Introduction to Game Theory
Lecture Note 8: Dynamic Bayesian Games

Haifeng Huang

University of California, Merced

Spring 2020

Basic terminology

- Now we study **dynamic Bayesian games**, or **dynamic/extensive games of incomplete information**, as opposed to the static (simultaneous-move) games of incomplete information in the last lecture note.
- **Incomplete information**: a player does not know another player's characteristics (in particular, preferences); **imperfect information**: a player does not know what actions another player has taken.
- Recall that in a dynamic game of perfect information, each player is perfectly informed of the history of what has happened so far, up to the point where it is her turn to move.

Harsanyi Transformation

- Following Harsanyi (1967), we can change a dynamic game of **incomplete** information into a dynamic game of **imperfect** information, by making nature as a mover in the game. In such a game, nature chooses player *i*'s type, but another player *j* is not perfectly informed about this choice.
- But first, let's look at a dynamic game of **complete but imperfect** information.

A dynamic game of complete but imperfect information

- An entry game: the challenger (she) may stay out, prepare for combat and enter (ready), or enter without preparation (unready). Each player's preferences are common knowledge.

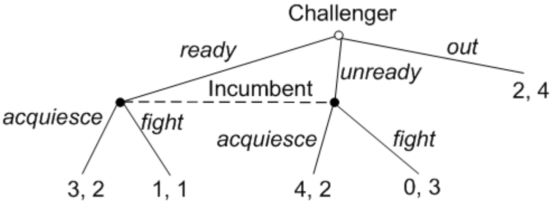

- The dashed line indicates that after the history "ready" and the history "unready", the incumbent does not know whether the challenger has chosen ready or unready.
- Whether the incumbent (he) should choose A or F depends on his belief about what the challenger has chosen.

Information set

- {ready, unready} is an information set of the incumbent.

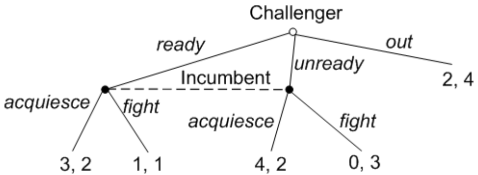

- Definition: An **information set** of a player is a collection of decision nodes (or histories) satisfying the following two conditions:
  1. the player has the move at every node in the information set;
  2. when the play of the game reaches a node in the information set, the player with the move does not know which node in the information set has been reached, *unless the information set is a singleton* (containing only one decision node).

Strategies and information set

- The incumbent has one information set in the game, and the challenger also has one information set, after history $\emptyset$.

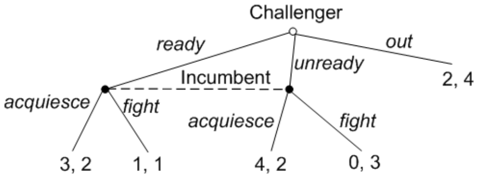

  - ▷ A game in which every information set of every player contains a singleton is a game of perfect information.
- A (pure) **strategy** of player *i* in a dynamic game is a function that assigns to each of *i*'s information sets an action in the set of actions available to player *i* at that information set.

A dynamic game of incomplete information

- Now suppose the challenger (she) can have two types: strong (with prior probability p) or weak (with prior probability $1 - p$). The incumbent (he) observes the challenger's action, but not her type.
- This is a game of *incomplete* information. But we can change it into a game of *imperfect* information by letting nature have the initial move of choosing the type of the challenger:

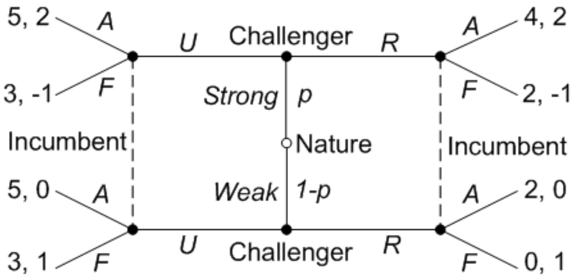

Information sets and strategies

- The challenger has two information sets, {Strong} and {Weak}, at each of which she has two possible actions: $U(nready)$ and $R(ready)$.
- So the challenger has four strategies: (1) R after Strong and R after Weak; (2) R after Strong and U after Weak; (3) U after Strong and R after Weak; (4) U after Strong and U after Weak.

