CUR 412: Game Theory and its Applications, Lecture 10

Prof. Ronaldo CARPIO

May 10, 2016

Announcements

- Homework #3 is due at the end of class.

Ch 7.2: Entry into a monopolized industry

- Consider a model that combines the Entry Game with Cournot Duopoly.
- There is an industry which is currently monopolized by the *Incumbent* firm.
- A *Challenger* firm is deciding whether to choose *In* or *Out*.
- If *In* is chosen, *Challenger* has to pay a cost of entry f; then both firms face Cournot duopoly.
- If *Out* is chosen, *Challenger* makes zero profit while *Incumbent* faces a monopoly problem.

Ch 7.2: Entry into a monopolized industry

- Assume the same parameters as in the Cournot duopoly we've seen before:
- Cost of production $C_i(q_i) = cq_i$, i.e. constant marginal cost, equal for both firms.
- Inverse market demand is given by:

$$P_d(Q) = \begin{cases} \alpha - Q & \text{if } \alpha \geq Q \\ 0 & \text{if } \alpha < Q \end{cases}$$

Subgame after *In*

- Consider the subgame if *In* is chosen. We know the outcome of the Cournot duopoly game in this case is:
- $q_{incumbent} = q_{challenger} = (\alpha - c)/3$. Both firms choose the same output
- *Incumbent*'s profit is: $\pi_{incumbent} = \frac{1}{9}(\alpha - c)^2$
- *Challenger*'s profit is: $\pi_{challenger} = \frac{1}{9}(\alpha - c)^2 - f$

Subgame after *Out*

- Consider the subgame if *Out* is chosen. This is a standard monopoly problem where the *Incumbent* firm chooses q to maximize

$$\pi(q) = q(\alpha - q) - cq$$

- The optimal value of q is $(\alpha - c)/2$.
- *Incumbent*'s profit is $\frac{1}{4}(\alpha - c)^2$
- *Challenger*'s profit is 0.

Challenger's first move

- Subgame after *In*:
  - *Incumbent*'s profit: $\pi_{incumbent} = \frac{1}{9}(\alpha - c)^2$
  - *Challenger*'s profit: $\pi_{challenger} = \frac{1}{9}(\alpha - c)^2 - f$
- Subgame after *Out*:
  - *Incumbent*'s profit is: $\pi_{incumbent} = \frac{1}{4}(\alpha - c)^2$
  - *Challenger*'s profit is: $\pi_{challenger} = \frac{1}{9}(\alpha - c)^2 - f$
- *Challenger* will choose *In* if $\frac{1}{9}(\alpha - c)^2 - f > 0$ and *Out* if $\frac{1}{9}(\alpha - c)^2 - f < 0$.
- If they are equal, there are two SPNE outcomes, one where *Challenger* chooses *In* and one where it chooses *Out*.

Ch 7.6: Allowing for exogenous uncertainty

- So far, all the extensive games we've seen have been deterministic: no randomness in outcomes.
- We can allow randomness by introducing an additional player, called "Chance" or "Nature"
- Nature makes choices randomly, according to a known probability distribution.
- Players' preferences are now over lotteries.
- As before, we will assume von Neumann-Morgenstern preferences (i.e. lotteries are valued by their expected payoff).

Ch 7.6: Allowing for exogenous uncertainty

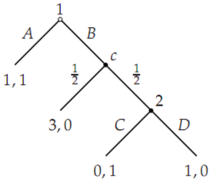

- Consider this game with chance moves.
- Here, "c" is the Chance player, who chooses randomly between two branches, each with probability $1/2$.
- This is still a game with *perfect information*: at each player's move, he knows exactly what sequence of moves has occurred in the past.
- In the last subgame, Player 2 chooses C, so this subgame has a payoff of $(0, 1)$.

Ch 7.6: Allowing for exogenous uncertainty

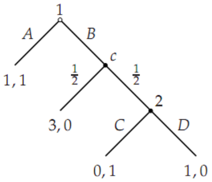

- In the subgame following B, Chance chooses each action with probability $\frac{1}{2}$.
- The expected payoff to Player 1 of this subgame is therefore $\frac{1}{2}3 + \frac{1}{2}0 = 1.5$.
- The expected payoff to Player 2 of this subgame is $\frac{1}{2}0 + \frac{1}{2}1 = 0.5$.
- At the beginning, Player 1 chooses B, since the expected payoff of B is greater than the expected payoff of A.

