

14.12 Game Theory Lecture Notes

Lectures 15-18

1 Dynamic Games with Incomplete Information

In these lectures, we analyze the issues arise in a dynamics context in the presence of incomplete information, such as how agents should interpret the actions the other parties take. We define perfect Bayesian Nash equilibrium, and apply it in a sequential bargaining model with incomplete information. As in the games with complete information, now we will use a stronger notion of rationality — sequential rationality.

2 Perfect Bayesian Nash Equilibrium

Recall that in games with complete information some Nash equilibria may be based on the assumption that some players will act sequentially irrationally at certain information sets off the path of equilibrium. In those games we ignored these equilibria by focusing on subgame perfect equilibria; in the latter equilibria each agent's action is sequentially rational at each information set. Now, we extend this notion to the games with incomplete information. In these games, once again, some Bayesian Nash equilibria are based on sequentially irrational moves off the path of equilibrium.

Consider the game in Figure 1. In this game, a firm is to decide whether to hire a worker, who can be hard-working (High) or lazy (Low). Under the current contract, if the worker is hard-working, then working is better for the worker, and the firm makes profit of 1 if the worker works. If the worker's lazy, then shirking is better for him, and the firm will lose 1 if the worker shirks. If the worker is sequentially rational, then he will work if he's hard-working and shirk if he's lazy. Since the firm finds the worker




Figure 1:

more likely to be hard-working, the firm will hire the worker. But there is another Bayesian Nash equilibrium: the worker always shirks (independent of his type), and therefore the firm does not hire the worker. This equilibrium is indicated in the figure by the bold lines. It is based on the assumption that the worker will shirk when he is hard-working, which is sequentially irrational. Since this happens off the path of equilibrium, such irrationality is ignored in the Bayesian Nash equilibrium—as in the ordinary Nash equilibrium.

We'll now require sequential rationality at each information set. Such equilibria will be called perfect Bayesian Nash equilibrium. The official definition requires more details.

For each information set, we must specify the beliefs of the agent who moves at that information set. Beliefs of an agent at a given information set are represented by a probability distribution on the information set. In the game in figure 1, the players' beliefs are already specified. Consider the game in figure 2. In this game we need to specify the beliefs of player 2 at the information set that he moves. In the figure, his beliefs are summarized by μ , which is the probability that he assigns to the event that player 1 played T given that 2 is asked to move.

Given a player's beliefs, we can define sequential rationality:

Definition 1 *A player is said to be sequentially rational iff, at each information set he is to move, he maximizes his expected utility given his beliefs at the information set (and given that he is at the information set) – even if this information set is precluded by his*




Figure 2:

own strategy.

In the game of figure 1, sequential rationality requires that the worker works if he is hard-working and shirks if he is lazy. Likewise, in the game of figure 2, sequential rationality requires that player 2 plays R.




Figure 3:

Now consider the game in figure 3. In this figure, we depict a situation in which player 1 plays T while player 2 plays R, which is not rationalizable. Player 2 assigns probability .9 to the event that player 1 plays B. Given his beliefs, player 2's move is sequentially rational. Player 1 plays his dominant strategy, therefore his move is sequentially rational. The problem with this situation is that player 2's beliefs are not

consistent with player 1's strategy. In contrast, in an equilibrium a player maximizes his expected payoff given the other players' strategies. Now, we'll define a concept of consistency, which will be required in a perfect Bayesian Nash equilibrium.

Definition 2 *Given any (possibly mixed) strategy profile s , an information set is said to be on the path of play iff the information set is reached with positive probability according to s .*

Definition 3 (Consistency on the path) *Given any strategy profile s and any information set I on the path of play of s , a player's beliefs at I is said to be consistent with s iff the beliefs are derived using the Bayes' rule and s .*

For example, in figure 3, consistency requires that player 2 assigns probability 1 to the event that player 1 plays T. This definition does not apply off the equilibrium path. Consider the game in Figure 4. In this game, after player 1 plays E, there