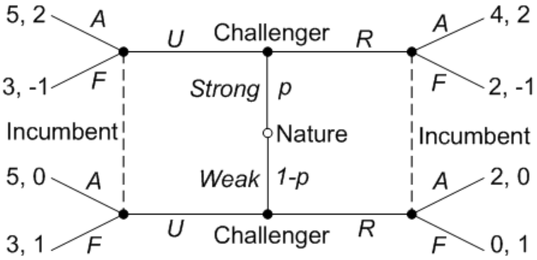

Information sets and strategies (cont.)

- The incumbent also has two information sets, {(Strong, R), (Weak, R)} and {(Strong, U), (Weak, U)}, at each of which he has two choices: *F*(*ight*) or *A*(*cquiesce*).
- So the incumbent also has four strategies: (1) *A* after *R* and *A* after *U*; (2) *A* after *R* and *F* after *U*; (3) *F* after *R* and *A* after *U*; (4) *F* after *R* and *F* after *U*.

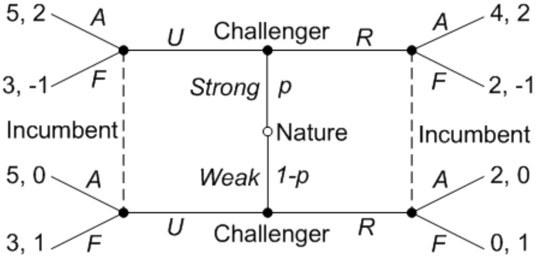

Definitions: belief system; behavioral strategy

- A **belief system** in an extensive game is a function that assigns to *each information set of each player* a probability distribution over the histories (or decision nodes) in that information set.
- A **behavioral strategy** of player i in an extensive game is a function that assigns to each of i's information set (denoted as I_i) a probability distribution over the set of actions to player i at that information set (denoted as $A(I_i)$), with the property that each probability distribution is independent of every other distribution.
  - Difference with mixed strategy: a mixed strategy refers to a probability distribution over pure strategies, whereas a behavioral strategy refers to the collection of probability distributions over the actions at the information sets.
  - A behavioral strategy that assigns probability 1 to a single action at every information set is equivalent to a pure strategy.

Assessment and equilibrium

- An **assessment** in an extensive game is a pair consisting of (1) a profile of (behavioral) strategies and (2) a belief system.
- An assessment constitutes an equilibrium if it satisfies the following two conditions:
  1. **Sequential rationality**: each player's strategy is optimal whenever she has to move, given her beliefs *and* the other players' strategies.
     - ▷ The strategy has to be optimal in every information set, regardless of whether that information set is reached if the players follow their strategies.
     - ▷ Similarity with and difference from SPNE.
  2. **Consistency of beliefs with strategies**: each player's belief is consistent with the strategy profile.
     - ▷ Each player's belief must be correct in equilibrium.

Sequential rationality

- In the game below, player 1 will select J after history (C, F). Suppose player 1's choice at the beginning is E, and player 2's belief at his information set is that with probability $\frac{2}{3}$ player 1 has chosen C and with probability $\frac{1}{3}$ she has chosen D.

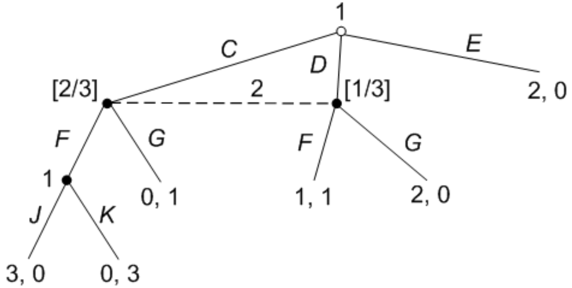

- Sequential rationality requires player 2 to select G over F at that information set since $\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 > \frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 1$.

(Weak) consistency of beliefs with strategies

- Each player's belief must be correct: the player's assignment of probability to any history must be the probability with which that history occurs if the players adhere to their strategies.
  - At an information set that is reached with probability 0 if the players follow their strategies, the player that moves at that information set can hold any belief.
  - Some equilibrium refinement notions would specify certain requirements for such information sets, but we will not consier them here.
- Denoting an information set by I_i and the strategy profile by β, then the probability player i assigns to a particular history h^* at I_i is

$$\frac{P(h^* \text{ according to } \beta)}{\sum_{h \in I_i} P(h \text{ according to } \beta)}. \tag{1}$$

(Weak) consistency of beliefs with strategies: example

- If player 1's strategy is *EJ*, player 2 can hold any belief at her information set.