Exercise 227.2: Firm-Union Bargaining

- Suppose a firm and a union are bargaining over how to split the "surplus", i.e. the profits before labor has been paid.
- Surplus is a random variable, that takes on the value H with probability p, and $L < H$ with probability $1 - p$.
- The sequence of the game is as follows:
  - First, the union (who does not know the outcome of the surplus, just its distribution) makes a demand $x \geq 0$.
  - The firm, who does know what the surplus (denoted z) is, can choose to *Accept* or *Reject*. If *Accept* is chosen, the union gets x and the firm gets $z - x$. If *Reject* is chosen, both players get 0.

Exercise 227.2: Firm-Union Bargaining

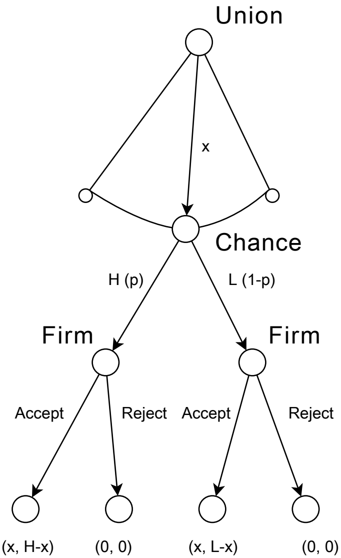

Exercise 227.2: Firm-Union Bargaining

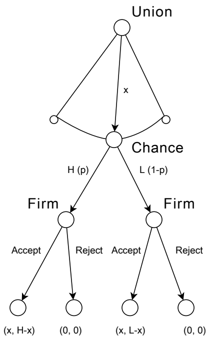

- Note that "Chance" moves after the union but before the firm; this is how we ensure that the firm makes its choice with the knowledge of the outcome of the surplus.
- Again, note that this game has perfect information.
- Assume that in case of a tie in payoffs, the firm chooses *Accept*.

Exercise 227.2: Firm-Union Bargaining

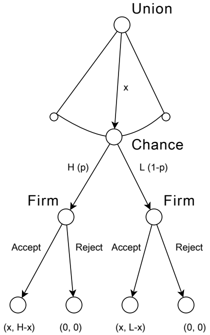

- We can solve this with backwards induction:
  - In the bottom left subgame, take x as given. The firm will choose *Accept* if $H - x \geq 0$.
  - In the bottom right subgame, take x as given. The firm will choose *Accept* if $L - x \geq 0$.

Exercise 227.2: Firm-Union Bargaining

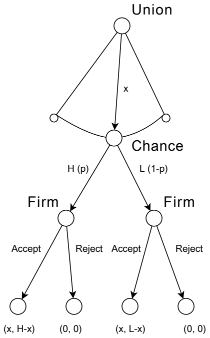

- Taking the firm's actions as given, the union chooses x. Let's consider the possible cases for x:
- Suppose $x \leq L$. Then $x < H$ as well, so the firm will *Accept* if either H or L is the outcome. The expected payoff to the union is: $px + (1 - p)x = x$.

Exercise 227.2: Firm-Union Bargaining

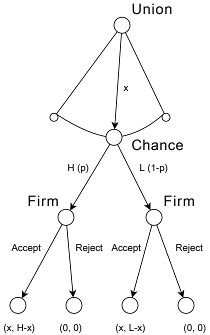

- Suppose $L < x \leq H$. The firm will *Accept* if H is the outcome, but will *Reject* if L is the outcome. The expected payoff is: $px + (1 - p)0 = px$.
- Suppose $H < x$. The firm will always *Reject*, so the expected payoff is 0.

Exercise 227.2: Firm-Union Bargaining

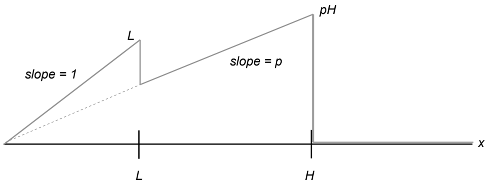

- We can plot the expected payoff as a function of x.
- If $L > pH$, then $x = L$ is the optimal action for the union.
- If $L < pH$, then $x = H$ is the optimal action.
- If $L = pH$, then either is optimal for the union, and there is more than one SPNE outcome.

