

Chapter 9

Correlated and Sequential Equilibria

In this lecture, I will cover two important equilibrium concepts, namely correlated equilibrium and sequential equilibrium. Correlated equilibrium relaxes the assumption in the Nash equilibrium that the players' mixed strategies are independent (hence the name). It is therefore weaker than Nash equilibrium. It is stronger than rationalizability. On the other hand, sequential equilibrium is an equilibrium refinement. Unlike other refinements, sequential equilibrium makes the players' beliefs about the other players' strategies as an explicit part of equilibrium, in addition to strategy profiles. It is one of the most commonly used solution concepts, especially in dynamic games with incomplete information. For a more detailed discussion of these topics, see Fudenberg and Tirole's chapters 2.2 and 8.3.

9.1 Correlated Equilibrium

In a mixed-strategy Nash equilibrium it is assumed that the strategies are independently distributed. As it is explained in the previous lecture, there is no reason to believe that a player's belief about the other players' strategies are independent. Likewise, from an econometrician's point of view the distribution of the strategy profiles may contain correlation. Correlated equilibrium drops the independence assumption.

There are two ways to define correlated equilibrium. One way is to describe each player's information structure explicitly and impose the assumption that every player is a best response. Another way is to consider the distribution induced by such a model

on the strategy profile. The latter distribution is then characterized by using a simpler reduced form structure. I will first present the first formulation, which makes the logic of the solution concepts and its relation to rationalizability clearer.

Definition 34 *A (common-prior) information structure is a list $(\Omega, I_1, \dots, I_n, p)$ where Ω is a (finite) state space, p is a probability distribution on Ω and I_i is the information partition of player i for each i .*

I will write $I_i(\omega)$ for the cell of the partition I_i that contains ω . Here, if the true state is ω , player i is informed that the true state is in $I_i(\omega)$, and he does not get any other information. Such an information structure arises if each player observes a state-dependent signal, where $I_i(\omega)$ is the set of states in which the value of the signal of player i is identical to the value of the signal at state ω .

Finally, p is a common prior on Ω . I will assume without loss of generality that each information set $I_i(\omega)$ has positive probability, i.e., $p(I_i(\omega)) > 0$. Hence, by Bayes' rule, observing that the true state is in $I_i(\omega)$, player i updates his belief to $p(\cdot | I_i(\omega))$, which is a probability distribution on $I_i(\omega)$, where

$$p(\omega' | I_i(\omega)) = \frac{p(\omega')}{p(I_i(\omega))} \quad (\forall \omega' \in I_i(\omega)).$$

Definition 35 *An adapted strategy profile $(\mathbf{s}_1, \dots, \mathbf{s}_n)$ with respect to information structure $(\Omega, I_1, \dots, I_n, p)$ is a list of mappings $\mathbf{s}_i : \Omega \rightarrow S_i$ such that $\mathbf{s}_i(\omega) = \mathbf{s}_i(\omega')$ whenever $I_i(\omega) = I_i(\omega')$.*

Here, the last condition guarantees that player i knows what strategy he is playing.

Definition 36 *A correlated equilibrium with respect to information structure $(\Omega, I_1, \dots, I_n, p)$ is a strategy profile $(\mathbf{s}_1, \dots, \mathbf{s}_n)$ with respect to $(\Omega, I_1, \dots, I_n, p)$ such that for each i and ω , $\mathbf{s}_i(\omega)$ is a best response to \mathbf{s}_{-i} under $p(\omega' | I_i(\omega))$, i.e., for all s_i ,*

$$\begin{aligned} E[u_i(\mathbf{s}_i(\omega), \mathbf{s}_{-i}) | I_i(\omega)] &\equiv \sum_{\omega' \in I_i(\omega)} u_i(\mathbf{s}_i(\omega), \mathbf{s}_{-i}(\omega')) p(\omega' | I_i(\omega)) \\ &\geq \sum_{\omega' \in I_i(\omega)} u_i(s_i, \mathbf{s}_{-i}(\omega')) p(\omega' | I_i(\omega)) \equiv E[u_i(s_i, \mathbf{s}_{-i}) | I_i(\omega)]. \end{aligned}$$