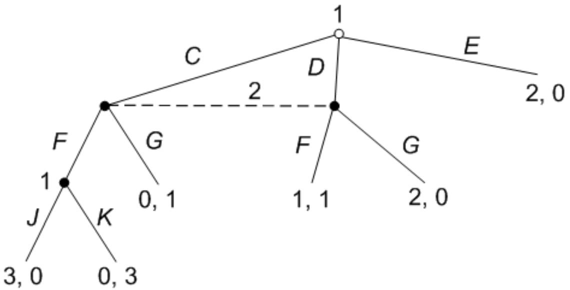

(Weak) consistency of beliefs with strategies: example (cont.)

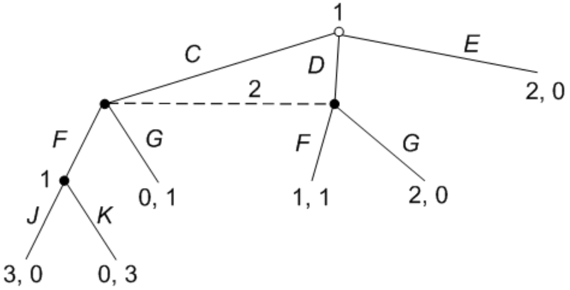

- If player 1's strategy at her first information set is to choose C with probability p, D with probability q, and E with probability $1 - p - q$, then player 2 must assign probability $\frac{p}{p+q}$ to history C and $\frac{q}{p+q}$ to history D.
  - If player 1 chooses D with probability 1, then player 2's belief must assign probability 0 to C and 1 to D.

Summary: weak sequential equilibrium

- Denote a behavioral strategy profile as β and a belief system as μ.
- Definition: an assessment (β, μ) is a **weak sequential equilibrium** if it satisfies the following two conditions:
  1. **Sequential rationality**: for each player i and each information set I_i of player i, her expected payoff to the probability distribution $O_{I_i}(\beta, \mu)$ over terminal histories generated by her belief μ_i and I_i and the behavior prescribed subsequently by the strategy profile β is at least as large as her expected payoff to the probability distribution $O_{I_i}((\gamma_i, \beta_{-i}), \mu)$ generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile (γ_i, β_{-i}), for each of her behavioral strategies γ_i.
  2. **Weak consistency of beliefs with strategies**

Summary: weak sequential equilibrium (cont.)

- Definition: an assessment (β, μ) is a **weak sequential equilibrium** if it satisfies the following two conditions:
  1. **Sequential rationality**
  2. **Weak consistency of beliefs with strategies**: for every information set I_i reached with positive probability given the strategy profile β, the probability assigned by the belief system to each history h^* in I_i is given by (1).

Weak sequential equilibrium: example

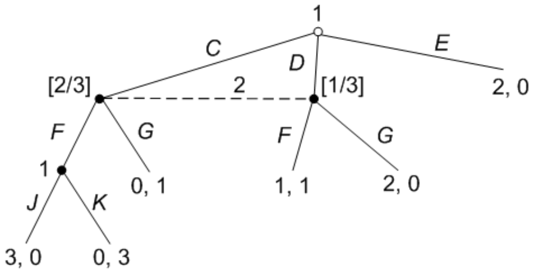

- Does the game have a weak sequential equilibrium in which the strategy profile is (EJ, G) and player 2's belief is that he assigns $\frac{2}{3}$ to history C and $\frac{1}{3}$ to history D? Yes.
  - Player 1's strategy EJ is sequentially rational given player 2's G. Player 2's strategy G is also sequentially rational given his beliefs and player 1's strategy EJ.
  - Player 2's belief is (weakly) consistent with the strategy profile (EJ, G).

Weak sequential equilibrium: example (cont.)

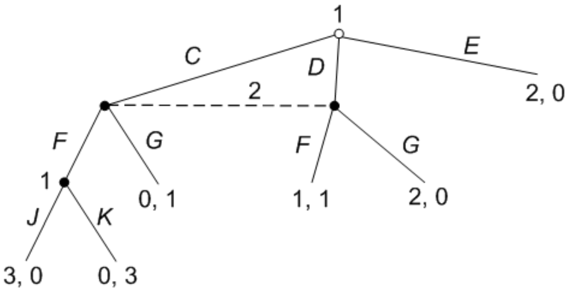

- But there is no weak sequential equilibrium in which the strategy profile is (DJ, G).
  - Player 1's strategy DJ is sequentially rational.
  - But given player 1's strategy, player 2 should believe the history is D with probability 1 at her information set, and should therefore choose F rather than G.