Interpretation of Subgame Perfect NE

- What does a SPNE correspond to in "real life"?
- When we studied NE of strategic games, we saw one possible interpretation:
- A steady state that occurs in a repeated situation where people are drawn from different populations (each population corresponds to a different player type).
- For mixed NE, people are randomly drawn from fractions of a population, with each fraction always choosing an action.
- Or: people randomly choose from their set of actions.
- These interpretations apply to masses of social interactions, but not an individual interaction.

Interpretation of Subgame Perfect NE

- It's possible to imagine an individual interaction as a NE, but that requires some assumptions:
  - All players are completely rational, and they know each others' payoff functions and the probability distributions of any sources of randomness
  - They also know all levels of *higher-order knowledge*: Player 1 knows what Player 2 knows, and Player 3 knows what Player 1 knows what Player 2 knows, ...
  - Players then simulate all possible strategies in their mind, and then choose one that happens to be a NE
- This becomes more implausible as the game gets more complicated.

Interpretation of Subgame Perfect NE

- For a SPNE of an extensive game, it's hard to apply the "steady state" interpretation.
- We can interpret the process of backward induction as reasoning about what will happen in the future.
- If all players are rational, they will end up choosing a SPNE strategy profile.
- But what if we find ourselves in a history that is *not* rational?
- Not clear what to believe at that point - you could continue as if all players were rational, but you have evidence that they are not.

Centipede Game

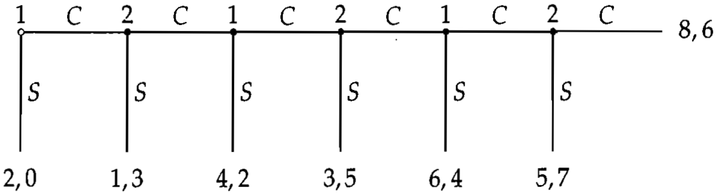

- Let's play this extensive game.
- I will choose two people to be Player 1 and Player 2.

Centipede Game

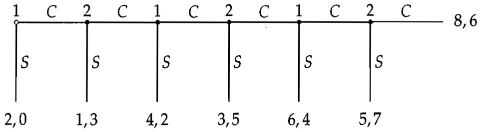

- The SPNE is for Player 1 to end the game by choosing *s* on the first move.
- However, that gives the lowest payoff to Player 2 and the second lowest payoff to Player 1.
- Empirically, people cooperate until they are near the end, even after they know what the SPNE is.
- One possible answer is that one of the assumptions of SPNE does not hold: *common knowledge of rationality*.
- In the real world, people are not completely certain that the other player will always choose to end the game as soon as possible.

Chapter 10: Extensive Games with Imperfect Information

- Up to now, the extensive games we have studied have been games with *perfect information*, i.e. players always know exactly what has happened in the past when making decisions.
- This eliminates any kind of uncertainty about the past, e.g:
  - What if moves by another player are unknown (as in simultaneous-move games)?
  - What if some parameter of the game or the players depends on a random variable?
- We can allow for uncertainty in this way: if information about the past is hidden from a player when he makes his decision, we'll say the player *cannot distinguish* between histories (nodes) that differ based on that information.

BoS as an Extensive Game with Imperfect Information

- Consider the simultaneous-move game BoS.
- We cannot model this as an extensive game with perfect information, because each player must not know what the other player chose.
- With imperfect information, we can model the game this way:
  - Assume Player 1 moves first, choosing B or S, but this information is hidden from Player 2.
  - Player 2 cannot distinguish between the history where Player 1 chose B and where he chose S.

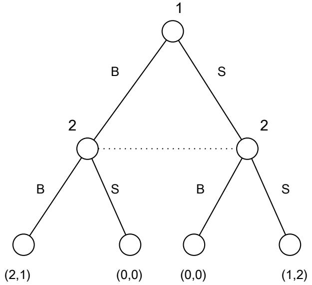

- On the tree diagram, we denote this with a dashed line connecting the two nodes.
- We call this group of nodes (i.e. a group of histories) an *information set* of Player 2.
- The actions available to a player at all histories (nodes) in an information set *must be the same*.
- Note that we could also represent BoS as an extensive game where Player 2 moves first, and Player 1's information set contains B and S.