The condition in the definition is, of course, equivalent to \mathbf{s}_i being a best response in the ex-ante stage. That is,

$$E [u_i(\mathbf{s}_i, \mathbf{s}_{-i})] \equiv \sum_{\omega \in \Omega} u_i(\mathbf{s}_i(\omega), \mathbf{s}_{-i}(\omega)) p(\omega) \geq \sum_{\omega \in \Omega} u_i(\mathbf{s}'_i(\omega), \mathbf{s}_{-i}(\omega)) p(\omega) \equiv E [u_i(\mathbf{s}'_i, \mathbf{s}_{-i})]$$

for any adapted strategy \mathbf{s}'_i .

Example 3 *As an example, study the correlated equilibria of the game in Figure 2.4 in Fudenberg and Tirole.*

Note that for any ω , $\mathbf{s}_i(\omega)$ is a best response to a correlated belief $p_{i,\omega}$ about the other players' strategy profiles where $p_{i,\omega}(s_{-i}) = \sum_{\omega' \in I_i(\omega), \mathbf{s}_{-i}(\omega') = s_{-i}} p(\omega' | I_i(\omega))$. By the same token, for each s_j with $p_{i,\omega}(s_j) > 0$, s_j is a best response to a belief $p_{j,\omega'}$, where $\mathbf{s}_j(\omega') = s_j$, and this is true ad infinitum. Hence, $\mathbf{s}_i(\omega)$ is rationalizable for player i . Therefore, correlated equilibrium is stronger than rationalizability. Note moreover that, unlike rationalizability, which does not put any restriction about the beliefs other than the above best response condition, the belief sequences obtained above exhibit stringent properties, as they are derived from a common prior p using the Bayes' rule. This is indeed the only distinction between the two concept.

A common-prior information structure assumes that the players' share a common prior belief. In a more general information/belief structure, each player would have his belief at each information set of his, and this can be represented by a list of probability distributions p_1, \dots, p_n on Ω , where p_i represent the (hypothetical) prior distribution of player i . The common-prior information structure assumes that $p_1 = \dots = p_n$. Rationalizability correspond to each player playing a best response at every information set in a general information structure.

In the above definition, we have an explicit information structure. One may be only interested in the probability distribution on the strategy profiles induced by $(\mathbf{s}, (\Omega, I_1, \dots, I_n, p))$. In that case, we can use a simpler formulation as follows.

Definition 37 *A correlated equilibrium is a probability distribution p on S such that for each s_i and s'_i ,*

$$\sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) p(s_{-i} | s_i) \geq \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) p(s_{-i} | s_i). \quad (9.1)$$

Note that this definition is a special case of the previous one in which the information structure is as follows

$$\begin{aligned}\Omega &= S \\ I_i(s) &= \{s_i\} \times S_{-i} = \{(s_i, s'_{-i}) \mid s'_{-i} \in S_{-i}\}.\end{aligned}\tag{9.2}$$

Conversely, in order to capture probability distributions induced by correlated equilibria with respect to arbitrary information structures, it suffices to consider this limited set of information structures. To see this, take any correlated equilibrium $(\mathbf{s}, (\Omega, I_1, \dots, I_n, p))$. The distribution \tilde{p} induced by $(\mathbf{s}, (\Omega, I_1, \dots, I_n, p))$ on S is given by $\tilde{p}(s) = \sum_{\omega \in \Omega, \mathbf{s}(\omega)=s} p(\omega)$. Now suppose that instead of letting i know that the true state is in $I_i(\omega)$, we only inform him that he needs to play $\mathbf{s}_i(\omega)$ according to \mathbf{s}_i . Since he did not have an incentive to deviate under any information (by definition of correlated equilibrium), by sure-thing principle, he does not have an incentive to deviate. Hence, the new information structure with limited information is also a correlated equilibrium. Since u_i does not depend on ω , the latter information structure can be represented by (9.2).