Signaling game

- Now we can analyze the incomplete information entry game, which is an example of **signaling game**, an important class of dynamic games of incomplete information.
- **Signaling game**: Some players are informed about variables that affect everyone while others are not. The informed players ("sender") take actions first, and the uninformed players ("receiver") take actions after observing the informed players' actions. The informed players' actions may "signal" their information (e.g., their types).

The entry game as a signaling game

- In the entry game, the incumbent does not know if the challenger is strong or weak. The challenger decides whether or not to prepare herself for entry (R or U). If U, the challenger receives a payoff of 5 if the incumbent acquiesces and 3 if the latter fights. *Preparations cost a strong challenger 1 and a weak challenger 3.*

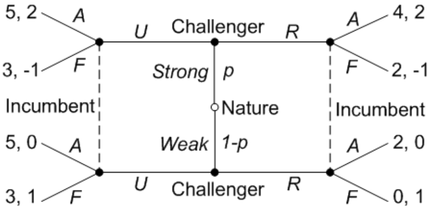

- Whether a challenger prepares herself or not *may* say something about her type.

Pure-strategy weak sequential equilibria of the entry game (1)

- First note that a weak challenger prefers U regardless of the incumbent's action. So in any equilibrium a weak challenger chooses U.

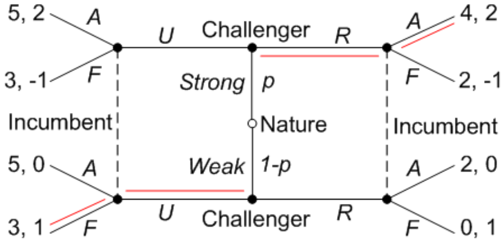

- If a strong challenger chooses R in equilibrium, then the incumbent knows that a challenger that chooses R is strong, and he knows that a challenger that chooses U is weak. So the incumbent chooses A after R and F after U.

Separating equilibrium of the entry game

- Note the incumbent's strategy is sequentially rational and her belief is consistent with strategy.
- Given that the incumbent will choose A after observing R and F after observing U, a strong challenger will not deviate from R $(4 > 3)$; a weak challenger will not deviate from U $(1 > 0)$.
    - ▷ The challenger's strategy is sequentially rational.
- Therefore there is a weak sequential equilibrium in which a weak challenger chooses U, a strong challenger chooses R, and the incumbent chooses A after observing R and chooses F after observing U.
- This is called a **separating equilibrium**: each type of the sender chooses a different action, so that upon observing the sender's action, the receiver knows the sender's type.

Pure-strategy weak sequential equilibria of the entry game (2)

- If a strong challenger chooses U too in equilibrium, then by consistency of the belief, the incumbent believes a challenger that has chosen U is strong with probability p and weak with probability $1-p$. So A is optimal for the incumbent if $2p + 0 \cdot (1-p) \geq -1 \cdot p + 1 \cdot (1-p)$, i.e., $p \geq \frac{1}{4}$; and F is optimal if $p < \frac{1}{4}$.

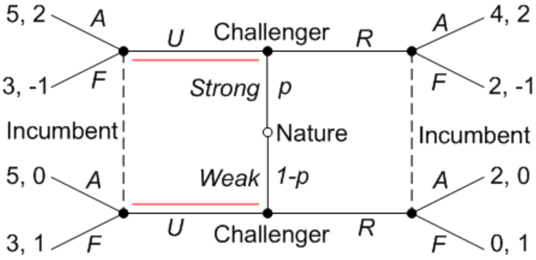

Pooling equilibria of the entry game (1)

- Does a strong challenger want to deviate?
- If $p \geq \frac{1}{4}$, a strong challenger gets 5 by sticking to the strategy U; if she deviates to R, she cannot get a higher payoff regardless of the incumbent's action.
  - ▷ This is indeed a weak sequential equilibrium.

