Advanced Microeconomic Analysis, Lecture 12

Prof. Ronaldo CARPIO

June 12, 2017

Administrative Stuff

- The final exam will be on June 23 from 13:30-16:30 in Boxue 504.
- It will cover the material after the midterm exam and is written jointly with the other two microeconomics classes.

Review of Last Week

- So far, the games we have studied have been *simultaneous-move* games: all players choose their strategy at the same time.
- This cannot model sequential situations, in which one player moves, then another player moves, etc...
- We model sequential situations with an *extensive form game*.
- An extensive form game with *perfect information* (that is, all players know what moves have been chosen in the past) is represented as a tree.
- Each node in the tree corresponds to some player's turn to choose an action.
- At each node, the branches from that node correspond to the player's actions.
- A player's strategy in an extensive form game specifies an action for every node where it is that player's turn to move.

Review of Last Week

- One way to solve an extensive form game is to convert it into a simultaneous-move game.
- We can list all possible strategies for each player, and form a payoff matrix just like a strategic form game.
- Then, find the Nash equilibria as usual.
- This treats each player's strategy as chosen once and for all at the beginning of the game.
- However, this can lead to NE that contain "incredible threats", that is, actions that are not optimal when played in sequence.
- To solve this problem, we introduce the concept of *subgame perfect* equilibria.

Review of Last Week

- A subgame is a subset of a game tree that results from choosing a node N in the tree, and removing everything except the branches and nodes that follow N.
- A subgame perfect equilibria is a NE that also induces a NE in every subgame.
- This ensures that all moves will be optimal when played in sequence.
- The backwards induction algorithm will find all subgame perfect equilibria in finite extensive form games.

Stackelberg Duopoly

- In Chapter 4, we saw models of duopoly (Cournot and Bertrand) that were simultaneous-move situations.
- Now, let's look at a duopoly model where one firm gets to move first. Does being first to move provide an advantage or disadvantage?
- Suppose firms are competing by choosing output quantities, as in Cournot duopoly.
- Each of two firms $i = 1, 2$ chooses to produce $q_i \geq 0$ units.
- Total cost of production: $C_i(q_i)$.
- Inverse market demand function: $P_d(Q)$, where $Q = q_1 + q_2$.
- Profits: $P_d(q_i + q_j)q_i - C_i(q_i)$.

Stackelberg Duopoly

- Assume:
  - constant unit cost: $C_i(q_i) = cq_i$
  - market demand:

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

- Best response of firm 2, taking q_1 as given, is:

$$b_2(q_1) = \begin{cases} \frac{1}{2}(\alpha - c - q_1) & \text{if } q_1 \leq \alpha - c \\ 0 & \text{if } q_1 > \alpha - c \end{cases}$$

- This will be Firm 2's strategy in the last subgame.

Stackelberg Duopoly

$$b_2(q_1) = \begin{cases} \frac{1}{2}(\alpha - c - q_1) & \text{if } q_1 \leq \alpha - c \\ 0 & \text{if } q_1 > \alpha - c \end{cases}$$

- Firm 1's problem is to choose q_1 that maximizes

$$q_1 P(q_1 + b_2(q_1)) - cq_1 = q_1(\alpha - c - (q_1 + \frac{1}{2}(\alpha - c - q_1))) = \frac{1}{2} q_1(\alpha - c - q_1)$$

- This is a quadratic that is maximized at $q_1 = \frac{1}{2}(\alpha - c)$.
- Unique SPNE outcome: Firm 1 chooses $q_1 = \frac{1}{2}(\alpha - c)$
- Firm 2's strategy is given by $b_2(q_1)$.
- The outcome for Firm 2 will be $q_2 = b_2(q_1) = b_2(\frac{1}{2}(\alpha - c)) = \frac{1}{4}(\alpha - c)$.

Stackelberg Duopoly

- Compared to the Cournot duopoly outcome where both firms choose $q_i = \frac{1}{3}(\alpha - c)$, Firm 1 produces more and Firm 2 produces less.
- Firm 1's profits are also higher: $\frac{1}{8}(\alpha - c)^2$ compared to $\frac{1}{9}(\alpha - c)^2$ in Cournot duopoly.
- In this case, going first is an advantage for the first mover.

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

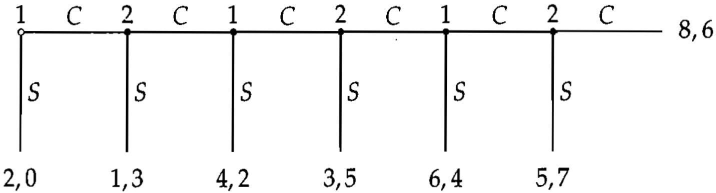

- Let's play this extensive game.
- I will choose two people to be Player 1 and Player 2.

Centipede Game

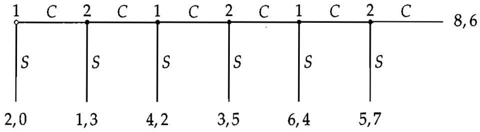

- The SPNE is for Player 1 to end the game by choosing *s* on the first move.
- However, that gives the lowest payoff to Player 2 and the second lowest payoff to Player 1.
- Empirically, people cooperate until they are near the end, even after they know what the SPNE is.
- One possible answer is that 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.

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).

Allowing for exogenous uncertainty

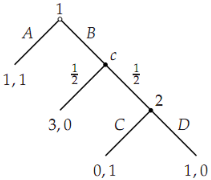

- Consider this game with chance moves.
- Here, "c" is the Chance player, who chooses randomly between two branches, each with probability $1/2$.
- This is still a game with *perfect information*: at each player's move, he knows exactly what sequence of moves has occurred in the past.
- In the last subgame,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

Allowing for exogenous uncertainty

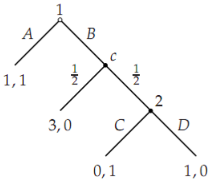

- In the subgame following B, Chance chooses each action with probability $\frac{1}{2}$.
- The expected payoff to Player 1 of this subgame is therefore $\frac{1}{2}3 + \frac{1}{2}0 = 1.5$.
- The expected payoff to Player 2 of this subgame is $\frac{1}{2}0 + \frac{1}{2}1 = 0.5$.
- At the beginning, Player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

Firm-Union Bargaining with Uncertainty

- 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.