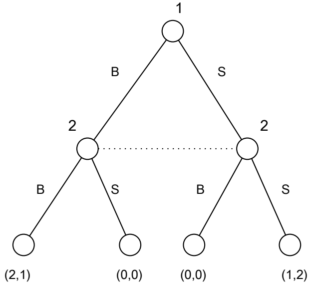

- With imperfect information, our definition of *subgame* must be modified.
- if a subgame contains a node in an information set, it must contain *all* the nodes in that information set (i.e. it makes no sense to have a game where not all the actions are specified).
- The extensive form of BoS has only one subgame, since we cannot split Player 2's information set.

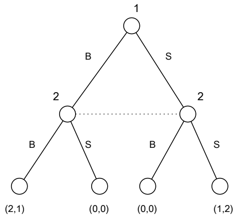

- We can convert extensive games with imperfect information into strategic form, just as before.
- We can also allow for pure and mixed strategies in the strategic form.
- A pure strategy must specify an action at every information set.
- A mixed strategy must specifiy a *probability distribution over actions* at every information set.
- In the extensive form of BoS, both players have a single information set, at which the set of possible actions is B and S; so a strategy would specify a single distribution over B, S.

- Recall the firm-union bargaining game with exogenous uncertainty from last week. In that example:
    1. The union offers w.
    2. Chance chooses the surplus to be L or H randomly.
    3. The firm observes Chance's choice, and chooses *Accept* or *Reject*.
- This is a game with *perfect information*, since every player who moves after a random choice, knows what the outcome was.
- Suppose we change the order of moves to the following:
    1. Chance chooses the surplus to be L or H randomly.
    2. The union offers w.
    3. The firm observes Chance's choice, and chooses *Accept* or *Reject*.
- This becomes a game with *imperfect information*.

Perfect Information

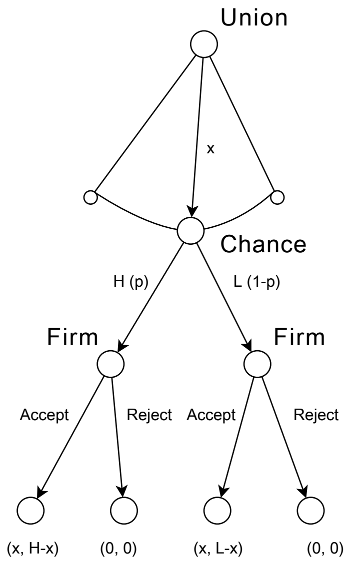

Imperfect Information

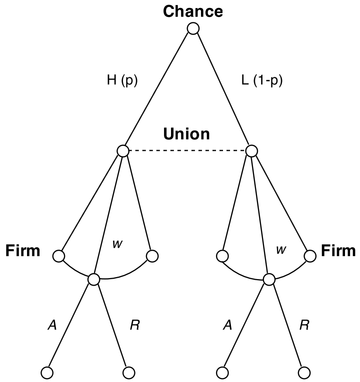

- In this diagram, the firm knows what the surplus is, so the firm's strategy can specify different actions depending on whether L or H occurred.
- The union does not know, so its strategy cannot *condition on* L or H.

Beliefs over Histories

- Many interesting situations involve uncertainty due to a random outcome in the past.
- In Chapter 7, we saw how to incorporate randomness by adding a player called "Chance" or "Nature", who chooses actions randomly, according to a known probability distribution.
- Imperfect information is a natural way to handle randomness in games.
- Note that with randomness, we must now assume preferences over lotteries; specifically, we will assume von Neumann-Morgenstern utilities (i.e. lotteries are valued based on expected payoff).

Beliefs over Histories

- Suppose it is a player's turn to move, and his information set has more than one history (node).
- His future payoffs may depend on which history he is actually in.
- In order to be able to calculate his payoffs to his actions, he needs some way to quantify his beliefs about which history has actually occurred.
- We will use *probability distributions* over histories in an information set to model beliefs about what has happened in the past.
- This will use concepts of conditional probability and Bayes' rule.