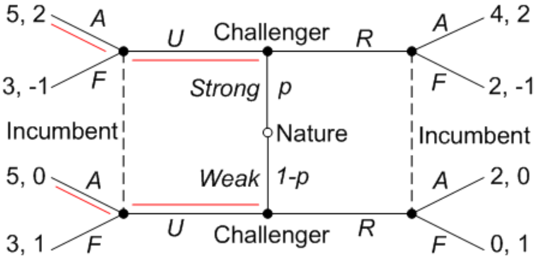

Pooling equilibria of the entry game (2)

- If $p < \frac{1}{4}$, a strong challenger gets 3 by sticking to the strategy U. If she deviates to R, she gets 4 if the incumbents *acquiesces* and 2 if he *fights*. Thus for the strong challenger to have no incentive to deviate, it must be the case that the incumbent will *fight* when he observes the challenger has somehow chosen R (even though the equilibrium says she should choose U regardless of her type).

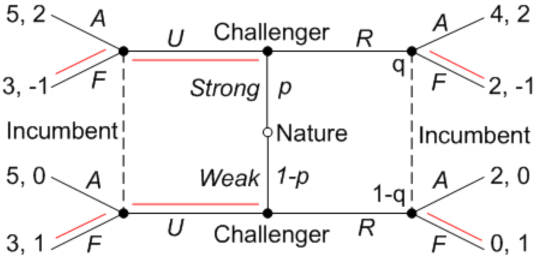

Pooling equilibria of the entry game (3)

- What makes the incumbent *fight* upon observing *R*?

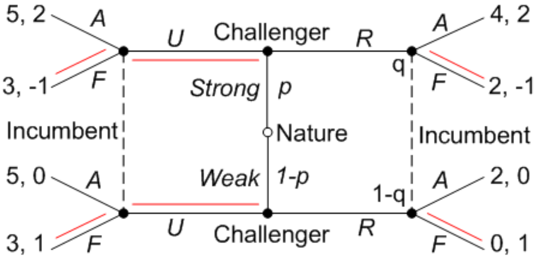

- The incumbent will choose *fight* upon observing *R* if he believes the probability that a challenger that has chosen *R* is strong, denoted as q, is such that $-1 \cdot q + 1 \cdot (1 - q) \geq 2q + 0 \cdot (1 - q) \Rightarrow q \leq \frac{1}{4}$.

Pooling equilibria of the entry game (4)

- So there is a weak sequential equilibrium in which both types of challenger choose U, and the incumbent chooses F upon observing U and F upon observing R, with $p < \frac{1}{4}$ and $q \leq \frac{1}{4}$.
- This is called a **pooling equilibrium**: all types of the sender choose the same action, so that the sender's action gives the receiver no information about the sender's type.
- There can also be semi-pooling, semi-separating equilibria in signaling games.

Job market signaling

- This is a famous signaling model, due to Spence (1973).
- There are a worker and two firms. The worker can either have high ability (H) or low ability (L), meaning the payoff she brings to her employer is either H or L, $H > L$. The worker's type is known to herself but not to the firms, and the prior probability of a worker being type H is π.
- The worker chooses the amount e of education to obtain. The cost of obtaining e of education is e/H for a high ability worker and e/L for a low ability worker.
- The firms, observing e, simultaneously offer wages w_1 and w_2. Finally, the worker chooses a firm.

Game tree

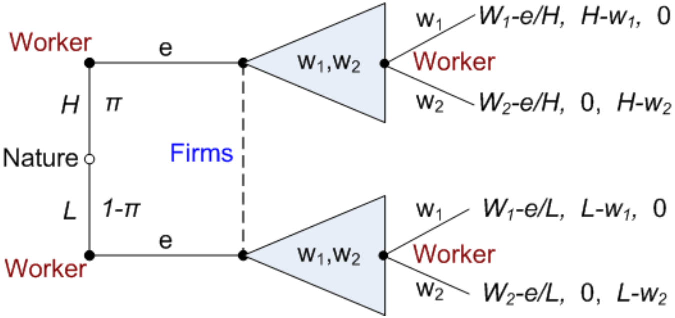

- Is there a separating equilibrium, in which the two types of workers choose different amount of education?
- Is there a pooling equilibrium in which both types of worker choose the same amount of education?

