Introduction to Game Theory

George Voutsadakis[1]

[1]Mathematics and Computer Science
Lake Superior State University

LSSU Math 500

Subsection 1

Strategies and Beliefs

Subgame Perfect Equilibrium Revisited

- Recall that a subgame perfect equilibrium of an extensive game with perfect information is a strategy profile for which every player's strategy is optimal (given the other players' strategies) at any history after which it is his turn to take an action, whether or not the history occurs if the players follow their strategies.
- The natural application of this idea to extensive games with imperfect information leads to the requirement that each player's strategy be optimal at each of his information sets.

Example

- Consider the game shown below.

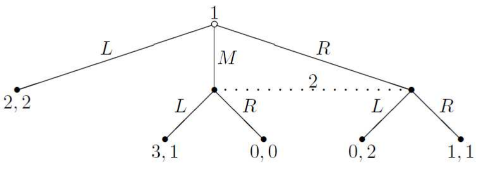

The requirement that each player's strategy be optimal at each of his information sets is substantial.

The pair of strategies (L, R) is a Nash equilibrium of this game.

If Player 1 adheres to this equilibrium, then Player 2's information set is not reached.

Example (Cont'd)

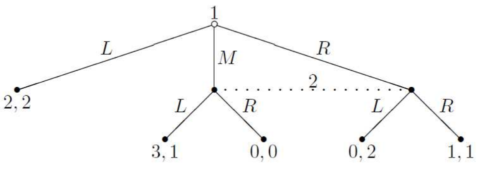

- The pair of strategies (L, R) is a Nash equilibrium of this game.
- If, for some reason, Player 2's information set is reached, then his action R is inferior to his action L whatever he thinks caused him to have to act (i.e., whether Player 1, despite her plan, chose M or R).
- For this game the natural extension of the idea of subgame perfect equilibrium is unproblematic.
  - The equilibrium (L, R) does not satisfy the conditions of this extension;
  - The equilibrium (M, L) satisfies those conditions.
- The games for which this is so are rare.

Another Example

- Consider the following game.

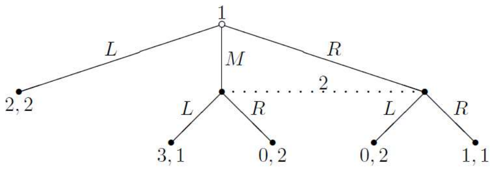

- The strategy profile (L, R) is a Nash equilibrium in which Player 2's information set is not reached.
- Now Player 2's optimal action in the event that his information set is reached depends on his belief about past history.
  - R is optimal if he assigns probability of at least $\frac{1}{2}$ to the history M.
  - L is optimal if he assigns probability of at most $\frac{1}{2}$ to history M.
- Thus his optimal action depends on his explanation of the cause of his having to act.
- His belief cannot be derived from the equilibrium strategy, since this strategy assigns probability zero to his information set being reached.

Sequential Equilibria: Including Belief Systems

- Our solutions for extensive games had strategy profile as its only component.
- We now study a solution, called sequential equilibrium, that consists of both a strategy profile and a belief system, where a belief system specifies, for each information set, the beliefs held by the players who have to move at that information set about the history that occurred.
- It is natural to include a belief system as part of the equilibrium, given the interpretation of the notion of subgame perfect equilibrium.
- To describe fully the players' reasoning about a game, we have to specify their expectations about the actions that will be taken after histories that will not occur if the players adhere to their plans, and these expectations should be consistent with rationality.
- Recall the interpretation of the components of a strategy that specify actions after histories that are not consistent with the strategy as beliefs about what will happen in these unexpected events.

Sequential Rationality

- The notion of sequential equilibrium should specify a pair, called an assessment consisting of:
  - The players' strategies;
  - The players' beliefs at each information set about the history that occurred.
- More formally, an **assessment** consists of:
  - (i) A profile of behavioral strategies;
  - (ii) A belief system consisting of a collection of probability measures, one for each information set.
- An assessment coincides with a strategy profile for an extensive game with perfect information, since, in such a game, all information sets are singletons and, hence, there is only one possible belief system.
- **Sequential rationality** stipulates that, for each information set of each player i, the strategy of Player i is a best response to the other players' strategies, given Player i's beliefs at that information set.

Consistency with Strategies

- One may impose certain restrictions on the players' beliefs.
- Here we consider three restrictions.
  - Consistency with strategies;
  - Structural consistency;
  - Commonality of beliefs.

1. Consistency with Strategies is the requirement that the belief system be consistent with the strategy profile.

   This means that, at any information set consistent with the players' strategies, the belief about the history that has occurred should be derived from the strategies using Bayes' rule.

Example

- Consider again the game

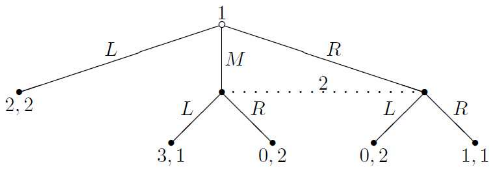

- We do not want (M, L) to be a solution supported by the belief of Player 2 that the history that led to his information set is R.
- Suppose Player 1's strategy is consistent with her choosing M or R.
- Then we want to require that Player 2's belief that the history M has occurred be derived from Player 1's strategy using Bayes' rule.
- Given Player 1's behavioral strategy β_1, Player 2 should assign probability $\frac{\beta_1(\emptyset)(M)}{\beta_1(\emptyset)(M)+\beta_1(\emptyset)(R)}$ to this event.