Exercise 14 Find all the correlated equilibria (as distributions on S) for the game of Figure 2.4 in Fudenberg and Tirole.

9.2 Sequential Equilibria

Consider the game in Figure 9.1. One can easily check that the strategy profile indicated with thick lines is a Nash equilibrium. Since the game does not have a proper subgame, it is also a subgame-perfect equilibrium. Nevertheless, the equilibrium prescribes the irrational move L for Player 2 at the information set she moves. At the information set she moves, she knows that Player 1 has played T or B. No matter what she believes about the likelihood of T or B, she finds R a better move than L, because conditional on T and B, R dominates L. Sequential equilibrium explicitly specifies the beliefs of players at each information set that they move and requires that the players act rationally according to these beliefs and that the beliefs are consistent with the solution.

Formally, consider an extensive form game. Consider an information set h at which a player $i(h)$ moves, where h is a collection of nodes that i is to move and cannot distinguish from each other. At h , player $i(h)$ knows that he is at one of the nodes h ,




Figure 9.1: A subgame-perfect equilibrium with sequentially irrational moves

but he does not anything more than that. Hence, being an expected utility maximizer, he has a belief about the nodes, a probability distribution $\mu(\cdot|h)$ on h . A *belief system* μ is a list of such probability distributions, one for each information set.

Recall also that a mixed strategy σ_i of a player i is a complete contingent plan that maps each information set h of player i to a mixed action $\sigma_i(\cdot|h)$ that is available at h . An *assessment* is a pair (σ, μ) of a strategy profile σ and a system of beliefs μ .

Definition 38 An assessment (σ, μ) is sequentially rational if at each information set h , playing according to $\sigma_{i(h)}$ in the continuation game is a best response for $i(h)$ to belief $\mu(\cdot|h)$ and the belief that the other players will play according to $\sigma_{-i(h)}$ in the continuation game, i.e., for any strategy $\sigma'_{i(h)}$,

$$\int u_i(\sigma_i, \sigma_{-i}) d\mu(\cdot|h) \geq \int u_i(\sigma'_i, \sigma_{-i}) d\mu(\cdot|h).$$

For example, in Figure 9.1, for player 2, given any belief μ , L yields

$$U_2(L; \mu) = 1 \cdot \mu(T|\{T, B\}) + 3 \cdot \mu(B|\{T, B\})$$

while R yields

$$U_2(R; \mu) = 2 \cdot \mu(T|\{T, B\}) + 5 \cdot \mu(B|\{T, B\}).$$

Hence, sequential rationality requires that player 2 plays R. Given player 2 plays R, the only best reply for player 1 is T. Therefore, for any belief assessment μ , the only sequentially rational strategy profile is (T, R) .




Figure 9.2: An inconsistent belief assessment

In order to have an equilibrium, we also need to require that μ is consistent with σ . Roughly speaking, consistency requires that players know which (possibly mixed) strategies are played by the other players. For a motivation, consider Figure 9.2 and call the node on the left n_T and the node on the right n_B , writing also $h_2 = \{n_T, n_B\}$. Given the beliefs $\mu(n_T|h_2) = 0.1$ and $b(n_B|h_2) = 0.9$, strategy profile (T, R) is sequentially rational. Strategy T is a best response to R . To check the sequential rationality for R , it suffices to note that, given the beliefs, L yields

$$(.1)(10) + (.9)(3) = 3.7$$

while R yields

$$(.1)(2) + (.9)(5) = 4.7.$$

(Note that there is no continuation game.) But (T, R) is not even a Nash equilibrium in this game. This is because in a Nash equilibrium player knows the other player's strategy. She would know that player 1 plays T , and hence she would assign probability 1 on n_T . In contrast, according to μ , she assigns only probability 0.1 on n_T .