Figure 4:

is a subgame with a unique rationalizable strategy profile: 2 plays T and 3 plays R. Anticipating this, player 1 must play E. Now consider the strategy profile (X,T,L), in which player 1 plays X, 2 plays T, and 3 plays L, and assume that, at his information set, player 3 assigns probability 1 to the event that 2 plays B. Players' moves are all sequentially rational, but player 3's beliefs are not consistent with what the other players play. Since our definition was valid only for the information sets that are on the path of equilibrium, we could not preclude such beliefs. Now, we need to extend our definition

of consistency to the information sets that are off the path of equilibrium. The difficulty is that the information sets off the path of equilibrium are reached with probability 0 by definition. Hence, we cannot apply Bayes' formula to compute the beliefs. To check the consistency we might make the players "tremble" a little bit so that every information sets is reached with positive probability. We can then apply Bayes rule to compute the conditional probabilities for such a perturbed strategy profile. Consistency requires that the players' beliefs must be close to the probabilities that are derived using Bayes' rule for some such small tremble (as the size of the tremble goes to 0). In figure 4, for any small tremble (for player 1 and 2), the Bayes rule yields a probability close to 1 for the event that player 2 plays T. In that case, consistency requires that player 3 assigns probability 1 to this event. **Consistency is required both on and off the equilibrium path.**

In the definition of sequential rationality above, the players' beliefs about the nodes of the information set are given but his beliefs about the other players' play in the continuation game are not specified. In order to have an equilibrium, we also need these beliefs to be specified consistently with the other players' strategies.

Definition 4 *A strategy profile is said to be sequentially rational iff, at each information set, the player who is to move maximizes his expected utility given*

1. *his beliefs at the information set, and*
2. *given that the other players play according to the strategy profile in the continuation game (and given that he is at the information set).*

Definition 5 *A Perfect Bayesian Nash Equilibrium is a pair (s,b) of strategy profile and a set of beliefs such that*

1. *s is sequentially rational given beliefs b , and*
2. *b is consistent with s .*

The only perfect Bayesian equilibrium in figure 4 is (E,T,R). This is the only subgame perfect equilibrium. Note that every perfect Bayesian equilibrium is subgame perfect.

3 Examples

Beer-Quiche Game Consider the game in figure 5. In this game, player one has two types: weak or strong. Player 2 thinks that player 1 is strong with probability .9. Player 2, who happens to be a bully, wants to fight with player 1 if player 1 is weak and would like to avoid a fight if player 1 is strong. Player 1 is about to order his breakfast, knowing that player 2 observes what player 1 orders. He prefers beer if he is strong, and he prefers quiche if he is weak. He wants to avoid a fight.




Figure 5:

This game has two equilibria. (For each equilibrium there is a continuum of mixed strategy equilibria off the path of equilibrium.) First, consider the perfect Bayesian Nash equilibrium depicted in figure 6. We need to check two things: sequential rationality and consistency. Let us first check that the strategy profile is sequentially rational. In his information set on the right, player 2 is sure that player 1 is weak, hence he chooses to duel. When he sees that player 1 is having beer for his breakfast he assigns probability .9 to the event that player 1 is strong. Hence, his expected payoffs from duel is $.9 \times 1 = .9$, and his expected payoff is .1 otherwise. Therefore, his moves are sequentially rational. Now consider the strong type of player 1. If he chooses beer, then he gets 3, and if he chooses quiche, then he gets 0. He chooses beer. Now consider the weak type. If he chooses beer, he gets 2, while he gets only 1 if he chooses quiche. He chooses beer. Therefore, player 1's moves are sequentially rational.




Figure 6:

Let's now check the consistency. The information set after the beer is on the path of equilibrium; hence we need to use the Bayes' rule. The probabilities .9 and .1 are indeed computed through Bayes' rule. The information set after the quiche is off the equilibrium path. In this game, any belief off the equilibrium path is consistent. For the present belief, which puts probability 1 to the weak type, consider a perturbation in which player 1 trembles and orders quiche with probability ε if he is weak, and he does not tremble if his strong. Now Bayes' rule yields $\varepsilon/\varepsilon = 1$ as the conditional probability of being weak given quiche. Therefore, the players beliefs are consistent, and we have a perfect Bayesian Nash equilibrium.

Note that we have a continuum of equilibria in which player 1 orders beer. After the quiche, player 2 assigns equal probabilities to each node and mixes between duel and not duel, where the probability of duel is at least .5. Check also that there is a perfect Bayesian Nash equilibrium in which player 1 orders quiche independent of this type, and player 2 fights when he observes a beer.

