

CUR 412: Game Theory and its Applications, Lecture 11

Prof. Ronaldo CARPIO

May 17, 2016

Announcements

- ▶ Homework #4 will be posted on the web site later today, due in two weeks.

Review of Last Week

- ▶ An extensive form game with *imperfect information* is one where players may not know the exact history when they move.
- ▶ This can model uncertainty about Nature or other player's moves.
- ▶ An *information set* is a set of histories (nodes) that are indistinguishable to a player.
- ▶ The set of actions at each node in an information set must be the same.
- ▶ A subgame cannot split the nodes in an information set.
- ▶ A player's strategy now specifies an action at each of his information sets.

Review of Last Week

- ▶ Suppose we have a random experiment.
- ▶ Suppose A, B are events (i.e. sets of possible outcomes).
- ▶ $P(A)$ is the probability that event A has occurred, i.e. that the random outcome is a member of the set A .
- ▶ The *conditional probability* of A given B , denoted $P(A|B)$, is the probability that the outcome is in A , after we take as given that B has occurred.

Examples

- ▶ Consider our 4-sided dice roll, with result X . Assume the dice is unbiased: each side occurs with equal probability $\frac{1}{4}$.
- ▶ Let A be the event that $X \leq 2$. Then $P(A) = \frac{1}{2}$.
- ▶ Let B be the event that X is even. Then $P(A) = P(2) + P(4) = \frac{1}{2}$.
- ▶ $P(A \cap B) = P(2) = \frac{1}{4}$, which is equal to $P(A)P(B)$. Therefore, events A and B are independent.

Conditional Probability

- ▶ If we know one event A has occurred, does that affect the probability that another event B has occurred?
- ▶ Suppose we are told that event A has occurred. Then $P(A) > 0$, and every outcome outside of A is no longer possible.
- ▶ Therefore, the universe is reduced to A . The only part of B which can occur is $A \cap B$.
- ▶ Since total probability of the universe must equal 1, the probability of $A \cap B$ must be scaled by $\frac{1}{P(A)}$

$$P(B|A) = \frac{P(B \cap A)}{P(A)}, P(B \cap A) = P(A)P(B|A)$$

Conditional Probability

$$P(B|A) = \frac{P(B \cap A)}{P(A)}, P(B \cap A) = P(A)P(B|A)$$

- ▶ $B \cap A$ and $\neg B \cap A$ are disjoint, and their union is A . Therefore, $P(B \cap A) + P(\neg B \cap A) = P(A)$.

$$P(B|A) = \frac{P(B \cap A)}{P(A)} = \frac{P(B \cap A)}{P(B \cap A) + P(\neg B \cap A)}$$

- ▶ $P(B|A)$ is the probability of B conditional on A occurring. Equivalently,

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A \cap B)}{P(A \cap B) + P(\neg A \cap B)}, P(A \cap B) = P(B)P(A|B)$$

Bayes' Theorem

$$P(B|A) = \frac{P(B \cap A)}{P(A)} = \frac{P(B \cap A)}{P(B \cap A) + P(\neg B \cap A)}, P(B \cap A) = P(A)P(B|A)$$

- ▶ Using the result $P(B \cap A) = P(B)P(A|B)$ and $P(\neg B \cap A) = P(\neg B)P(A|\neg B)$, we get

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)}$$

- ▶ This is known as Bayes' Theorem. It is simply substituting in different formulas for $\frac{P(B \cap A)}{P(A)}$.
- ▶ Using this formula, we can calculate the probability of any event conditional on any other event, if we know the all of the other probabilities in the formula.

Example

- ▶ Consider again the 4-sided dice that gives value $X \in \{1, 2, 3, 4\}$. Suppose that the dice is no longer unbiased:

$$P(1) = P(2) = P(3) = 0.2, P(4) = 0.4$$

- ▶ A is the event that $X \leq 2$. $P(A) = 0.4$
- ▶ B is the event that X is even. $P(B) = 0.6$
- ▶ $P(A \cap B) = 0.2$, which is different from $P(A)P(B) = 0.24$. Therefore, A and B are not independent.