Firm-Union Bargaining with Uncertainty

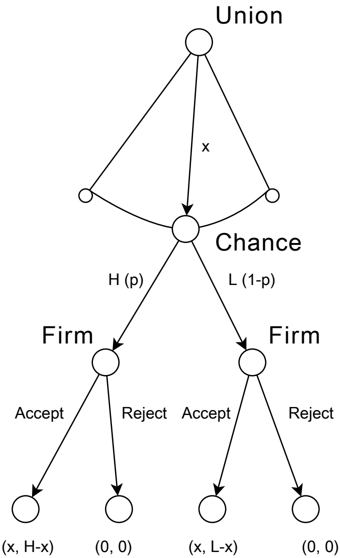

Firm-Union Bargaining with Uncertainty

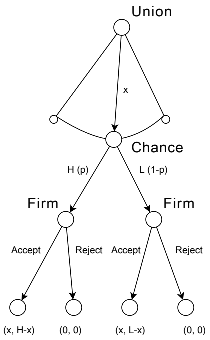

- Note that "Chance" moves after the union but before the firm; this is how we ensure that the firm makes its choice with the knowledge of the outcome of the surplus.
- Again, note that this game has perfect information.
- Assume that in case of a tie in payoffs, the firm chooses *Accept*.

Firm-Union Bargaining with Uncertainty

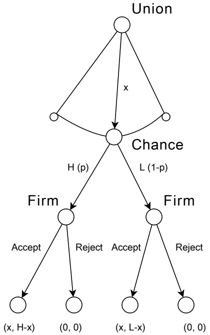

- We can solve this with backwards induction:
  - In the bottom left subgame, take x as given. The firm will choose *Accept* if $H - x \geq 0$.
  - In the bottom right subgame, take x as given. The firm will choose *Accept* if $L - x \geq 0$.

Firm-Union Bargaining with Uncertainty

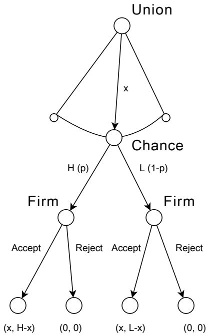

- Taking the firm's actions as given, the union chooses x. Let's consider the possible cases for x:
- Suppose $x \leq L$. Then $x < H$ as well, so the firm will *Accept* if either H or L is the outcome. The expected payoff to the union is: $px + (1 - p)x = x$.

Firm-Union Bargaining with Uncertainty

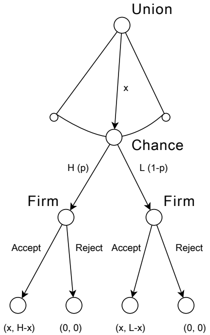

- Suppose $L < x \leq H$. The firm will *Accept* if H is the outcome, but will *Reject* if L is the outcome. The expected payoff is: $px + (1-p)0 = px$.
- Suppose $H < x$. The firm will always *Reject*, so the expected payoff is 0.

Firm-Union Bargaining with Uncertainty

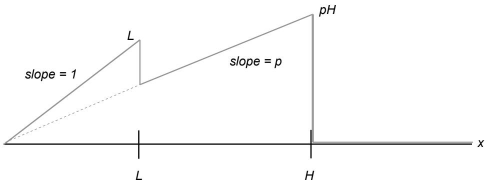

- We can plot the expected payoff as a function of x.
- If $L > pH$, then $x = L$ is the optimal action for the union.
- If $L < pH$, then $x = H$ is the optimal action.
- If $L = pH$, then either is optimal for the union, and there is more than one SPNE outcome.

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

	B	S
B	3,1	0,0
S	0,0	1,3

- Consider the simultaneous-move coordination game BoS.
- We cannot model this as an extensive game with perfect information, because each player does not know what the other player chose when choosing his own action.
- 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.

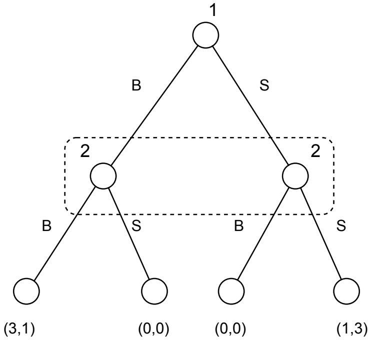

- We denote this with a dashed line connecting or containing the two nodes.
- We call this group of nodes (i.e. a group of histories) an *information set* of Player 2.
- The actions available to a player at all histories (nodes) in an information set *must be the same*.

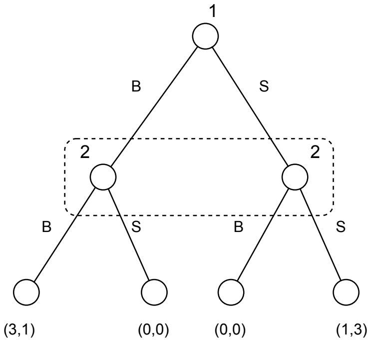

- Note that we could also represent BoS as an extensive game where Player 2 moves first, and Player 1's information set contains B and S.

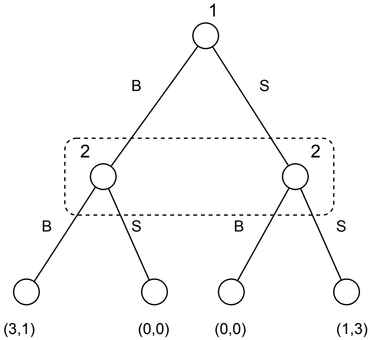

- With imperfect information, our definition of *subgame* must be modified.
- if a subgame contains a node in an information set, it must contain *all* the nodes in that information set (i.e. it makes no sense to have a game where not all the actions are specified).
- In other words, the boundary of a subgame cannot split an information set.

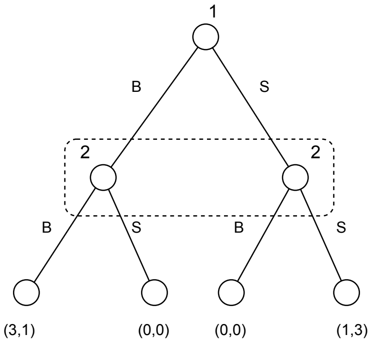

- The extensive form of BoS has only one subgame (the game itself), since we cannot split Player 2's information set.