Separating equilibria (1)

- Let's suppose that in a separating equilibrium a high type worker chooses $e^* > o$ and a low type worker chooses 0 (why?).
- Then each firm believes that a worker is type H if she chooses e^* education and type L otherwise (such a belief is weakly consistent with the strategy profile).
- Each firm then offers the wage H to a worker with e^* education and the wage L to a worker with any other value of education (why?).
- The worker chooses whichever firm that offers a higher wage, or randomly chooses one if the wages are equal.
- The firms' strategy is clearly optimal. What will be value of a type H worker's e^*?

Separating equilibria (2)

- For type H not to deviate to 0, $w - e/H = H - e/H \geq L$, which means
$$e^* \leq H(H-L).$$
- For type L not to deviate to e^*, $w - e/L = L \geq H - e/L$,which means
$$e^* \geq L(H-L).$$
- Therefore, there is a separating weak sequential equilibrium in which type H chooses a education level e^* such that
$$L(H-L) \leq e^* \leq H(H-L).$$

Separating equilibria (3)

- Again, there is a separating equilibrium in which $L(H-L) \leq e^* \leq H(H-L)$.
- Intuition: type H chooses an education level low enough so that she is still profitable but high enough that type L is not willing to imitate her signal.
- In the model education itself does not add to the worker's productivity, but a high type worker obtains it in order to distinguish herself from a low type worker.

Pooling equilibria (1)

- Suppose that both types of worker choose the same education level e^*. Then the firms believe a worker with e^* is type H with probability π and type L with probability $1-\pi$.
- Thus the firms will each offer a worker with e^* a wage equal to $\pi H + (1-\pi)L$.
- A type H worker's payoff is then $\pi H + (1-\pi)L - e^*/H$, and a type L worker's payoff is then $\pi H + (1-\pi)L - e^*/L$.
- Suppose firms believe that a worker that selects any other level of education is type L (this is weakly consistent with the worker's strategy and supports the widest range of equilibrium values of e^* since this belief makes it least profitable for a worker to deviate).

Pooling equilibria (2)

- If $e^* = 0$, obviously neither type of worker will deviate.
- If $e^* > 0$, then the most profitable level of education a worker can deviate to is 0.
- For neither type to have an incentive to deviate, it must be that $\pi H + (1-\pi)L - e^*/H \geq L$, and $\pi H + (1-\pi)L - e^*/L \geq L$.
- Therefore $\pi H + (1-\pi)L - e^*/L \geq L \Rightarrow e^* \leq \pi L(H-L)$.

Pooling equilibria (3)

- In other words, there is a pooling equilibrium in which both types of worker choose a level of education $\leq \pi L(H-L)$, and the worker is paid $\pi H+(1-\pi)L$. In such an equilibrium, the firms believe that a worker selecting any other value of education (including an $e>e^*$) is type L.
  - There is something unnatural about this belief for $e>e^*$, which some refinement of the weak sequential equilibrium can deal with.
- Since $\pi L(H-L)<L(H-L)$, the education levels in a pooling equilibrium is lower than those in a separating equilibrium.

A model of cheap talk

- In the signaling models we have just seen, signals are costly, and so can sometimes be used to distinguish different types of players. What if signals are costless to send (cheap talk)?
- Consider the following legislative game: the amount of military expenditure (w) that is objectively needed for national defense depends on the state of the world, which is uniformly distributed on $[0, 1]$. A committee (C) knows the state of the world, but the floor of the House (F) does not.
- C sends a message to F about the state of the world, and then F chooses the level of expenditure in $[0, 1]$.

Payoffs

- F prefers the expenditure to be what is objectively required (w), but C wants a higher level of defense expenditure since it is captured by the military-defense industry complex, so it has incentive to exaggerate the needed expenditure.
- More precisely, let $w + c$ be the amount of expenditure that C prefers (c measures the difference between the committee and the floor), and f be the amount of expenditure that F chooses in the end, then F's payoff function is

$$u_F = -(f - w)^2,$$

and C's payoff function is

$$u_C = -(f - (w + c))^2.$$

Payoffs in graph

- Again, $u_F = -(f - w)^2$, and $u_C = -(f - (w + c))^2$.

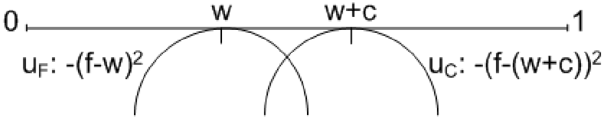

- Note that the payoffs depend solely on what the final expenditure is. There is no cost for C to send any signal.