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{0.2}{0.6} = \frac{1}{3}$$

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{0.2}{0.4} = \frac{1}{2}$$

Example

- ▶ Here's an real-world example of using Bayes' Theorem:
- ▶ Suppose in the general population, we observe that 1.4% of women whose age 40-49 develop breast cancer.
- ▶ A mammogram is a diagnostic test to predict breast cancer. We know that it has a 10% false positive rate: when the test is used on someone who *does not* have cancer, it says they do 10% of the time.
- ▶ The test has a 25% false negative rate: when the test is used on someone who *does* have cancer, it says they do 75% of the time.
- ▶ Suppose a woman gets a mammogram and it gives a positive result. What is the probability that she has cancer, knowing nothing else?
- ▶ Let B be the event that she has cancer, and A the event that the test gives a positive result.
- ▶ We want to find $P(B|A)$, which is:

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)} = \frac{1.4\% \times 0.75}{1.4\% \times 0.75 + 98.6\% \times 0.1}$$

Example

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)} = \frac{1.4\% \times 0.75}{1.4\% \times 0.75 + 98.6\% \times 0.1}$$

- ▶ $P(B|A)$ is 0.096242, about 10 %.
- ▶ B is our *prior distribution*: our beliefs *before* we learn any new information.
- ▶ After we observe new information from A , we update our beliefs, giving us a *posterior distribution*.
- ▶ "Prior" means "before", "posterior" means "after".


Example

- ▶ Suppose we are the person in this situation, and we got a positive result. What can we do to be sure that this is not a false positive?
- ▶ We can take a *second* test. Our beliefs after the second test will depend on the correlation between the two tests.
- ▶ Suppose the two tests are *independent*: the outcome of the first test is completely unrelated to the second test.
- ▶ If the second test gives a positive result for cancer, we can use Bayes' Theorem again, using 0.0962 as our *prior* distribution.

$$P(B|A) = \frac{9.62\% \times 0.75}{9.62\% \times 0.75 + (1.0 - 9.62\%) \times 0.1}$$

- ▶ $P(B|A)$ is 0.4439, about 44 %.


Beliefs as a Probability Distribution




- ▶ Let's return to the description of BoS as an extensive game with imperfect information.
- ▶ Consider Player 2's information set. He knows that there are two possibilities for Player 1's action: B or S . Some possible beliefs of Player 2 are:
 - ▶ B and S are equally likely.
 - ▶ B is more likely to have occurred than S .
 - ▶ It is impossible for B to have occurred. S has occurred with certainty.

Beliefs as a Probability Distribution

- ▶ We can formulate these statements in a precise way with a probability distribution.
- ▶ From the player's point of view, he is adopting a subjectivist view of probability.
- ▶ Player 2 places a probability on B and S that must add up to 1; the relative sizes of the probabilities reflects Player 2's opinion on how likely it is that B vs. S has occurred.
- ▶ The first number is the probability on B . For example:
 - ▶ If Player 2's opinion is that B and S are equally likely, beliefs are modeled by a probability distribution $(\frac{1}{2}, \frac{1}{2})$.
 - ▶ If Player 2's opinion is that B is more likely to have occurred, the probability distribution is of the form $(p, 1 - p)$ where $p > \frac{1}{2}$.
 - ▶ If Player 2's opinion is that S is impossible, the probability distribution is $(1, 0)$.



- ▶ What is the probability that history L occurs? p
- ▶ What is the probability that history R occurs? $1 - p$
- ▶ The probability that information set $\{L\}$ is reached is p .
- ▶ The probability that information set $\{R\}$ is reached is $1 - p$.