Review of Probability

- Suppose we have a random experiment with a range of possible outcomes.
- For example: we roll a 4-sided dice, with sides labeled 1, 2, 3, 4. Call the result X.
- The *universe* U is the set of all possible outcomes of the world.
- In this case, $U = \{1, 2, 3, 4\}$.
- An *event* is a *set* of outcomes.
- Examples of events
  - The set of all possible outcomes, U.
  - The empty set $\emptyset$.
  - The set of outcomes where $X \leq 2$: $\{1, 2\}$.
  - The set of outcomes where X is even:: $\{2, 4\}$.

Review of Probability

- The *probability* of an event (that is, set of states) is a number between 0 and 1 assigned to each event. What determines this number?
- First, let's talk about the *frequentist* view of probability.
- In the frequentist view, the probability of event A, denoted $P(A)$, is the frequency at which we would observe an outcome in A, if we ran the experiment an infinite number of times.

Subjectivist View of Probability

- In the *subjectivist* view, probability is a belief in an individual's mind, based on past experience, logic, and whatever else leads the individual to believe that something is likely or unlikely to happen.
- In the subjectivist view, probabilities apply to any uncertain situation, not just one that can be replicated an infinite number of times. For example:
  - What is the probability that you will get married before age 30?
  - What is the probability that there is life on other planets?
- These are one-time events in the future or even the past. It makes no sense to talk about replicating them, but we have varying levels of belief in them.
- Regardless of one's view of where probability numbers come from, they must satisfy certain rules.

Laws of Probability

- For any event A, $P(A) \geq 0$.
- $P(U) = 1$.
- $\neg A$ is the set of all outcomes that are *not* in A. $P(\neg A) = 1 - P(A)$
- $A \cup B$ is the event that the outcome is in A *or* in B.
- If A and B are disjoint (i.e. no outcome is in both A and B), then $P(A \cup B) = P(A) + P(B)$.
- This extends to $A_1, A_2, ...A_n$: if these are all disjoint, the $P(A_1 \cup ... \cup A_n) = P(A_1) + ...P(A_n)$.
- $A \cap B$ is the event that the outcome is in A *and* in B.
- Two events are *independent* if $P(A \cap B) = P(A)P(B)$.

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.

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

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

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 positve?
- 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.v
- 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

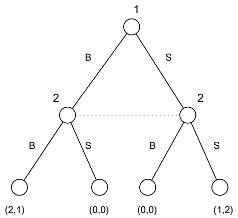

- Let's return to the description of BoS as an extensive game with imperfect information.
- Consider Player 2's information set. He knows that there are two possibilities for Player 1's action: B or S. Some possible beliefs of Player 2 are:
  - B and S are equally likely.
  - B is more likely to have 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 $(1, 0)$.

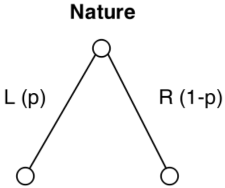

- What is the probability that history L occurs? p
- What is the probability that history R occurs? $1 - p$
- The probability that information set $\{L\}$ is reached is p.
- The probability that information set $\{R\}$ is reached is $1 - p$.

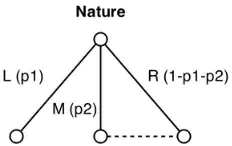

- What is the probability that history M occurs? p_2
- What is the probability that history R occurs? $1 - p_1 - p_2$
- The probability that information set $\{L\}$ is reached is p_1.
- The probability that information set $\{M, R\}$ is reached is $P(M) + P(R) = 1 - p_1$.
- Suppose we know we have reached 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 $\frac{1-p_1-p_2}{1-p_1}$.

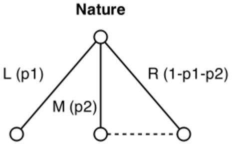

- At information set $\{M, R\}$, the player will have beliefs over the possible histories M, R.
- This can be any probability distribution over M, R.
- We say that beliefs are *consistent* with the given probability distribution of Nature if they match the true probability

$$\left(\frac{p_2}{1-p_1}, \frac{1-p_1-p_2}{1-p_1}\right)$$

Next Class

- Please read the appendix (Chapter 17) in the textbook dealing with conditional probability and Bayes' rule.
- Please read Chapter 10.
- If you don't have the official copy of the book, please email me.