Structural Consistency

2. Structural Consistency is the requirement that, even at an information set that is not reached if all players adhere to their strategies, a player's belief be derived from some (alternative) strategy profile using Bayes' rule.

   This constraint on the beliefs is referred to as "structural" since it does not depend on the players' payoffs or on the equilibrium strategy.

Common Beliefs

3. Common Beliefs requires that all players share the same belief about the cause of any unexpected event.

   The rationale is that all asymmetries are included in the description of the game.

   So every player analyzes the situation in the same way.

   Therefore, in the context of subgame perfect equilibrium, all the players' beliefs about the plans of some player i, in case an unexpected event occurs, are the same.

Subsection 2

Sequential Equilibrium

Assessments

- We restrict attention throughout to games with perfect recall in which every information set contains a finite number of histories.
- A candidate for a sequential equilibrium of such a game is an assessment, defined formally as follows.

Definition (Assessment)

An **assessment** in an extensive game is a pair (β, μ), where:

- β is a profile of behavioral strategies;
- μ is a function that assigns to every information set a probability measure on the set of histories in the information set.

- Let (β, μ) be an assessment in $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i) \rangle$.
- The interpretation of μ, which we refer to as a **belief system**, is that $\mu(I)(h)$ is the probability that player $P(I)$ assigns to the history $h \in I$, conditional on I being reached.

Sequential Rationality

- An assessment is **sequentially rational** if, for every player i and every information set $I_i \in \mathcal{I}_i$, the strategy of Player i is a best response to the other players' strategies given i's beliefs at I_i.
- We define the **outcome** $O(\beta, \mu \mid I)$ **of** (β, μ) **conditional on** I to be the distribution over terminal histories determined by β and μ, conditional on I being reached.
- More precisely, let:
  - $h^* = (a^1, \ldots, a^K)$ be a terminal history;
  - $h = (a^1, \ldots, a^L) \in I$, for some $L < K$.

  Then:
  - $O(\beta, \mu \mid I)(h^*) = 0$, if there is no subhistory of h^* in I;
    I.e. the information that the game has reached I rules out h^*;
  - $O(\beta, \mu \mid I)(h^*) = \mu(I)(h) \cdot \prod_{k=L}^{K-1} \beta_{P(a^1,\ldots,a^k)}(a^1, \ldots, a^k)(a^{k+1})$.

Sequential Rationality (Cont'd)

- Perfect recall implies that there is at most one subhistory of h^* in I.
- We take the product in the second case, since:
  - By perfect recall, the histories $(a^1, \ldots, a^k)$, for $k = L, \ldots, K-1$, lie in different information sets;
  - Thus, for $k = L, \ldots, K-1$, the events

    $$\{a^{k+1} \text{ follows } (a^1, \ldots, a^k) \text{ conditional on } (a^1, \ldots, a^k) \text{ occurring}\}$$

    are independent.

A Troublesome Setup

- This definition of $O(\beta, \mu \mid I)$ has undesirable features in a game in which there are:
  - Two information sets I and I';
  - Histories $h \in I$ and $h' \in I'$,

  with the property that:
  - A subhistory of h is in I';
  - A subhistory of h' is in I.

Example Showcasing a Problem

- Consider the following game form.

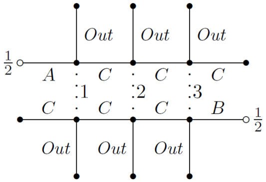

- We represent the initial history by several small circles.
- The number adjacent to each is the probability assigned by chance to one of its actions at the initial history.
- In an assessment (β, μ) in which $\beta_1 = \beta_3 =$ Out, Player 2's information set is not reached.
- If he is called upon to move, then an unexpected event must have occurred.

An Example Showcasing a Problem (Cont'd)

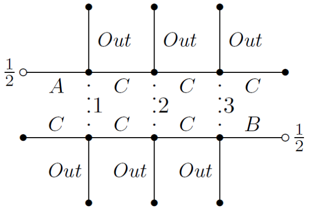

- Suppose that his belief at his information set, I, satisfies:
  - $\mu(I)(A, C) > 0$;
  - $\mu(I)(B, C) > 0$.
- In deciding the action to take in the event that I is reached, he must calculate $O(\beta, \mu \mid I)$.
- The definition of this distribution assumes that he continues to hold expectations about the moves of Players 1 and 3 derived from β.
- However, any strategy profile that generates the belief $\mu(I)$ must differ from β, since it must assign positive probability to both player 1 and player 3 choosing C.
- That is, if his belief is derived from an alternative strategy profile, then his explanation of the past is inconsistent with his expectation of the future.

Sequentially Rationality

- We formally define sequential rationality.