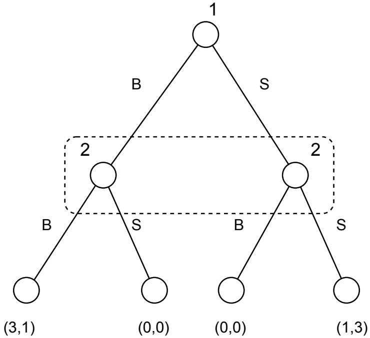

- We can convert extensive games with imperfect information into strategic form, just as before.
- We can also allow for pure and mixed strategies in the strategic form.

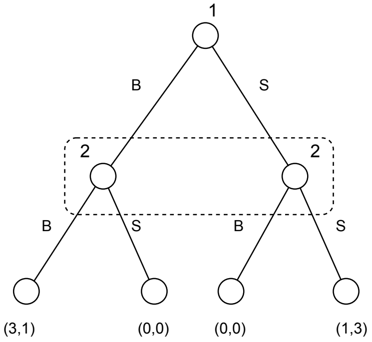

- A pure strategy must specify an action at every information set.
- A mixed strategy must specifiy a *probability distribution over actions* at every information set.
- In the extensive form of BoS, both players have a single information set, at which the set of possible actions is B and S; so a strategy would specify a single distribution over B, S.

- Recall the firm-union bargaining game with exogenous uncertainty we saw earlier. 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

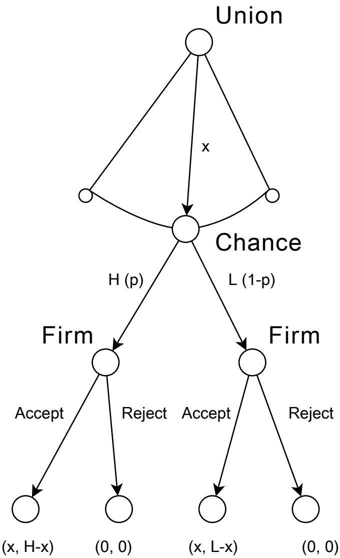

Imperfect Information

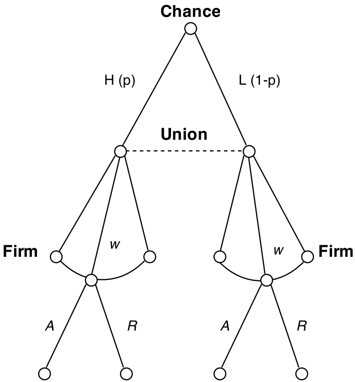

- In this diagram, the firm knows what the surplus is, so the firm's strategy can specify different actions depending on whether L or H occurred.
- The union does not know, so its strategy cannot *condition on* L or H.

Beliefs over Histories

- Many interesting situations involve uncertainty due to a random outcome in the past.
- We can incorporate exogenous uncertainty 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.

Beliefs as a Probability Distribution

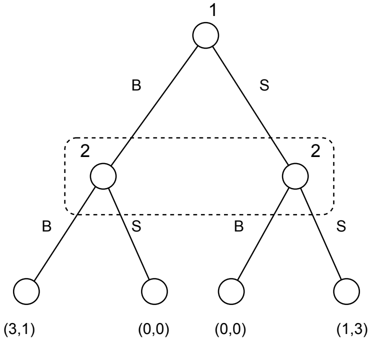

- Let's return to the description of BoS as an extensive game with imperfect information.

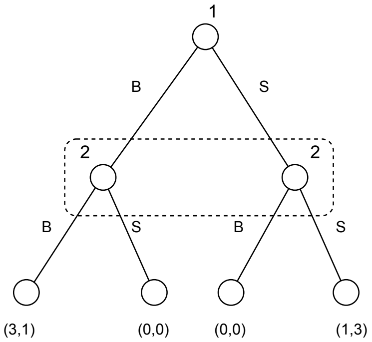

- Consider Player 2's information set. He knows that there are two possibilities for Player 1's action: B or S. Some possible beliefs of Player 2 are:
  - B and S are equally likely.
  - B is more likely to have occurred than S.
  - It is impossible for B to have occurred. S has occurred with certainty.

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 $(1, 0)$.

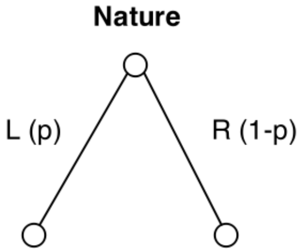

- Let's consider a trivial example: there is only one player (Nature), who chooses L with probability p and R with probability $1 - p$.
- What is the probability that history L occurs? p
- What is the probability that history R occurs? $1 - p$
- The probability that information set $\{L\}$ is reached is p.
- The probability that information set $\{R\}$ is reached is $1 - p$.

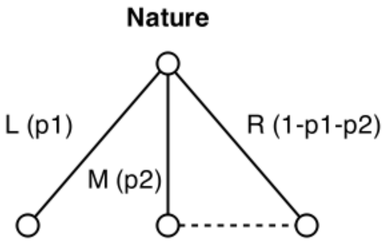

- Now, suppose there are three choices for Nature, but the next player cannot distinguish between M and R.
- What is the probability that history M occurs? p_2
- What is the probability that history R occurs? $1 - p_1 - p_2$
- The probability that information set $\{L\}$ is reached is p_1.
- The probability that information set $\{M, R\}$ is reached is $P(M) + P(R) = 1 - p_1$.

