(\omega)) = p_i\}.$$

Theorem: It is given Ω finite and p probability on Ω (the common prior). We are given two information functions P_1 and P_2.
Assume that it is CK between 1 and 2 *in some state* $\omega^\star \in \Omega$ that 1 assigns probability p_1 to some event E and that 2 assigns probability p_2 to E.
Then, $p_1 = p_2$
Proof: The event "1 assignes probability p_1 to E and 2 assigns probability p_2 to E " is:

$$\{\omega \in \Omega : p(E|P_1(\omega)) = p_1\} \cap \{\omega \in \Omega : p(E|P_2(\omega) = p_2\}$$

Since it is assumed to be CK, there is a self evident F s.t. :

$$\omega^\star \in F \subseteq \{\underbrace{\omega \in \Omega : p(E|P_1(\omega))}_{*} = p_1\} \cap \{\omega \in \Omega : p(E|P_2(\omega)) = p_2\}$$

Thanks to the theorem proved above, we have that F is a union of members of the partition P_1 and P_2.
So, $F = \cup_{\alpha \in A} P_{1,\alpha} = \cup_{\beta \in B} P_{2,\beta}$.
Now, notice that $p(E|P_{1,\alpha}) = p_1$. To be sure of that, it is enough to notice that $P_{1,\alpha}$ is one of the $P_1(\omega)$ that appear in $*$.

In more detail:

Take $\omega \in F$.
Because $\omega \in *$, we have that $p(E|P_1(\omega)) = p_1$.
But $\omega \in F$, so $P_1(\omega)$ is one of the elements of the info partition whose union gives F. That is, $P_1(\omega) = P_{1,\alpha}$ for some $\alpha \in A$.

So, $p(E|P_{1,\alpha}) = p_1$ for every $\alpha \in A$.
Hence, $p(E|U_{\alpha \in A} P_{1,\alpha}) = p_1$.

(Namely, $[p(E|P_{1,\alpha'}) = p_1$ and $p(E|P_{1,\alpha''}) = p_1$ IMPLIES that $p(E|P_{\alpha'} \cup P_{1,\alpha''}) = p_1])$.

So, $p(E|F) = p_1$.
But the same reasoning can be repeated for $p_2 \cdots$ So we get $p(E|F) = p_2$.
But $p(E|F)$ is a well defined number ...

5.5 Correlated equilibria as expression of Bayesian rationality

Last remark on correlated equilibria.
The paper by Aumann in 1987 has an interesting title.
the assertion is that with correlated equilibria it should be possible to reconcile two different approaches:

- the GT approach, which tries to incorporate rationality and intelligence of the players into a solution concept, from which are derived the choices that should be made by a player.
- the subjectivistic approach, in which a player (as decision maker under uncertainty) will choose an action which maximizes his expected payoff. The expectation is based on the subjective probability assessments of the decision maker over the elements which are not under his control but that will influence the outcome. Notably, among them are the action(s) chosen by the other player(s).

It is intuitively clear that the amount of subjectivity has to be somehow constrained, if we want to get this "reconciliation" approach. We have seen that diverging subjective assessments on uncertain events can provide results quite different from the classical game theoretic predictions.
There should be some kind of "common ground".
This common ground is found in what has been named by Aumann as the "Harsanyi doctrine". That is, players may have different subjective probability distributions over uncertain events. But this difference should be ascribed only to different information status. Players should share a *common*

prior. From which different posteriori can arise, due to different exposure to experience, that is to different streams of information.
The result of Aumann is the following:

- if every player is bayesian rational *at each state of the world*, then the distribution of the action profile is a correlated equilibrium distribution.

To interpret the result, we need a set Ω of states of the world. Notice that an element of Ω gives a very detailed description of the situation.
I will do this for the case of two players .
A state of the world is $\omega \in \Omega$, where:

$\omega = (P_I(\omega), P_{II}(\omega), a_I(\omega), a_{II}(\omega), \mu_I(\omega), \mu_{II}(\omega))$

$P_I(\omega)$ is the set of states ω' that I cannot distinguish from ω, at ω (the information partition of I...).
$a_I(\omega)$ is the action chosen by I at ω.
$\mu_I(\omega)$ is the probability distribution (at ω) on the actions available to II (represents the believes of I w.r.t. the choices of II)
Now, if we assume that there is a "common prior" that is: a probability distribution P on Ω, and that:

- the believes of players are derived by P taking into account their information partitions (via Bayes' rule).
- the actions of players are constant on their information partitions (a quite reasonable assumption).
- at every ω players are rational (so, rationality is CK; rationality means that $a_I(\omega)$ is a best reply to $\mu_I(\omega)$.)

Then, the probability distribution induced by $\omega \to (a_I(\omega), a_{II}(\omega))$ is a correlated equilibrium.
See Osborne and Rubinstein, cap 5, for all of the details