Definition (Sequentially Rational Assessment)

Let $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i)\rangle$ be an extensive game with perfect recall. The assessment (β, μ) is **sequentially rational** if, for every player $i \in N$ and every information set $I_i \in \mathcal{I}_i$, we have, for every strategy β_i' of Player i,

$$O(\beta, \mu \mid I_i) \succsim_i O((\beta_{-i}, \beta_i'), \mu \mid I_i).$$

Consistent Assessments

- Define a behavioral strategy profile to be **completely mixed** if it assigns positive probability to every action at every information set.

Definition (Consistent Assessment)

Let $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i) \rangle$ be a finite extensive game with perfect recall. An assessment (β, μ) is **consistent** if there is a sequence

$$((\beta^n, \mu^n))_{n=1}^{\infty}$$

of assessments that:

- Converges to (β, μ) in Euclidian space;
- Has the properties that:
  - Each strategy profile β^n is completely mixed;
  - Each belief system μ^n is derived from β^n using Bayes' rule.

Sequential Equilibrium

Definition (Sequential Equilibrium)

An assessment is a **sequential equilibrium** of a finite extensive game with perfect recall if it is sequentially rational and consistent.

- We shall see that every finite extensive game with perfect recall has a sequential equilibrium.
- It is clear that if (β, μ) is a sequential equilibrium, then β is a Nash equilibrium.
- Further, in an extensive game with perfect information, (β, μ) is a sequential equilibrium if and only if β is a subgame perfect equilibrium.

An Example

- Consider again the following game.

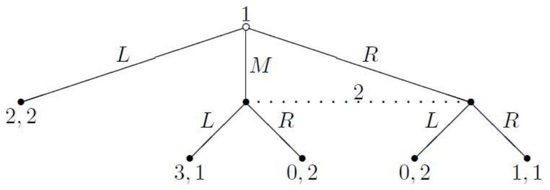

- Define the assessment (β, μ) by:
  - $\beta_1(L) = 1$ and $\beta_2(R) = 1$;
  - $\mu(\{M, R\})(M) = \alpha$, for any $\alpha \in [0, 1]$.
- (β, μ) is consistent.
  It is the limit as $\epsilon \to 0$ of assessments $(\beta^\epsilon, \mu^\epsilon)$, where:
  - $\beta_1^\epsilon = (1 - \epsilon, \alpha\epsilon, (1 - \alpha)\epsilon)$, $\beta_2^\epsilon = (\epsilon, 1 - \epsilon)$;
  - $\mu^\epsilon(\{M, R\})(M) = \alpha$, for every ϵ.
- For $\alpha \geq \frac{1}{2}$ this assessment is also sequentially rational.
- So, it is a sequential equilibrium.

Selten's Horse

- Consider the following game.

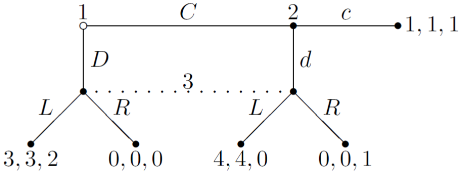

- Player 3's information set is $I = \{(D), (C, d)\}$.
- It has two types of Nash equilibria:
  - One in which $\beta_1(\emptyset)(D) = 1$, $\frac{1}{3} \leq \beta_2(C)(c) \leq 1$ and $\beta_3(I)(L) = 1$;
  - One in which $\beta_1(\emptyset)(C) = 1$, $\beta_2(C)(c) = 1$ and $\frac{3}{4} \leq \beta_3(I)(R) \leq 1$.
- A Nash equilibrium of the first type is not part of any sequential equilibrium, since the associated assessment violates sequential rationality at Player 2's (singleton) information set.

Selten's Horse(Cont'd)

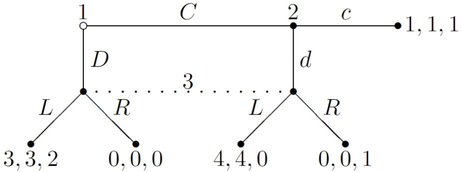

- Consider a Nash equilibrium β of the second type,

$$\beta_1(\emptyset)(C) = 1, \quad \beta_2(C)(c) = 1, \quad \frac{3}{4} \leq \beta_3(I)(R) \leq 1.$$

There is a sequential equilibrium (β, μ) in which $\mu(I)(D) = \frac{1}{3}$.
For consistency, take (β^ϵ), with:

$$\begin{array}{rcl} \beta_1^\epsilon(\emptyset)(C) & = & 1 - \epsilon; \\ \beta_2^\epsilon(C)(c) & = & \frac{1-2\epsilon}{1-\epsilon}; \\ \beta_3^\epsilon(I)(R) & = & \beta_3(I)(R) - \epsilon. \end{array}$$

Structurally Consistent Belief Systems

- The condition of structural consistency relates the beliefs and strategies in an assessment.

Definition (Structural Consistency)

Consider an extensive game with perfect recall.
The belief system μ is **structurally consistent** if, for each information set I, there is a strategy profile β, with the properties that:

- I is reached with positive probability under β;
- $\mu(I)$ is derived from β using Bayes' rule.

- Different strategy profiles may justify the beliefs at different information sets.