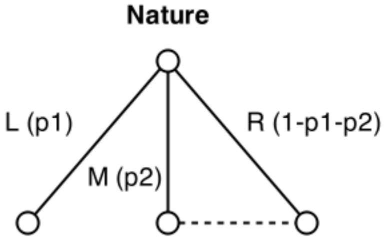

- Suppose we know we have reached information set $\{M, R\}$. What is the probability that we are in history M?
- This is the conditional probability $P(M|\{M, R\})$. From Bayes' Rule, this is:

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

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

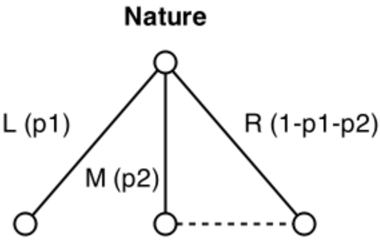

- At information set $\{M, R\}$, the player will have beliefs over the possible histories M, R.
- This can be any probability distribution over M, R.
- We say that beliefs are *consistent* with the given probability distribution of Nature if they match the true probability

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

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

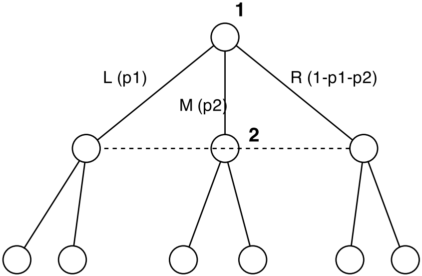

- Here, suppose Player 1 chooses L, M, R randomly according to a known distribution: $P(L) = p_1, P(M) = p_2, P(R) = 1 - p_1 - p_2$.
- Player 1 may be a human player, or it may be Nature.
- Player 2 has one information set, consisting of the histories $\{L, M, R\}$.

Example 1

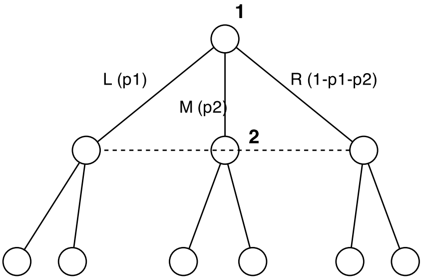

- What is the correct belief at Player 2's information set $\{L, M, R\}$?
- The true probabilities generated by Player 1's behavioral strategy is simply the strategy itself. So, the only correct belief at Player 2's information set is $(p_1, p_2, 1 - p_1 - p_2)$.

Example 2

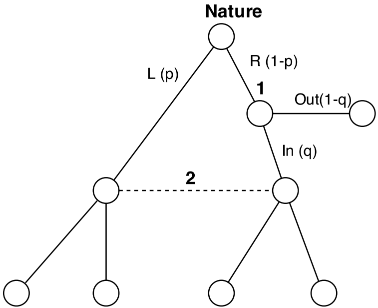

- In this game, Nature goes first and chooses L with probability p, and R with probability $1 - p$.
- If R is chosen, then it is Player 1's move. He can choose *Out* and end the game, or choose *In*.
- Assume his behavioral strategy is to choose *In* with probability q, and *Out* with probability $1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

Example 2

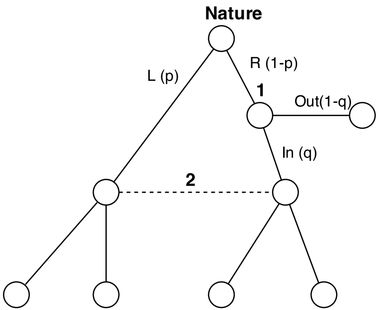

- Player 2 has one information set, consisting of the histories $\{L, (R, In)\}$.
- What is the probability that history L occurs? p
- What is the probability that history R, In occurs?
- Since Nature's choice and Player 1's choice are independent, this is the product of the probabilities of R and In: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

Example 2

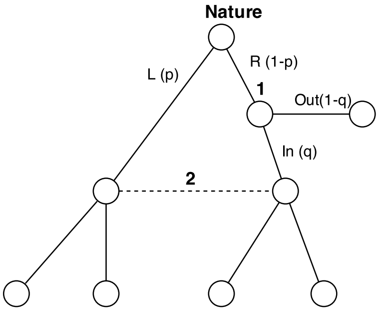

- What is the probability that history R, Out occurs? $(1-p)(1-q)$
- What is the probability that Player 2's information set $\{L, (R, In)\}$ is reached? $P(L) + P(R, In) = p + (1-p)q$

Example 2

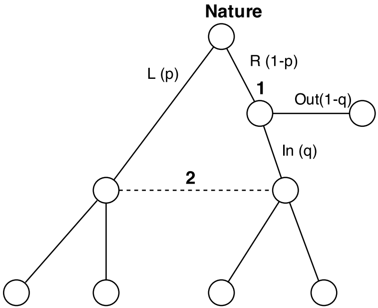

- Conditional on reaching Player 2's information set, what is the probability that L has occurred?

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

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

$(1-p)q$

Example 2

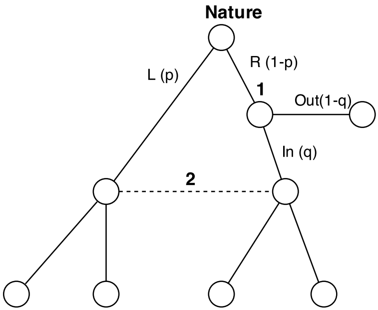

- Correct beliefs for Player 2 would be:

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

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

Example 2

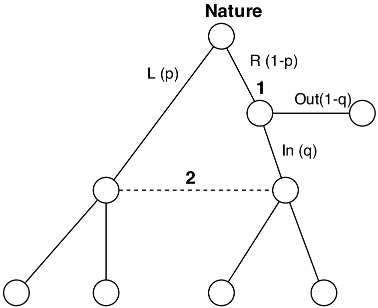

- There is a possible problem here. What if $p = 0$ and $q = 0$?
- The probability of reaching Player 2's information set is 0.
- In this case, Player 2 has no basis on which to form beliefs.

Example 2

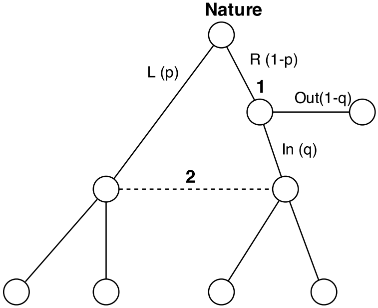

- We will say that *any* probability distribution is a consistent belief at an information set that is reached with zero probability.
- So, any probability distribution $r, 1-r$ over $L, (R, In)$ is consistent with a behavioral strategy profile where $p = 0$ and $q = 0$.