- ▶ What is the probability that history M occurs? p_2
- ▶ What is the probability that history R occurs? $1 - p_1 - p_2$
- ▶ The probability that information set $\{L\}$ is reached is p_1 .
- ▶ The probability that information set $\{M, R\}$ is reached is $P(M) + P(R) = 1 - p_1$.
- ▶ Suppose we know we have reached information set $\{M, R\}$. What is the probability that we are in history M ?
- ▶ This is the conditional probability $P(M|\{M, R\})$. From Bayes' Rule, this is:

$$P(M|\{M, R\}) = \frac{P(M \cap \{M, R\})}{P(\{M, R\})} = \frac{p_2}{1 - p_1}$$

- ▶ Similarly, the conditional probability $P(R|\{M, R\})$ is $\frac{1-p_1-p_2}{1-p_1}$.


Behavioral strategy in extensive games

- ▶ The equivalent of a mixed strategy for extensive games is a *behavioral strategy*.
- ▶ A **behavioral strategy** of player i in an extensive game is a function that assigns to each of i 's information sets I_i , a probability distribution over the actions available at I_i , with the property that each probability distribution is independent of every other distribution.
- ▶ A *behavioral strategy profile* is a list of behavioral strategies for all players.
- ▶ In the extensive form of BoS, each player has one information set, so a behavioral strategy would specify a probability distribution over actions in that information set.
- ▶ As before, a behavioral strategy may place 100% probability on one action, in which case it is a pure strategy.

Beliefs Consistent With Behavioral Strategies


- ▶ When we defined the Nash equilibrium solution concept, we said that there were two parts:
 - ▶ Players' actions were optimal given their beliefs about the game situation;
 - ▶ Their beliefs were correct.
- ▶ So far, we have not had to explicitly model beliefs; they were implicitly derived from the strategy profile. Now we will do so.
- ▶ We want players' beliefs about possible histories to be *correct*, given a behavioral strategy profile.
- ▶ That is, the true probability distribution generated by the behavioral strategy profile should be equal to the players' beliefs at all information sets.

Example 1




- ▶ Here, suppose Player 1 chooses L, M, R randomly according to a known distribution: $P(L) = p_1, P(M) = p_2, P(R) = 1 - p_1 - p_2$.
- ▶ Player 1 may be a human player, or it may be Nature.
- ▶ Player 2 has one information set, consisting of the histories $\{L, M, R\}$.
- ▶ What is the correct belief at Player 2's information set?

Example 1




- ▶ The true probabilities generated by Player 1's behavioral strategy is simply the strategy itself. So, the only correct belief at Player 2's information set is $(p_1, p_2, 1 - p_1 - p_2)$.

Example 2




- ▶ In this game, Nature goes first and chooses L with probability p , and R with probability $1 - p$.
- ▶ If R is chosen, then it is Player 1's move. He can choose Out and end the game, or choose In .
- ▶ Assume his behavioral strategy is to choose In with probability q , and Out with probability $1 - q$, independent of Nature's outcome.

Example 2




- ▶ Player 2 has one information set, consisting of the histories $\{L, (R, In)\}$.
- ▶ What is the probability that history L occurs? p
- ▶ What is the probability that history R, In occurs?
- ▶ Since Nature's random variable and Player 1's random variable are independent, this is the product of the probabilities of R and In : $(1 - p)q$

Example 2



- ▶ What is the probability that history R, Out occurs? $(1 - p)(1 - q)$
- ▶ What is the probability that Player 2's information set $\{L, (R, In)\}$ is reached? $P(L) + P(R, In) = p + (1 - p)q$

Example 2




- ▶ Conditional on reaching Player 2's information set, what is the probability that L has occurred?

$$\frac{P(L \cap \{L, (R, In)\})}{P(\{L, (R, In)\})} = \frac{P(L)}{P(\{L, (R, In)\})} = \frac{p}{p + (1-p)q}$$

- ▶ Conditional on reaching Player 2's information set, what is the probability that R, In has occurred?

$$\frac{(1-p)q}{p + (1-p)q}$$

Example 2




- ▶ Correct beliefs for Player 2 would be:

$$\left(\frac{p}{p + (1-p)q}, \frac{(1-p)q}{p + (1-p)q} \right)$$


- ▶ Then, given these beliefs, Player 2 can then calculate the expected payoff to his actions at his information set.