Consistency and Structural Consistency

- In many games, for any assessment (β, μ) that is consistent, the belief system μ is structurally consistent.
- However, in some games there are consistent assessments (β, μ) (even sequential equilibria) in which μ is not structurally consistent.

  I.e., they cannot be derived from any alternative strategy profile.

Consistent but Not Structurally Consistent Assessments

- Consider the game

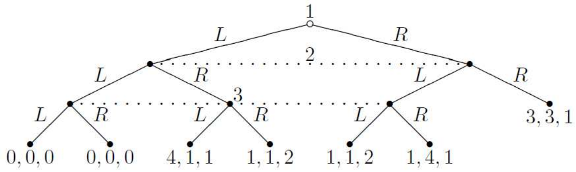

- It has a unique Nash equilibrium outcome, in which Players 1 and 2 choose R.

  Suppose, to the contrary, that Player 3's information set is reached with positive probability.

  Let the strategy profile used be β. Assume

$$\beta_i(I_i)(R) = \alpha_i, \quad i = 1, 2, 3,$$

  where I_i is the single information set of player i.

Nash Equilibrium (Case 1)

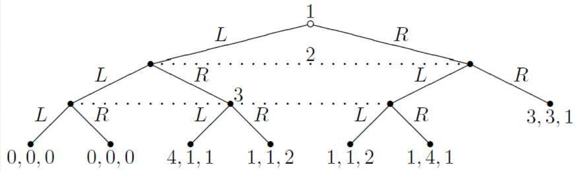

a. Suppose $\alpha_3 \leq \frac{1}{2}$. Then:
   - L yields Player 2 a payoff of $\alpha_1(1-\alpha_3)+\alpha_1\alpha_3 4=\alpha_1(1+3\alpha_3)$;
   - R yields Player 2 a payoff of $(1-\alpha_1)(1-\alpha_3)+(1-\alpha_1)\alpha_3+\alpha_1 3 = 1+2\alpha_1$.

   Note that
   $$\alpha_1(1+3\alpha_3) \leq \frac{5}{2}\alpha_1 < 1+2\alpha_1.$$
   Thus, Player 2 chooses R. But then, $\mu(I_3)((L,R))=1$.

   Hence, Player 3 chooses R with probability 1.

   This contradicts $\alpha_3 \leq \frac{1}{2}$.

Nash Equilibrium (Case 2)

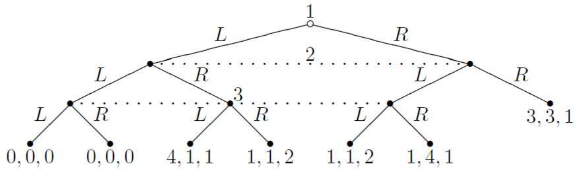

b. Suppose $\alpha_3 \geq \frac{1}{2}$. Then:
   - L yields Player 1 a payoff of $\alpha_2(4 - 3\alpha_3)$;
   - R yields Player 1 a payoff of $1 + 2\alpha_2$.

   Note that $\alpha_2(4 - 3\alpha_3) \leq \frac{5}{2}\alpha_2 < 1 + 2\alpha_2$.

   Thus, Player 1 chooses R.

   If Player 2 chooses L with positive probability, then $\mu(I_3)((R, L)) = 1$.

   So Player 3 chooses L with probability 1. This contradicts $\alpha_3 \geq \frac{1}{2}$.

   Thus, Player 2 chooses R with probability 1.

   This contradicts our assumption that Player 3's information set is reached with positive probability.

Nash Equilibrium

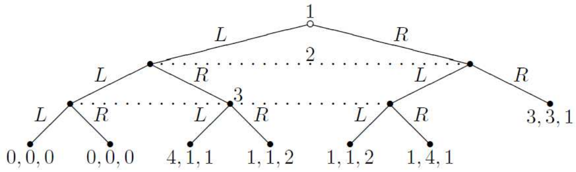

- We conclude that, in any Nash equilibrium, $\alpha_1 = \alpha_2 = 1$.
- In addition we need

$$\alpha_3 \in \left[\frac{1}{3}, \frac{2}{3}\right],$$

since, otherwise, either Player 1 or Player 2 can profitably deviate.

Let β be such an equilibrium.

We show that, for every assessment (β, μ) that is sequentially rational and consistent, the belief system μ is not structurally consistent.

Example (Cont'd)

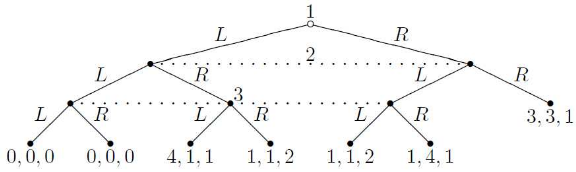

- For an assessment (β, μ) to be sequentially rational, Player 3's belief $\mu(I_3)$ must assign equal probabilities to the histories (L, R) and (R, L). Thus, it must take the form $(1 - 2\gamma, \gamma, \gamma)$.
  Such an assessment is consistent if and only if $\gamma = \frac{1}{2}$.
  For consistency in case $\gamma = \frac{1}{2}$, take (β^ϵ), with

$$\beta_1^\epsilon(I_1)(R) = \beta_2^\epsilon(I_2)(R) = 1 - \epsilon, \quad \beta_3^\epsilon(I_3)(R) = \alpha_3.$$

  However, the belief $(0, \frac{1}{2}, \frac{1}{2})$ of Player 3 violates structural consistency.
  Any strategy profile that yields (L, L) with probability zero also yields either (L, R) or (R, L) with probability zero.