Example 3

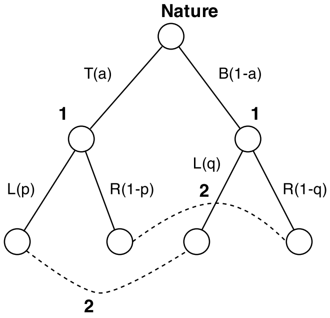

- In this game, Nature goes first and chooses T with probability a, B with probability $1 - a$.
- Player 1 observes this choice, and chooses L or R.
- Player 2 does not observe Nature's choice, but only Player 2's choice.
- This setup is quite common in signaling games. Nature's choice determines the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

Example 3

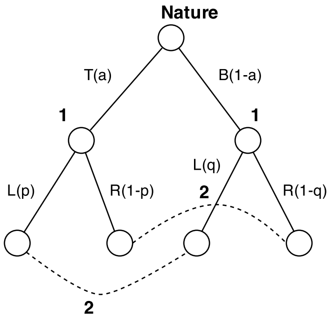

- Player 1 has two information sets: $\{T\}$ and $\{B\}$. At each of these information sets, he has the same two actions L, R.
- Player 2 has two information sets: the information set after L, $\{TL, BL\}$, and the information set after R, $\{TR, BR\}$.

Example 3

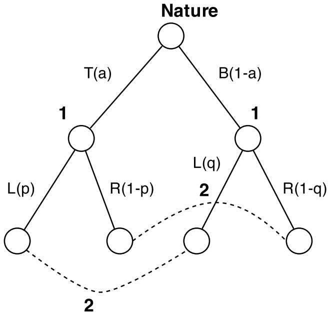

- Suppose that Player 1's behavioral strategy specifies this probability distribution at each of his two information sets:
- $P(L|T) = p, P(R|T) = 1 - p, P(L|B) = q, P(R|B) = 1 - q$
- Here, $P(L|T)$ means "if T occurs, choose L with probability p".

Example 3

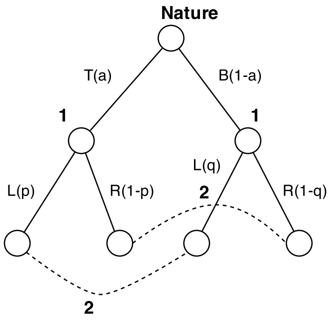

- What is the probability of reaching Player 2's information set $\{TL, BL\}$?

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

- 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

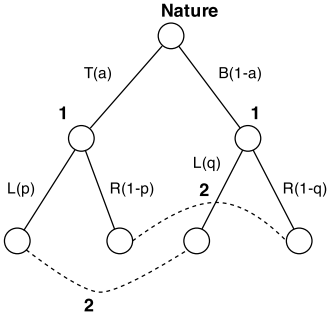

- Suppose we are at information set $\{TL, BL\}$. What is the probability that history TL has occurred?

$$P(TL|\{TL, BL\}) = \frac{P(TL \cap \{TL, BL\})}{P(\{TL, BL\})} = \frac{P(TL)}{P(\{TL, BL\})}$$

$$= \frac{ap}{ap + (1-a)q}$$

Example 3

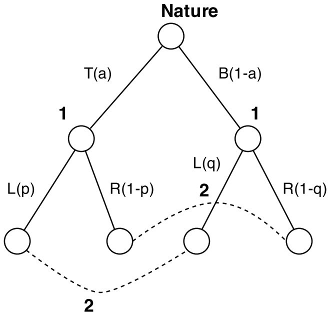

- Therefore, the beliefs at information set $\{TL, BL\}$ that would be consistent with the behavioral strategy of Player 1 are:

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

Example 3

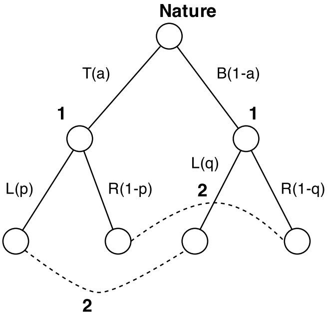

- Similarly, the beliefs at information set $\{TR, BR\}$ that would be consistent with the behavioral strategy of Player 1 are:

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

Example 3

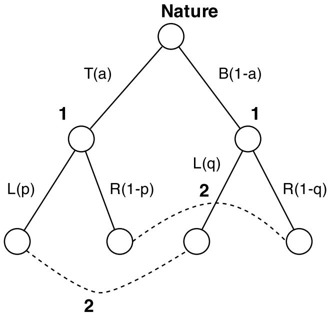

- Information set $\{TL, BL\}$ is reached with zero probability if $p = 0, q = 0$.
- Information set $\{TR, BR\}$ is reached with zero probability if $p = 1, q = 1$.
- Any probability distribution is a consistent belief at an information set reached with zero probability.

Belief System

- 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:
- A **belief system**, or *system of beliefs*, 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.

Consistent Assessment

- An **assessment** is a pair (p, b) consisting of a belief system p and a behavioral strategy profile b.
- **Def. 7.20** An assessment (p, b) is called **consistent** if there is a sequence of completely mixed behavioral strategies, $b^1, b^2, \ldots$ converging to b, such that the associated sequence of Bayes' rule-induced belief systems, $p^1, p^2, \ldots$ converges to p.
- More simply, an assessment (p, b) is consistent if both of these conditions hold:
  - When the players' beliefs are specified by p, then the behavioral strategy profile b is optimal.
  - When the players choose to play the behavioral strategy profile b, then the belief system p is consistent with Bayes' rule.

Sequential Rationality

- **Def 7.21** An assessment (p, b) is **sequentially rational** if for every player i, every information set I belonging to player i, and every behavioral strategy b_i' of player i,

$$v_i(p, b|I) \geq v_i(p, (b_i', b_{-i})|I)$$

- That is: no player *at any point in the game*, ever has an incentive to change his strategy.
- This holds for all information sets, even the ones reached with zero probability.
- Note that beliefs are considered fixed in this definition.
- Finally, we will combine the concepts of *consistency* (behavior determines beliefs) and *rationality* (beliefs determine optimal behvaior).

Sequential Equilibrium

- **Def 7.22**: An assessment (p, b) for a finite extensive form game is a **sequential equilibrium** if it is both *consisntent*, and *sequentially rational*.
- This combines the concepts of *consistency* (behavior determines beliefs) and *rationality* (beliefs determine optimal behavior).
- 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 can still randomize).
- To specify a sequential equilibrium, we need two things:
  - The strategy profile (actions at every information set);
  - The beliefs at every information set.