Example 2




- ▶ There is a possible problem here. What if $p = 0$ and $q = 0$?
- ▶ The probability of reaching Player 2's information set is 0.
- ▶ In this case, Player 2 has no basis on which to form beliefs.
- ▶ We will say that *any* probability distribution is a consistent belief at an information set that is reached with zero probability.
- ▶ So, any probability distribution $r, 1 - r$ over $L, (R, In)$ is consistent with the behavioral strategy profile.

Example 3




- ▶ In this game, Nature goes first and chooses T with probability a , B with probability $1 - a$.
- ▶ Player 1 observes this choice, and chooses L or R .
- ▶ Player 2 does not observe Nature's choice, but only Player 2's choice.
- ▶ This setup is quite common in signaling games. Nature's choice determines the *type* of Player 1, which may be T or B .

Example 3




- ▶ Player 1 has two information sets: $\{T\}$ and $\{B\}$. At each of these information sets, he has the same two actions L, R .
- ▶ Player 2 has two information sets: the information set after L , $\{TL, BL\}$, and the information set after R , $\{TR, BR\}$.

Example 3



- ▶ Suppose that Player 1's behavioral strategy specifies this probability distribution at each of his two information sets:
- ▶ $P(L|T) = p, P(R|T) = 1 - p, P(L|B) = q, P(R|B) = 1 - q$
- ▶ Here, $P(L|T)$ means "if T occurs, choose L with probability p ".

Example 3




- ▶ What is the probability of reaching Player 2's information set $\{TL, BL\}$?

$$P(TL) + P(BL) = P(T)P(L|T) + P(B)P(L|B) = ap + (1 - a)q$$

- ▶ What is the probability of reaching Player 2's information set $\{TR, BR\}$?

$$P(TR) + P(BR) = P(T)P(R|B) + P(B)P(R|B) = a(1 - p) + (1 - a)(1 - q)$$

Example 3




- Suppose we are at information set $\{TL, BL\}$. What is the probability that history TL has occurred?

$$\begin{aligned}P(TL|\{TL, BL\}) &= \frac{P(TL \cap \{TL, BL\})}{P(\{TL, BL\})} = \frac{P(TL)}{P(\{TL, BL\})} \\ &= \frac{ap}{ap + (1-a)q}\end{aligned}$$

- Therefore, the beliefs at information set $\{TL, BL\}$ that would be consistent with the behavioral strategy of Player 1 are:

$$\left(\frac{ap}{ap + (1-a)q}, \frac{(1-a)q}{ap + (1-a)q} \right)$$


Example 3



- ▶ Similarly, the beliefs at information set $\{TR, BR\}$ that would be consistent with the behavioral strategy of Player 1 are:

$$\left(\frac{a(1-p)}{a(1-p) + (1-a)(1-q)}, \frac{(1-a)(1-q)}{a(1-p) + (1-a)(1-q)} \right)$$

Example 3




- ▶ Information set $\{TL, BL\}$ is reached with zero probability if $p = 0, q = 0$.
- ▶ Information set $\{TR, BR\}$ is reached with zero probability if $p = 1, q = 1$.
- ▶ Any probability distribution is a consistent belief at an information set reached with zero probability.

- ▶ Now, we know how to calculate beliefs at all information sets, which should allow us to calculate expected payoffs for each player's action.
- ▶ Let's state some definitions which will lead to our concept of equilibrium:
- ▶ **Definition 324.1:** A **belief system** is a function that assigns to each information set, a probability distribution over the histories in that information set.
- ▶ A belief system is simply a list of beliefs for every information set.