Structurally Consistent but Not Consistent Beliefs

- Consider the following game.

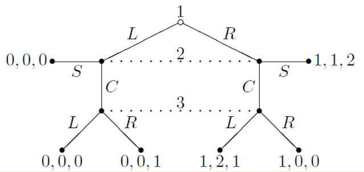

- We look at the assessment (β, μ), with:
  - β is the pure strategy profile (R, S, R);
  - Player 2's belief assigns probability 1 to the history R;
  - Player 3's belief assigns probability 1 to the history (L, C).

  (β, μ) is sequentially rational.

  The belief μ is structurally consistent.

  Player 3's belief is supported by the alternative pure strategy profile in which Player 1 chooses L and Player 2 chooses C.

  I.e., if Player 3 has to move, then she believes that Player 1, as well as Player 2, deviated from her equilibrium strategy.

Belief is not Consistent

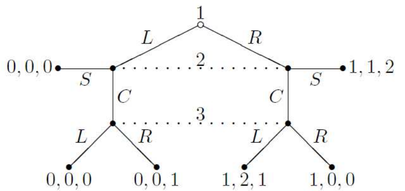

- On the other hand, (β, μ) is not consistent.

  Consider a sequence of assessments that involves strategies that:
  - Are completely mixed;
  - Converge to β.

  Such a sequence generates beliefs of Player 3 that converge to the belief that assigns probability 1 to the history (R, C).

  However, μ assigns probability 1 to (L, C)).

  Thus, (β, μ) is not a sequential equilibrium.

Subsection 3

Games with Observable Actions: Perfect Bayesian Equilibrium

Introducing Bayesian Extensive Games

- Bayesian extensive games constitute a family of games in which we can define a notion of equilibrium that is closely related to sequential equilibrium but is simpler.
- A Bayesian extensive game with observable actions models a situation in which every player observes the action of every other player.
- The only uncertainty is about an initial move of chance that distributes payoff-relevant personal information among the players in such a way that the information θ_i received by each player does not reveal any information about any of the other players.
- We say that chance selects types for the players.
- We refer to Player i after he receives the information θ_i as type θ_i.

Bayesian Extensive Games with Observable Actions

Definition (Bayesian Extensive Game with Observable Actions)

A **Bayesian extensive game with observable actions** is a tuple $\langle \Gamma, (\Theta_i), (p_i), (u_i) \rangle$, where:

- $\Gamma = \langle N, H, P \rangle$ is an extensive game form with perfect information and simultaneous moves;

and, for each player $i \in N$:

- Θ_i is a finite set (possible **types** of Player i); write $\Theta = \times_{i \in N} \Theta_i$;
- p_i is a probability measure on Θ_i for which $p_i(\theta_i) > 0$, for all $\theta_i \in \Theta_i$, and the measures p_i are stochastically independent;

  $p_i(\theta_i)$ is the probability that Player i is selected to be of type θ_i;
- $u_i : \Theta \times Z \to \mathbb{R}$ is a von Neumann-Morgenstern utility function;

  $u_i(\theta, h)$ is player i's payoff when the profile of types is θ and the terminal history of Γ is h.

Associated Extensive Game

- Suppose we start with a Bayesian extensive game with observable actions $\langle \Gamma, (\Theta_i), (p_i), (u_i) \rangle$.
- We can associate with it an extensive game (with imperfect information and simultaneous moves).
  - Its set of histories is $\{\emptyset\} \cup (\Theta \times H)$;
  - For each player i, each information set of Player i has the form

$$I(\theta_i, h) = \{((\theta_i, \theta'_{-i}), h) : \theta'_{-i} \in \Theta_{-i}\},$$

    for $i \in P(h)$ and $\theta_i \in \Theta_i$.
- Note that the number of histories in $I(\theta_i, h)$ is the number of members of Θ_{-i}.

Equilibrium Candidates

- Consider again a Bayesian extensive game with observable actions $\langle \Gamma, (\Theta_i), (p_i), (u_i) \rangle$.
- A candidate for an equilibrium is

$$((\sigma_i), (\mu_i)) = ((\sigma_i(\theta_i))_{i \in N, \theta_i \in \Theta_i}, (\mu_i(h))_{i \in N, h \in H-Z}),$$

  where:
  - Each $\sigma_i(\theta_i)$ is a behavioral strategy of player i in Γ (the strategy used by type θ_i of player i);
  - Each $\mu_i(h)$ is a probability measure on Θ_i (the common belief, after the history h, of all players other than i about player i's type).
- Such a pair is closely related to an assessment:
  - The profile (σ_i) rephrases the information in a profile of behavioral strategies in the associated extensive game;
  - The profile (μ_i) reflects the uncertainty about the other players' types.