Another example Consider the game in figure 7. sequential rationality requires that at the last node in the upper branch player 1 goes down, and at the last node of the lower branch player 1 goes across. Moreover, it requires that player 1 goes across at the first node of the lower branch. Therefore, player 1 must go across throughout the lower branch and go down at the last node of the upper branch at any perfect Bayesian Nash equilibrium. We now show that, in any perfect Bayesian Nash equilibrium, the players




Figure 7:

must play mixed strategies at the remaining information sets (i.e., at the first node of the upper branch, and at the information set of player 2). Suppose that player 1 goes down with probability 1 at the first node on the upper branch. Then, by Bayes' rule, player 2 must assign probability 1 to the lower branch at his information set and must go down with probability 1. In that case, it is better for player 1 to go across and get 5, rather than going down and getting 4—a contradiction. Therefore, player 1 must go across with positive probability at the first node of a upper branch. Now, suppose that player 1 goes across with probability 1 at this node. Then by Bayes' rule, player 2 must assign probability .9 to the upper branch in his information set. If he goes down, he gets 2; if he goes across, he gets $.9 \times 3 + .1 \times (-5) = 2.2$. Then, he must go across with probability 1. In that case, player 1 must go down with probability 1 at the node at hand—another contradiction. Therefore, player 1 must mix at the present node. In order to have this, player 1 must be indifferent between going across and going down. Let's write β for the probability that 2 goes across. For indifference, we must have

$$4 = 5(1 - \beta) + 3\beta = 5 - 2\beta,$$

i.e.,

$$\beta = 1/2.$$

Player 2 must also play a mixed strategy. Since player 2 plays a mixed strategy, he must be indifferent. Let's write μ for the probability he assigns to the upper branch at his information set. For indifference, we must have

$$2 = 3\mu + (1 - \mu)(-5) = 8\mu - 5,$$

i.e.,

$$\mu = 7/8.$$

If player 1 goes across with probability α , then by Bayes' rule, we must have

$$\mu = \frac{.9\alpha}{.9\alpha + .1} = \frac{7}{8},$$

hence

$$\alpha = 7/9.$$

Therefore, there is a unique perfect Bayesian Nash equilibrium as depicted in figure 8.




Figure 8:

4 Sequential bargaining

4.1 A one-period model with two types

We have a seller S with valuation 0 and a buyer B with valuation v . B knows v , S does not; S believes that $v = 2$ with probability π , and $v = 1$ with probability $1 - \pi$. We have the following moves. First, S sets a price $p \geq 0$. Knowing p , B either buys, yielding $(p, v - p)$ (where the first entry is the payoff of the seller), or does not, yielding $(0,0)$. The game is depicted in Figure 9.




Figure 9:




Figure 10:

The perfect Bayesian Nash equilibrium is as follows. B buys iff $v \geq p$. If $p \leq 1$, both types buy, and S gets p . If $1 < p \leq 2$, only H-type buys, and S gets πp . If $p > 2$, no one buys. The expected payoff of S is plotted in Figure 10. S offers 1 if $\pi < \frac{1}{2}$, and he offers 2 if $\pi > \frac{1}{2}$. He is indifferent between the prices 1 and 2 when $\pi = \frac{1}{2}$.

4.2 A two-period model with two types

Consider the same buyer and the seller, but allow them to trade at two dates $t = 0, 1$. The moves are as follows. At $t = 0$, S sets a price $p_0 \geq 0$. B either buys, yielding $(p_0, v - p_0)$, or does not. If he does not buy, then at $t = 1$, S sets another price $p_1 \geq 0$; B either buys, yielding $(\delta p_1, \delta(v - p_1))$, or does not, yielding $(0, 0)$.

The equilibrium behavior at $t = 1$ is the same as above. Let's write

$$\mu = \Pr(v = 2 | \text{history at } t = 1).$$

B buys iff $v \geq p_1$. If $\mu > \frac{1}{2}$, $p_1 = 2$; if $\mu < \frac{1}{2}$, $p_1 = 1$. If $\mu = \frac{1}{2}$, S is indifferent between 1 and 2.