Example: Entry Game with Preparation

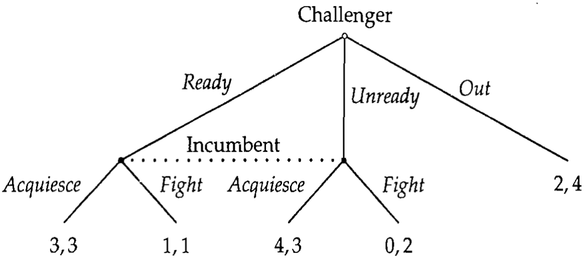

- Consider a modified entry game. The challenger has three choices: stay out, prepare for a fight and enter, or enter without preparation.
- Preparation is costly, but reduces the loss from a fight.
- As before, the incumbent can *Fight* or *Acquiesce*.
- A fight is less costly to the incumbent if the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

Example: 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 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 sequential equilibrium.

Strategic Form

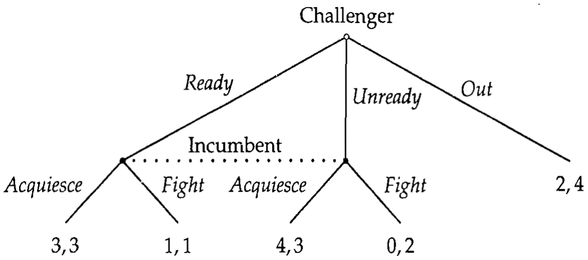

	Acquiesce	*Fight*
Ready	3,3	1,1
Unready	4,3	0,2
Out	2,4	2,4

	Acquiesce	*Fight*
Ready	3,3	1,1
Unready	4,3	0,2
Out	2,4	2,4

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

Adverse Selection

- Let's examine a famous 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.

Adverse Selection

- 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.

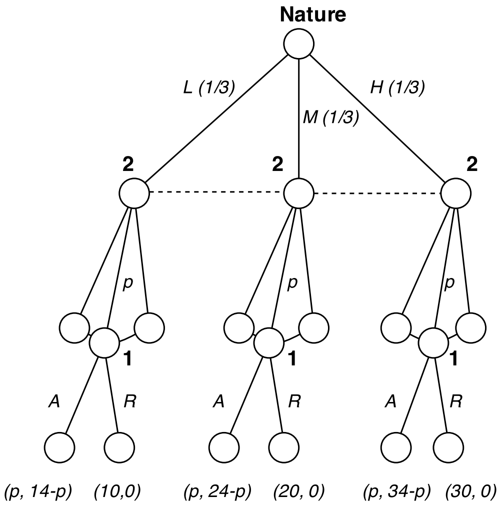
Nature
L (1/3)
M (1/3)
H (1/3)
2
2
2
p
p
p
1
1
1
A
R
A
R
A
R
(p, 14-p)
(10,0)
(p, 24-p)
(20, 0)
(p, 34-p)
(30, 0)

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, let's examine a method by which the seller can credibly communicate the quality of the car.

Signaling Games

- In many situations, Player 1 may know something that Player 2 does not, which would affect Player 2's choice if he knew. For example:
  - When you buy an item from a seller, the seller knows the item's quality, but you do not.
  - When a firm hires an employee, the employee knows his skill level, but the firm does not.
  - When two people get into a competition, each person knows how strong he is, but the other person does not.
- This is called a situation with *asymmetric information*.

Signaling Games

- We can model this kind of situation by assuming there are two (or more) *types* of Player 1, which is chosen by Nature with a known probability distribution.
- Suppose there are two types of Player 1, the "high" type H, and the "low" type L. H-types are more preferable to Player 2.
- Supose Player 1 is a H-type. Then he would like to somehow let Player 2 know that he is a H-type.
- On the other hand, suppose Player 2 is a L-type. Then he would like to imitate the H-type, and make Player 2 believe that he is type H.
- A H-type Player 1 could simply say that he is H-type, but why should Player 2 believe it? A L-type player could do exactly the same thing.
- However, if there was some test that L-types found more costly to pass than H-types, then perhaps L-types would rationally choose not to take the test at all.

Signaling Games

- These games are called *signaling games*.
- At the beginning, Nature chooses Player 1's type according to a known distribution.
- Player 1 can choose to send a *costly signal* that Player 2 observes.
- Player 2 then chooses his action, taking Player 1's choice into account.
- In a *pooling equilibrium*, H and L types of Player 1 behave the same way, and Player 2 cannot distinguish between them.
- In a *separating equilibrium*, H and L types behave differently, and Player 2 can tell them apart.

Entry Game with Signaling

- Consider this variant of the Entry Game.
- There are two types of *Challenger*, *Strong* and *Weak*, which occur with probabilities p and $1 - p$.
- The *Challenger* knows his type, but the *Incumbent* does not.
- *Challenger* can choose to be *Ready* or *Unready* for competition. Choosing *Ready* is more costly for a *Weak*-type.
- The *Incumbent* can choose to fight, F, or acquiesce, A.
- *Incumbent* prefers to fight a *Weak* challenger, but prefers to acquiesce to a *Strong* challenger.
- If the *Incumbent* chooses to fight, both types of challenger suffer the same cost.

Entry Game with Signaling

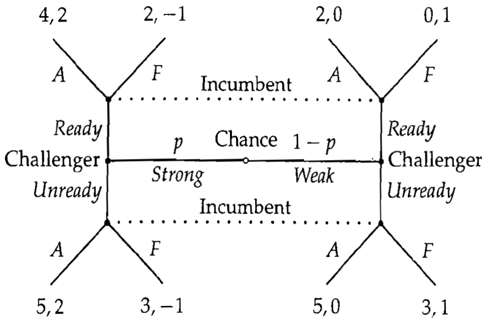

- The start of the game is in the center.
- The first number is the *Challenger*'s payoff. *Challenger*'s payoff is 5 if *Incumbent* chooses *A*.
- Cost of *Ready* is 1 for *Strong*-type, 3 for *Weak*-type.