Probability Measures on Terminal Histories

- Let s be a profile of behavioral strategies in Γ.
- Define $O_h(s)$ to be the probability measure on the set of terminal histories of Γ generated by s given that the history h has occurred.
- Define $O(\sigma_{-i}, s_i, \mu_{-i} \mid h)$ to be the probability measure on the set of terminal histories of Γ given that:
  - Player i uses the strategy s_i in Γ;
  - Each type θ_j of each Player j uses the strategy $\sigma_j(\theta_j)$;
  - The game has reached h;
  - The probability that i assigns to θ_{-i} is derived from $\mu_{-i}(h)$.
- That is, $O(\sigma_{-i}, s_i, \mu_{-i} \mid h)$ is the compound lottery in which the probability of the lottery $O_h((\sigma_j(\theta_j))_{j \in N - \{i\}}, s_i)$ is $\prod_{j \in N - \{i\}} \mu_j(h)(\theta_j)$, for each $\theta_{-i} \in \Theta_{-i}$.

Perfect Bayesian Equilibrium

Definition (Perfect Bayesian Equilibrium)

Let $\langle \Gamma, (\Theta_i), (p_i), (u_i) \rangle$ be a Bayesian extensive game with observable actions, where $\Gamma = \langle N, H, P \rangle$. A pair

$$((\sigma_i), (\mu_i)) = ((\sigma_i(\theta_i))_{i \in N, \theta_i \in \Theta_i}, (\mu_i(h))_{i \in N, h \in H - Z}),$$

where:

- $\sigma_i(\theta_i)$ is a behavioral strategy of Player i in Γ;
- $\mu_i(h)$ is a probability measure on Θ_i,

is a **perfect Bayesian equilibrium** if the following hold:

- Sequential Rationality For all $h \in H - Z$, $i \in P(h)$, $\theta_i \in \Theta_i$, the probability measure $O(\sigma_{-i}, \sigma_i(\theta_i), \mu_{-i} \mid h)$ is at least as good for type θ_i as $O(\sigma_{-i}, s_i, \mu_{-i} \mid h)$, for any strategy s_i of Player i in Γ.

Perfect Bayesian Equilibrium (Cont'd)

Definition (Perfect Bayesian Equilibrium)

- Correct Initial Beliefs $\mu_i(\emptyset) = p_i$, for each $i \in N$.
- Action-Determined Beliefs
  - If $i \notin P(h)$, $a \in A(h)$, then $\mu_i(h, a) = \mu_i(h)$;
  - If $i \in P(h)$, $a \in A(h)$, $a' \in A(h)$, $a_i = a'_i$, then $\mu_i(h, a) = \mu_i(h, a')$.
- Bayesian Updating If $i \in P(h)$ and a_i is in the support of $\sigma_i(\theta_i)(h)$, for some θ_i in the support of $\mu_i(h)$, then for any $\theta'_i \in \Theta_i$, we have

$$\mu_i(h, a)(\theta'_i) = \frac{\sigma_i(\theta'_i)(h)(a_i) \cdot \mu_i(h)(\theta'_i)}{\sum_{\theta_i \in \Theta_i} \sigma_i(\theta_i)(h)(a_i) \cdot \mu_i(h)(\theta_i)}.$$

Interpretations

- The first condition requires that the strategy $\sigma_i(\theta_i)$ of each type θ_i of each player i be optimal for type θ_i after every sequence of events.
- The second condition requires that initially the other players' beliefs about the type of each player i be given by p_i.
- The condition of action-determined beliefs requires that only a player's actions influence the other players' beliefs about his type:
  - (i) If Player i does not have to move at the history h, then the actions taken at h do not affect the other players' beliefs about Player i's type;
  - (ii) If Player i is one of the players who takes an action at h, then the other players' beliefs about Player i's type depend only on h and the action taken by Player i, not on the other players' actions.

Interpretations (Bayesian Updating)

- The condition of Bayesian updating relates to a case in which player i's action at the history h is consistent with the other players' beliefs about player i at h, given σ_i.
- In such a case the condition requires not only that the new belief depend only on Player i's action, but also that the players' beliefs be derived via Bayes' rule from their observation of Player i's actions.

Perfect Bayesian Equilibrium and Sequential Equilibrium

Proposition

Let $\langle\langle N, H, P\rangle, (\Theta_i), (p_i), (u_i)\rangle$ be a finite Bayesian extensive game with observable actions. Let (β, μ) be a sequential equilibrium of the extensive game associated with the Bayesian game. For every $h \in H$, $i \in P(h)$ and $\theta_i \in \Theta_i$, let

$$\sigma_i(\theta_i)(h) = \sigma_i(I(\theta_i, h)).$$

Then, there is a collection $(\mu_i(h))_{i\in N, h\in H}$, where $\mu_i(h)$ is a probability measure on Θ_i, such that:

- $\mu(I(\theta_i, h))(\theta, h) = \prod_{j\in N-\{i\}} \mu_j(h)(\theta_j)$, for all $\theta \in \Theta$ and $h \in H$;
- $((\sigma_i), (\mu_i))$ is a perfect Bayesian equilibrium of the Bayesian game.

- Perfect Bayesian equilibrium is easier to work with than sequential equilibrium, but applies to a significantly smaller set of situations.
- Even when they both apply, the two notions are not equivalent.