Given this, B with $v = 1$ buys at $t = 0$ if $p_0 \leq 1$. Hence, by Bayes' rule, if $p_0 > 1$,

$$\mu = \Pr(v = 2 | p_0, t = 1) \leq \pi.$$

When $\pi < 1/2$, this determines the equilibrium. This is because

$$\mu = \Pr(v = 2 | p_0, t = 1) \leq \pi < 1/2,$$

and thus

$$p_1 = 1.$$

Hence, B with $v = 2$ buys at $t = 0$ if

$$(2 - p_0) \geq \delta(2 - 1) = \delta.$$

This is true iff

$$p_0 \leq 2 - \delta.$$

Now S has two options: either set $p_0 = 1$ and sell the good with probability 1, yielding payoff 1, or set $p_0 = 2 - \delta$, and sell to the high-value buyer at $t = 0$ and sell the low-value buyer at $t = 1$. The former is better, and thus $p_0 = 1$:

$$\pi(2 - \delta) + (1 - \pi)\delta = 2\pi(1 - \delta) + \delta < (1 - \delta) + \delta = 1.$$

Consider the case $\pi > 1/2$. In that case, after any price $p_0 \in (2 - \delta, 2)$, the players must mix (see the slides). At any $p_0 > 2 - \delta$, since B mixes at $t = 1$, we must have

$$\mu(p_0) = \Pr(v = 2 | p_0, t = 1) = 1/2.$$

Write $\beta(p_0)$ for the probability that high-value buyer does not buy at price p_0 . Then, by Bayes' rule,

$$\mu(p_0) = \frac{\beta(p_0)\pi}{\beta(p_0)\pi + (1 - \pi)} = \frac{1}{2},$$

i.e.,

$$\beta(p_0) = (1 - \pi) / \pi.$$

Since the buyer with $v = 2$ mixes (i.e., $\beta(p_0) \in (0, 1)$), he must be indifferent towards buying at p_0 . That is, writing $\gamma(p_0) = \Pr(p_1 = 1|p_0)$, we have

$$2 - p_0 = \delta\gamma(p_0)$$

i.e.,

$$\gamma(p_0) = (2 - p_0) / \delta.$$

4.3 A one-period model with continuum of types

Modify the one-period model above by letting v be distributed uniformly on some interval $[0, a]$. In equilibrium, again B buys at price p iff $v \geq p$. S gets

$$U(p) = p \Pr(v \geq p) = p(a - p)/a.$$

Therefore, S sets

$$p = a/2.$$

4.4 A two-period model with continuum of types

Modify the two-period model above by letting v be distributed uniformly on $[0, 1]$. B buys at p_0 iff

$$v - p_0 \geq \delta(v - E[p_1|p_0]), \tag{1}$$

where $E[p_1|p_0]$ is the expected value of p_1 given p_0 . This inequality holds iff

$$v \geq \frac{p_0 - E[p_1|p_0]}{1 - \delta} \equiv a(p_0).$$

Hence, if B does not buy at price p_0 , S's posterior belief will be that v is uniformly distributed on $[0, a(p_0)]$, in which case he will set the price at $t = 1$ to

$$p_1(p_0) = a(p_0) / 2$$

as shown above. Substituting this into the previous definition we obtain

$$a(p_0) = \frac{p_0 - \delta E[p_1|p_0]}{1 - \delta} = \frac{p_0 - \delta a(p_0) / 2}{1 - \delta},$$

i.e.,

$$a(p_0) = \frac{p_0}{1 - \delta/2}.$$

(One could obtain this, simply by observing that (1) is an equality when $v = a(p_0)$ and $E[p_1|p_0] = a(p_0)/2$.) Notice that, if S offers p_0 , in equilibrium, he sells to the types $v \geq a(p_0)$ at price p_0 (at date $t = 0$), to the types with $a(p_0)/2 \leq v \leq a(p_0)$ at $p_1 = a(p_0)/2$ at date 1, and does not sell to the types $v < a(p_0)/2$ at all. His expected payoff is

$$\begin{aligned} U_S(p_0) &= \Pr(v > a(p_0)) p_0 + \delta \Pr(p_1 \leq v < a(p_0)) p_1 \\ &= (1 - a(p_0)) p_0 + \delta (a(p_0)/2) (a(p_0)/2) \\ &= \left(1 - \frac{p_0}{1 - \delta/2}\right) p_0 + \delta \left(\frac{p_0}{2 - \delta}\right)^2. \end{aligned}$$

The first order condition yields

$$0 = U'_S(p_0) = 1 - \frac{2p_0}{1 - \delta/2} + \frac{2\delta p_0}{(2 - \delta)^2},$$

i.e.,

$$p_0 = \frac{(1 - \delta/2)^2}{2(1 - 3\delta/4)}.$$