Entry Game with Signaling

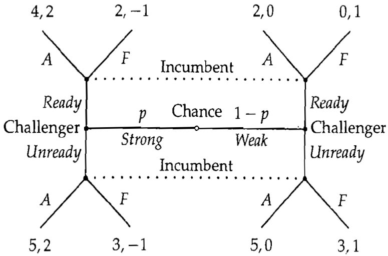

- *Challenger* has two information sets: $\{Strong\}$ and $\{Weak\}$.
- *Incumbent* has two information sets: $\{(Strong, Ready), (Weak, Ready)\}$ and $\{(Strong, Unready), (Weak, Unready)\}$.
- Each player has 2 actions at each information set, so a total of 4 pure strategies.

Entry Game with Signaling

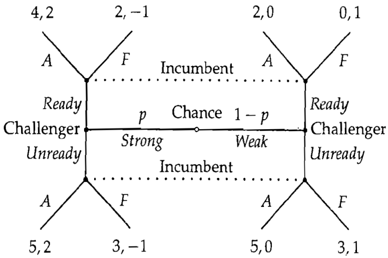

- *Challenger*'s pure strategies: RR, RU, UR, UU
- *Incumbent*'s pure strategies: $(A|R, A|U), (A|R, F|U), (F|R, A|U), (F|R, F|U)$ where the notation $(A|R, A|U)$ means: choose A conditional on R, choose A conditional on U
- Suppose $0 < p < 1$, so both of *Incumbent*'s information sets are reached with positive probability. Let's calculate the expected payoffs to pure strategy profiles.

Entry Game with Signaling

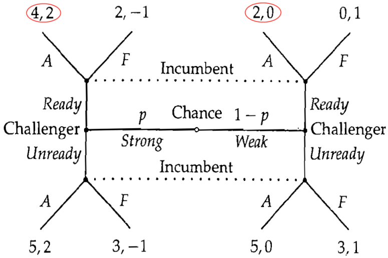

- For the strategy profile $(RR, (A|R, A|U))$, the outcome will be $(Strong, R, A)$ with probability p, and $(Weak, R, A)$ with probability $(1 - p)$.
- Expected payoffs are $(4p + 2(1 - p), 2p)$.

Entry Game with Signaling

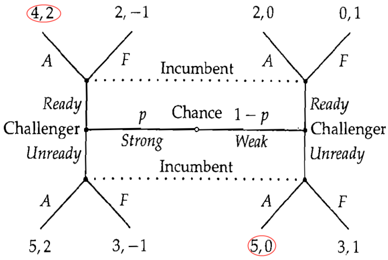

- For the strategy profile $(RU, (A|R, A|U))$, the outcome will be $(Strong, R, A)$ with probability p, and $(Weak, U, A)$ with probability $(1 - p)$.
- Expected payoffs are $(4p + 5(1 - p), 2p)$.

Expected Payoffs

	RR	RU
$A\|R, A\|U$	p(4,2) + (1-p)(2,0)	p(4,2) + (1-p)(5,0)
$A\|R, F\|U$	p(4,2) + (1-p)(2,0)	p(4,2) + (1-p)(3,1)
$F\|R, A\|U$	p(2,-1) + (1-p)(0,1)	p(2,-1) + (1-p)(5,0)
$F\|R, F\|U$	p(2,-1) + (1-p)(0,1)	p(2,-1) + (1-p)(3,1)

	UR	UU
$A\|R, A\|U$	p(5,2) + (1-p)(2,0)	p(5,2) + (1-p)(5,0)
$A\|R, F\|U$	p(3,-1) + (1-p)(2,0)	p(3,-1) + (1-p)(3,1)
$F\|R, A\|U$	p(5,2) + (1-p)(0,1)	p(5,2) + (1-p)(5,0)
$F\|R, F\|U$	p(3,-1) + (1-p)(0,1)	p(3,-1) + (1-p)(3,1)

Expected Payoffs with $p = 0.5$

	$A|R, A|U$	$A|R, F|U$	$F|R, A|U$	$F|R, F|U$
RR	3,1	3,1	1,0	1,0
RU	4.5, 1	3.5, 1.5	3.5, -0.5	2.5, 0
UR	3.5, 1	2.5, -0.5	2.5, 1.5	1.5, 0
UU	5, 1	3, 0	5, 1	3, 0

- The NE are: $(RU, (A|R, F|U)), (UU, (A|R, A|U))$, and $(UU, F|R, A|U))$. Let's go through each of these and see which ones can be part of a SE.

Check if NE is a SE

- Consider the NE $(RU, (A|R, F|U))$.
- The beliefs at *Incumbent*'s information set after R must be $(1, 0)$ (probability 1 on *Strong*) since *Challenger*'s strategy will only choose R if Nature chose *Strong*.
- The beliefs at *Incumbent*'s information set after U must be $(0, 1)$ (probability 0 on *Strong*).
- The actions specified for *Incumbent* at this NE are clearly optimal given these beliefs, so this is a SE.

Check if NE is a SE

- Consider the NE $(UU, (A|R, A|U))$.
- The beliefs at *Incumbent*'s information set after U must be $(p, 1-p) = (0.5, 0.5)$.
- Clearly, *Challenger*'s actions specified by this NE are optimal at this information set, since it is a NE.
- The information set after R is reached with zero probability, so any beliefs are consistent.
- Denote the beliefs at this information set as $(q, 1-q)$, where q is the probability on *Strong*.
- We want to find the range of q that makes A optimal, i.e. $E(A) \geq E(F)$, which is true if:

$$q2 + (1-q)0 \geq q(-1) + (1-q)1$$

- or if $q \geq \frac{1}{4}$.

Check if NE is a SE

- Consider the NE $(UU, F|R, A|U)$.
- For the information set after U, everything is the same as in the previous case.
- The information set after R is reached with zero probability, so any beliefs are consistent.
- Denote the beliefs at this information set as $(q, 1-q)$, where q is the probability on *Strong*.
- We want to find the range of q that makes F optimal, i.e. $E(A) \leq E(F)$, which is true if:

$$q2 + (1-q)0 \leq q(-1) + (1-q)1$$

- or if $q \leq \frac{1}{4}$.

Separating & Pooling Equilibria

- There is one *separating* equilibrium: $(RU, (A|R, F|U))$, and *Incumbent*'s beliefs are $(1, 0)$ after R and $(0, 1)$ after U.
- In a separating equilibrium the different types of Player 1 choose different actions, so Player 2 knows which type they are by their choice.
- There is a set of *pooling* equilibria: $(UU, (A|R, A|U))$ and $(UU, F|R, A|U))$.
- In both of these cases, the information set after R is reached with zero probability, so any beliefs are consistent.
- If the beliefs are such that A to be optimal, then the SE is $(UU, (A|R, A|U))$; otherwise, it is $(UU, F|R, A|U))$.
- In a pooling equilibrium, the different types of Player 1 behave the same way, and Player 2 cannot distinguish them.