Weak Sequential Equilibrium

- ▶ **Definition 325.1:** An **assessment** is a pair consisting of a profile of behavioral strategies, and a belief system. An assessment is called a **weak sequential equilibrium** if these two conditions are satisfied:
 - ▶ All players' strategies are optimal when they have to move, given their beliefs and the other players' strategies. This is called the *sequential rationality* condition.
 - ▶ All players' beliefs are consistent with the strategy profile, i.e. their beliefs are what is computed by Bayes' Rule for information sets that are reached with positive probability. This is called the *consistency of beliefs* condition.
- ▶ Note that this generalizes the two conditions of Nash equilibrium.
- ▶ Clearly, this may be difficult to find for complex games. We will usually simplify things by assuming that players only use pure strategies (Nature still randomizes).
- ▶ To specify a weak sequential equilibrium, we need two things:
 - ▶ The strategy profile (actions at every information set);
 - ▶ The beliefs at every information set.

Example 317.1: Entry Game with Preparation




- ▶ Consider a modified entry game. The challenger has three choices: stay out, prepare for a fight and enter, or enter without preparation.
- ▶ Preparation is costly, but reduces the loss from a fight.
- ▶ As before, the incumbent can *Fight* or *Acquiesce*.
- ▶ A fight is less costly to the incumbent if the entrant is unprepared, but the incumbent still prefers to *Acquiesce*.
- ▶ *Incumbent* does not know if *Challenger* is prepared or unprepared.

Example 317.1: Entry Game with Preparation

- ▶ First, let's convert this game to strategic form and find the NE.
- ▶ Then, we will see which NE can also be part of a weak sequential equilibrium.
- ▶ For each NE, we will find the beliefs that are consistent with the strategy profile.
- ▶ Then, check if the strategies are optimal, given beliefs.
- ▶ If both conditions are satisfied, then we have found a weak sequential equilibrium.

Strategic Form



	<i>Acquiesce</i>	<i>Fight</i>
<i>Ready</i>	3,3	1,1
<i>Unready</i>	4,3	0,2
<i>Out</i>	2,4	2,4

	<i>Acquiesce</i>	<i>Fight</i>
<i>Ready</i>	3,3	1,1
<i>Unready</i>	4,3	0,2
<i>Out</i>	2,4	2,4

- ▶ NE are: (*Unready*, *Acquiesce*) and (*Out*, *Fight*).
- ▶ Consider (*Unready*, *Acquiesce*): at Player 2's information set, only consistent belief over $\{Ready, Unready\}$ is $(0, 1)$.
 - ▶ Given this belief, the optimal action for Player 2 is *Acquiesce*. This matches the NE, so this is a WSE.
- ▶ Consider (*Out*, *Fight*): Player 2's information set is not reached, so any beliefs over $\{Ready, Unready\}$ are consistent.
 - ▶ However, *Acquiesce* is optimal given any belief over $\{Ready, Unready\}$. Therefore, this NE cannot be a WSE.

Adverse Selection

- ▶ Let's examine one more application of imperfect information.
- ▶ Suppose Player 1 owns a car, and Player 2 is considering whether to buy the car from Player 1.
- ▶ The car has a level of quality α , which can take on three types, L, M, H .
- ▶ Player 1 knows the type of the car, but Player 2 does not.
- ▶ Player 1's valuation of the car is:

$$v_1(\alpha) = \begin{cases} 10 & \text{if } \alpha = L \\ 20 & \text{if } \alpha = M \\ 30 & \text{if } \alpha = H \end{cases}$$

- ▶ The higher the quality, the higher is Player 1's valuation of the car.