Therefore, as an equilibrium condition, one would also like to impose that the beliefs $\mu(h)$ are consistent with the strategy profile σ , in that the beliefs are derived from σ using Bayes' rule. That is, when $\sigma(h) > 0$, for each node $x \in h$, $\mu(x|h) = \sigma(x)/\sigma(h)$, where $\sigma(x)$ is the probability of reaching node x under σ and $\sigma(h) = \sum_{x \in h} \sigma(x)$. For example, in order a belief assessment μ to be consistent with (T, R) , we need

$$\mu(n_T|h_2) = \frac{\Pr(n_T|(T, R))}{\Pr(n_T|(T, R)) + \Pr(n_B|(T, R))} = \frac{1}{1+0} = 1.$$

Unfortunately, in general, there can be information sets that are not supposed to be reached according to the strategy profile, i.e., $\sigma(h) = 0$. In that case, Bayes' rule does not apply, and conditional beliefs are arbitrary. For such information sets, we perturb the strategy profile slightly, by assuming that players may "tremble", and apply the Bayes rule using the perturbed strategy profile. To see the general idea, consider the game in Figure 9.3. The information set of player 3 is off the path of the strategy profile (X, T, L) . Hence, we cannot apply the Bayes rule. But we can still see that the beliefs the figure are inconsistent. Let us perturb the strategies of players 1 and 2 assuming that players 1 and 2 tremble with probabilities ε_1 and ε_2 , respectively, where ε_1 and ε_2 are small but positive numbers. That is, we put probability ε_1 on E and $1 - \varepsilon_1$ on X (instead of 0 and 1, respectively) and $1 - \varepsilon_2$ on T and ε_2 on B (instead of 1 and 0, respectively). Under the perturbed beliefs,

$$\Pr(n_T|h_3, \varepsilon_1, \varepsilon_2) = \frac{\varepsilon_1(1 - \varepsilon_2)}{\varepsilon_1(1 - \varepsilon_2) + \varepsilon_1\varepsilon_2} = 1 - \varepsilon_2,$$

where n_T is the node that follows T , and h_3 is the information set Player 3 moves. As $\varepsilon_2 \rightarrow 0$, $\Pr(n_T|h_3, \varepsilon_1, \varepsilon_2) \rightarrow 1$. Therefore, for consistency, we need $\mu(n_T|h_3) = 1$. Formally, consistency is defined as follows.

Definition 39 *An assessment (σ, μ) is consistent if there exists a sequence (σ^m, μ^m) of assessment converging to (σ, μ) such that for each m ,*

- σ^m is completely mixed (i.e. $\sigma_{i(h)}^m(a|h) > 0$ for every h and every action a available at h),
- and $\mu^m(\cdot|h)$ is derived from σ^m using Bayes' rule at each h :

$$\mu^m(x|h) = \frac{\sigma^m(x)}{\sigma^m(h)} \quad \forall x \in h.$$

Definition 40 *A sequential equilibrium is a sequentially rational and consistent assessment.*

Example 4 *In the game in Figure 9.3, the unique subgame-perfect equilibrium is $s^* = (E, T, R)$. Let us check that (s^*, μ^*) where $\mu^*(n_T|h_3) = 1$ is a perfect Bayesian Nash equilibrium. We need to check that*




Figure 9.3: A belief assessment that is inconsistent off the path

1. s^* is sequentially rational (at all information sets) under b^* , and
2. μ^* is consistent with s^* .

At the information set of player 3, given $\mu^*(n_T|h_3) = 1$, action L yields 1 while R yields 3, and hence R is sequentially rational. At the information set of player 2, given the other strategies, T and B yield 3 and 1, respectively, and hence playing T is sequentially rational. At the information set of player 1, E and X yield 3 and 2, respectively, and hence playing E is again sequentially rational.

Since all the information sets are reached under s^* , we just need to use the Bayes rule in order to check consistency:

$$\Pr(n_T|h_3, s^*) = \frac{1}{1+0} = \mu^*(n_T|h_3).$$

MIT OpenCourseWare
<http://ocw.mit.edu>

14.123 Microeconomic Theory III
Spring 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.