- Why do students obtain a college degree?
- One reason is that the knowledge they gain in college will increase their skills and abilities.
- However, there is another possible reason: perhaps students use degrees to *differentiate* themselves from other students when applying for jobs.
- This can hold even if the degree itself does not increase ability.
- We model this as a signaling game.

Education as a signal of ability

- Suppose the ability level of a worker can be measured by a single number
- There are two types of workers: "high" and "low"-ability workers, denoted H and L, with $L < H$.
- Type is known to the workers, but cannot be directly observed by employers.
- Workers can choose to obtain some amount of education, which has no effect on ability, but costs less for the H-type worker.

Education as a signal of ability

- The sequence of the game is as follows:
  - Chance chooses the type of the worker at random; the probability of H is p.
  - The worker, who knows his type, chooses an amount of education $e \geq 0$. The cost of education is different according to type; for a L-type worker, the cost is e/L; for a H-type worker, it is e/H.
  - Two firms observe the worker's choice of e (but not his type), and simultaneously offer two wages, w_1 and w_2.
  - The worker chooses one of the wage offers and works for that firm. The worker's payoff is his wage minus the cost of education. The firm that hires the worker gets a payoff of the worker's ability, minus the wage. The other firm gets a payoff of 0.

Game Tree

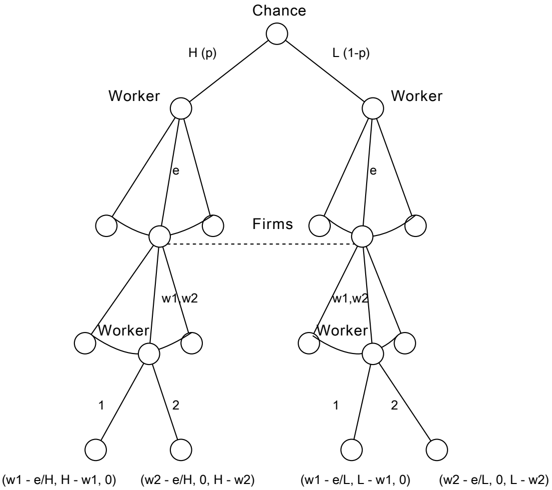
Chance
H (p)
L (1-p)
Worker
Worker
e
e
Firms
w1,w2
w1,w2
Worker
Worker
1
2
1
2
(w1 - e/H, H - w1, 0)
(w2 - e/H, 0, H - w2)
(w1 - e/L, L - w1, 0)
(w2 - e/L, 0, L - w2)

Finding SE

- We claim there is a sequential equilibrium in which a H-type worker chooses a positive amount of education, and a L-type worker chooses zero education.
- Consider this assessment (i.e. beliefs plus strategies), where e^* is a positive number (to be determined):
  - *Worker's strategy*: Type H chooses $e = e^*$ and type L chooses $e = 0$. After observing $w1, w2$, both types choose the highest offer if $w1, w2$ are different, and firm 1 if they are the same.
  - *Firms' belief*: Each firm believes that a worker is type H if he chooses $e = e^*$, and type L otherwise.
  - *Firms' strategies*: Each firm offers the wage H to a worker who chooses $e = e^*$, and L to a worker who chooses any other value of e.

Finding SE

- Let's check that the conditions for consistency of beliefs and optimality of strategies are satisfied.
  - Consistency of beliefs: take the worker's strategy as given.
  - The only information sets of the firm that are reached with positive probability are after $e = 0$ and $e = e^*$; at all the rest, the firms' beliefs may be anything.
  - At the information set after $e = 0$, the only correct belief is $P(H|e = 0) = 0$.
  - At the information set after $e = e^*$, the only correct belief is $P(H|e = e^*) = 1$. So these beliefs are consistent.
  - Optimality of firm's strategy: Each firm's payoff is 0, given its beliefs and strategy.
  - If a firm deviates by offering a higher wage, it will make a negative profit.
  - If it deviates by offering a lower wage, it gets a payoff of 0 since the worker will choose the other firm. So, there is no incentive to deviate.

Finding SE

- Optimality of worker's strategy: In the last subgame, the worker's strategy of choosing the higher wage is clearly optimal. Let's consider the worker's choice of e:
  - Type H: If the worker maintains the strategy and chooses $e = e^*$, he will get a wage offer of H and his payoff will be $H - \frac{e^*}{H}$.
  - If the worker deviates and chooses any other e, he will get a wage offer of L and his payoff will be $L - \frac{e}{H}$.
  - The highest possible payoff when deviating is when $e = 0$, which gives a payoff of L.
  - Therefore, in order for our hypothetical equilibrium to be optimal, we need $H - \frac{e^*}{H} \geq L$, or

$$e^* \leq H(H - L)$$

Finding SE

- Type L: If the worker maintains the strategy and chooses $e = 0$, he will get a wage offer of L and his payoff will be L.
- If the worker deviates and chooses anything but e^*, he still gets a wage offer of L and a lower payoff of $L - \frac{e}{L}$.
- If the worker deviates and chooses e^* (i.e. imitates a H-type) then he gets a wage offer of H, for a total payoff of $H - \frac{e^*}{L}$.
- For our hypothetical equilibrium to be optimal, we need $L \geq H - \frac{e^*}{L}$ or

$$e^* \geq L(H - L)$$

Conditions for Equilibria

- Combining these requirements, the condition for this equilibrium to be optimal is:
$$L(H - L) \leq e^* \leq H(H - L)$$
- If this is satisfied, then separating equilibria exist in which H-type workers can be distinguished from L-type workers by their choice of e.
- This is not the only type of equilibrium that exists: there may also exist pooling equilibria, given the same values of H and L, in which both types of workers choose the same amount of education.

Administrative Stuff

- The final exam will be on June 23 from 13:30-16:30 in Boxue 504.
- It will cover the material after the midterm exam and is written jointly with the other two microeconomics classes.