Adverse Selection

- ▶ Likewise, suppose that Player 2 has a similar valuation of the car:

$$v_2(\alpha) = \begin{cases} 14 & \text{if } \alpha = L \\ 24 & \text{if } \alpha = M \\ 34 & \text{if } \alpha = H \end{cases}$$

- ▶ Player 2 knows that the probability distribution of quality levels in the general population is $\frac{1}{3}$ for each type.
- ▶ Note that Player 2's valuation of the car is higher than Player 1's valuation, for all quality levels of the car.
- ▶ Therefore, if the type were common knowledge, a trade should occur: it is always possible to find a Pareto-efficient trade (that is, both players are not worse off, and at least one player is better off).
- ▶ For example, if it were known that quality was L , a trade at a price between 10 and 14 would make both players better off.

- ▶ Consider the following game:
 - ▶ Nature chooses the type of the car with $P(L) = P(M) = P(H) = \frac{1}{3}$. Player 1 observes the type; Player 2 does not.
 - ▶ Player 2 makes a price offer $p \geq 0$ to Player 1 for the car.
 - ▶ Player 1 can accept (A) or reject (R).
 - ▶ If Player 1 accepts, he gets the price offered, and the car is transferred to Player 2.
 - ▶ If trade occurs, Player 1's payoff is the price offered. Player 2's payoff is his valuation of the car, minus the price paid.
 - ▶ If trade does not occur, Player 1's payoff is his valuation of the car. Player 2's payoff is zero.



Adverse Selection

- ▶ Consider the subgame after p has been offered.
- ▶ Player 1's best response is:
 - ▶ If $\alpha = L$, accept if $p \geq 10$, reject otherwise.
 - ▶ If $\alpha = M$, accept if $p \geq 20$, reject otherwise.
 - ▶ If $\alpha = H$, accept if $p \geq 30$, reject otherwise.
- ▶ Now, consider Player 2's decision. His beliefs match Nature's probability distribution: the probability on L, M, H is $1/3$ each.
- ▶ Let's find the expected payoff $E_2(p)$ of choosing p , for the range $p \geq 0$.

Adverse Selection

- ▶ If $p < 10$, Player 1 will reject in all cases. $E_2(p) = 0$.
- ▶ If $10 \leq p \leq 14$, $E_2(p) = \frac{1}{3}(14 - p) + \frac{1}{3}(0) + \frac{1}{3}(0) = \frac{14-p}{3}$.
 - ▶ This is non-negative if p is in this range.
- ▶ If $14 < p < 20$, $E_2(p) = \frac{1}{3}(14 - p) + \frac{1}{3}(0) + \frac{1}{3}(0) = \frac{14-p}{3} < 0$.
- ▶ If $20 \leq p \leq 24$, $E_2(p) = \frac{1}{3}(14 - p) + \frac{1}{3}(24 - p) + \frac{1}{3}(0) = \frac{38-2p}{3} < 0$.
- ▶ If $24 < p < 30$, $E_2(p) = \frac{1}{3}(14 - p) + \frac{1}{3}(24 - p) + \frac{1}{3}(0) = \frac{38-2p}{3} < 0$.
- ▶ If $30 \leq p$, $E_2(p) = \frac{1}{3}(14 - p) + \frac{1}{3}(24 - p) + \frac{1}{3}(34 - p) = \frac{72-3p}{3} < 0$.

Adverse Selection

- ▶ The optimal choice is for Player 2 to offer $p = 10$. If $\alpha = L$, Player 1 will accept; otherwise Player 1 will reject.
- ▶ Note that in only the lowest quality case does trade occur.
- ▶ This is clearly inefficient, since trades that would benefit both parties are not taking place.
- ▶ This is an example of *adverse selection*, in which the low-quality type "drives out" the high-quality type from the market, due to uncertainty.
- ▶ One way to overcome this problem is if the buyer could get some information about the true quality of the car.
- ▶ The seller could simply say that the car is high-quality. But why should the buyer believe him?
- ▶ Next week, we'll see a method by which the seller can credibly communicate the quality of the car.

Announcements

- ▶ We are currently covering Chapter 10.
- ▶ Homework #4 will be posted on the web site later today, due in two weeks.