Impossibility of perfect information transmission

- Suppose there is an equilibrium in which the committee accurately reports the true state of the world (w) to the floor.
- Given this strategy of C, F should believe the state of the world is whatever C reports, and choose $f = w$.
- But given this strategy of F, the committee should report that the state of the world is $w + c$, since that would lead the floor to select $f = w + c$, which maximizes the committee's payoff.
- Therefore in a cheap talk game there is no completely separating equilibrium, in which the sender perfectly reveal her information (i.e., the state of the world) to the receiver, as long as there is a difference between the two players' preferences.

Babbling equilibrium: no information transmission

- On the other hand, there is a completely pooling equilibrium in the game, in which the committee's message to the floor is constant, regardless of the true state of the world.
    - If the committee always reports the same thing, say r, the floor's optimal strategy is to ignore its report, and sets $f = 1/2$, which is the mean of the uniform distribution on $[0, 1]$.
    - If the floor always ignores the committee's message and sets $f = 1/2$, then any message is optimal for the committee, including r.
- Such a completely pooling equilibrium, called "**babbling equilibrium**", always exists in a cheap talk game.

Partial information transmission: $K = 2$

- Does the game have equilibria in which *some* information is transmitted?
- Suppose C can send one of two messages: r_1 if $0 \leq w < w_1$ and r_2 if $w_1 \leq w \leq 1$. What will be the threshold w_1?
- Given the strategy of C, consistency requires that F believes the true state of the world is uniformly distributed between 0 and w_1 if C reports r_1 and uniformly distributed between w_1 and 1 if C reports r_2.
- Therefore F optimally chooses $f = \frac{w_1}{2}$ if C reports r_1 and $f = \frac{w_1+1}{2}$ if C reports r_2.
    - Consistency does not restrict F's belief if the report is something other than r_1 and r_2. We can assume in that case F believes that the state is uniformly distributed either between 0 and w_1 or between w_1 and 1.

Partial information transmission: $K = 2$ (cont.)

- For this to be an equilibrium, $\frac{w_1}{2}$ must be at least as good as $\frac{w_1+1}{2}$ for C when $0 \leq w < w_1$, and $\frac{w_1+1}{2}$ must be at least as good as $\frac{w_1}{2}$ for C when $w_1 \leq w \leq 1$.
- In particular, when the true state is w_1, C should be indifferent between reporting r_1 and r_2, which means $w_1 + c$ is midway between $\frac{w_1}{2}$ and $\frac{w_1+2}{2}$. Therefore

$$w_1 = \frac{1}{2} - 2c.$$

- Because $w_1 > 0$, we must have $c < \frac{1}{4}$. In other words, if $c \geq \frac{1}{4}$, the game has no equilibrium in which the committee can send two different messages depending on the state.
- But if $c < \frac{1}{4}$, the committee can credibly transmit some partial information to the floor through cheap talk!

Partial information transmission: the general case

- Now suppose the committee can send one of K messages, $K \geq 2$. Can there be an equilibrium in which C reports r_1 if $0 \leq w < w_1$, r_2 if $w_1 \leq w < w_2$, ..., r_K if $w_{K-1} \leq w \leq 1$?
- The logic of the $K = 2$ case applies. When F receives the message r_k, it believes the state is uniformly distributed between w_{k-1} and w_k, and so it optimally chooses $f = \frac{w_{k-1}+w_k}{2}$.
- Given F's strategy, C is indifferent between r_{k-1} and r_k when $w = w_k$, which means

$$w_k + c = \frac{1}{2}\left(\frac{w_{k-1} + w_k}{2} + \frac{w_k + w_{k+1}}{2}\right),$$

which is equivalent to

$$w_{k+1} - w_k = w_k - w_{k-1} + 4c.$$

Partial information transmission: the general case (cont.)

- In other words, every interval is $4c$ longer than the previous interval.
- The length of the first interval is w_1, and all the intervals add up to 1. So

$$w_1 + (w_1 + 4c) + ... + (w_1 + 4(K-1)c) = 1,$$

which means

$$Kw_1 + 4c(1 + 2 + ... + (K-1)) = Kw_1 + 2cK(K-1) = 1.$$

- Therefore $w_1 = \frac{1-2cK(K-1)}{K}$, provided that c is sufficiently small that $2cK(K-1) < 1$.
- If $2cK(K-1) \geq 1$, there is no equilibrium in which C can send one of K messages.