Example

- Consider a Bayesian extensive game with observable actions.

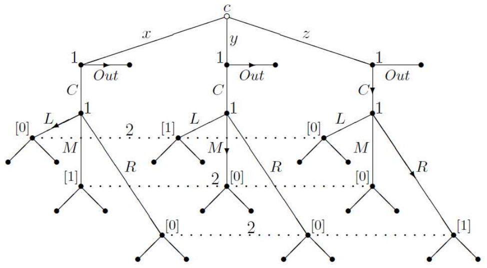

- Player 1 has three equally likely possible types, x, y and z.
- Player 2 has a single type.

Example (Cont'd)

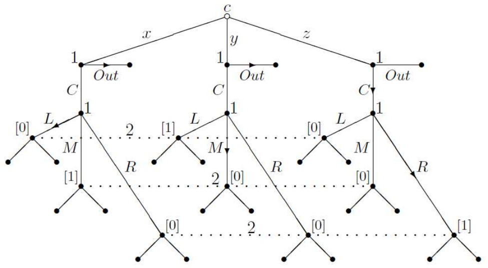

- Consider a perfect Bayesian equilibrium $((\sigma_i),(\mu_i))$ in which:
  - $\sigma_1(x) = (\text{Out}, L)$, $\sigma_1(y) = (\text{Out}, M)$, $\sigma_1(z) = (C, R)$;
  - $\mu_1(C, L)(y) = 1$, $\mu_1(C, M)(x) = 1$, and $\mu_1(C, R)(z) = 1$.
- Player 2 believes that Player 1 is:
  - Certainly of type y, if he observes the history (C, L);
  - Certainly of type x, if he observes (C, M);
  - Certainly of type z, if he observes (C, R) (the only history that is consistent with σ_1).

Example (Cont'd)

- Claim: $((\sigma_i), (\mu_i))$ may (depending on the payoffs) be a perfect Bayesian equilibrium of such a game, since it satisfies the conditions of action-determined beliefs and Bayesian updating.

  The probabilities of (C, L) and (C, M) are both zero, given σ_1.

  So $\mu_1(C, L)$ and $\mu_1(C, M)$ are not constrained by Bayesian updating.

  On the other hand,

$$\mu_1((C, R))(x), \quad \mu_1((C, R))(y), \quad \mu_1((C, R))(z)$$

  all obey Bayesian updating. E.g.,

$$\begin{aligned} \mu_1((C, R))(z) &= \frac{\sigma_1(z)(C)(R)\mu_1(C)(z)}{\sum_{w \in \{x,y,z\}} \sigma_1(w)(C)(R)\mu_1(C)(w)} \\ &= \frac{1}{0 + 0 + 1} = 1. \end{aligned}$$

Example (Cont'd)

- The associated assessment (β, μ) is not consistent, and hence is not a sequential equilibrium of any associated extensive game.

  To see this, let (β^n, μ^n) be a sequence of assessments that converges to (β, μ), with the properties that:
  - Each β^n assigns positive probability to each choice at every information set;
  - Each μ^n is derived from β^n using Bayes' rule.

  Denote by:
  - c_θ^n the probability, according to β_1^n, that Player 1 chooses C after the history θ.
  - ℓ_θ^n and m_θ^n the probabilities, according to β_1^n, that she chooses L and M respectively, after the history (θ, C).

  For $K = L, M, R$, let $I_2^K = \{(x, C, K), (y, C, K), (z, C, K)\}$ be the information set of Player 2 that is reached if Player 1 chooses C and then K.

Example (Cont'd)

- By Bayes' Rule, using the fact that the three types of Player 1 are equally likely, we have

$$\mu^n(I_2^L)(y,C,L) = \frac{c_y^n \ell_y^n}{c_x^n \ell_x^n + c_y^n \ell_y^n + c_z^n \ell_z^n}.$$

This converges (by assumption) to $\mu(I_2^L)(y,C,L) = 1$.

Dividing the numerator and denominator of $\mu_n(I_2^L)(y,C,L)$ by c_y^n:

$$\mu^n(I_2^L)(y,C,L) = \frac{\ell_y^n}{\frac{c_x^n}{c_y^n}\ell_x^n + \ell_y^n + \frac{c_z^n}{c_y^n}\ell_z^n}.$$

Since $\ell_y^n \to \beta_1(y,C)(L) = 0$ and $\ell_x^n \to \beta_1(x,C)(L) = 1$, we conclude $\frac{c_x^n}{c_y^n} \to 0$. Via a similar calculation for the belief at I_2^M, $\frac{c_y^n}{c_x^n} \to 0$.

This is contradictory. Thus, (β,μ) is not consistent.

Explanation

- The notion of sequential equilibrium requires that the beliefs of Player 2 at two information sets not reached in the equilibrium not be independent.
- They must be derived from the same sequence of perturbed strategies of Player 1.
- The notion of perfect Bayesian equilibrium imposes no such restriction on beliefs.

Another Example

- In all perfect Bayesian equilibria of the following game, a player who at some point assigns probability zero to some history later assigns positive probability to this history.

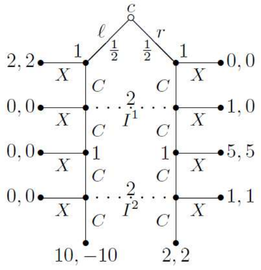

- Player 1 chooses C after r;
- Player 1 chooses X after (r, C, C);
- Player 2 chooses C at information I^1;
- Player 2 chooses X with probability at least $\frac{4}{5}$ at his information set I^2

Otherwise, Player 1 chooses C after ℓ and (ℓ, C, C), so that Player 2 assigns probability 1 to (ℓ, C, C, C) at information I^2, making C inferior to X.

Another Example (Cont'd)

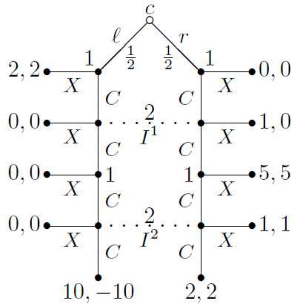

- Player 1 chooses X after ℓ.
- Thus, Player 2's belief at I^1 assigns probability 1 to r.
- On the other hand, his belief at I^2 assigns positive probability to chance having chosen ℓ. (Otherwise C is better than X.)

Introducing Signaling Games

- A **signaling game** is a Bayesian extensive game with observable actions that has the following simple form.
  - There are two players, a "sender" and a "receiver".
  - The sender is informed of the value of an uncertain parameter θ_1 and then chooses an action m (referred to as a **message**).
  - The receiver observes m (but not θ_1) and takes an action a.
  - Each player's payoff depends upon:
    - The value of θ_1;
    - The message m sent by the sender;
    - The action a taken by the receiver.

Signaling Games

- Formally, a **signaling game** is a Bayesian extensive game with observable actions $\langle \Gamma, (\Theta_i), (p_i), (u_i) \rangle$ in which:
  - Γ is a two-player game form in which:
    - First, Player 1 takes an action;
    - Then, Player 2 takes an action;
  - Θ_2 is a singleton.
- The tension in such a game arises from the fact that:
  - The receiver controls the action;
  - The sender controls the information.
- The receiver has an incentive to try to deduce the sender's type from the sender's message, and the sender may have an incentive to mislead the receiver.

A Signaling Game: Spence's Model of Education

- A worker (the sender) knows her talent θ_1 while her employer (the receiver) does not.
  - The value of the worker to the employer is the expectation of θ_1.
  - The employer pays a wage w that is equal to this expectation.

  To model this behavioral assumption, we assume that:
  - The payoff of the employer is $-(w-\theta_1)^2$ (the expectation of which is maximized when $w = E(\theta_1)$).
  - The worker's message is the amount e of education that she obtains.
  - Her payoff is $w - \frac{e}{\theta}$ (reflecting the assumption that the larger is θ the easier it is for a worker to acquire education).
- Assume that the worker's talent is either θ_1^L or $\theta_1^H > \theta_1^L$.
- Denote the probabilities of these values by p^L and p^H.
- Restrict attention to pure strategy equilibria.
- Denote the choices (messages) of the two types by e^L and e^H.
- This game has two types of perfect Bayesian equilibrium.

Pooling Equilibrium

- In one type of equilibrium both types choose the same level of education, $e^L = e^H = e^*$.

  Then the wage is $w^* = p^H\theta_1^H + p^L\theta_1^L$.

  We determine the possible values of e^*.

  Suppose a worker chooses a value of e different from e^*.

  Then in an equilibrium the employer must pay her a wage $w(e)$, such that, for $K = L, H$,

$$w(e) - \frac{e}{\theta_1^K} \leq w^* - \frac{e^*}{\theta_1^K}.$$

  The easiest way to satisfy this inequality is by making the employer believe that every deviation originates from a type θ_1^L worker, so that $w(e) = \theta_1^L$, for $e \neq e^*$.

  The most profitable deviation for the worker is to choose $e^L = 0$.

  So we need $\theta_1^L \leq w^* - \frac{e^*}{\theta_1^L}$. Equivalently, $e^* \leq \theta_1^L p^H(\theta_1^H - \theta_1^L)$.

Separating Equilibrium

- In another type of equilibrium the two types of worker choose different levels of education.

  In this case $e^L = 0$ (since the wage paid to a type θ_1^L worker is θ_1^L, independent of e^L).

  For it to be unprofitable for either type to mimic the other we need:
  - $\theta_1^L \geq \theta_1^H - \frac{e^H}{\theta_1^L}$;
  - $\theta_1^H - \frac{e^H}{\theta_1^H} \geq \theta_1^L$.

  These are equivalent to

  $$\theta_1^L(\theta_1^H - \theta_1^L) \leq e^H \leq \theta_1^H(\theta_1^H - \theta_1^L).$$

  Since $\theta_1^H > \theta_1^L$, a separating equilibrium thus always exists.
- The messages $e^L = 0$ and $e^H \in [\theta_1^L(\theta_1^H - \theta_1^L), \theta_1^H(\theta_1^H - \theta_1^L)]$ are supported as a part of an equilibrium in which any action other than e^H leads the employer to conclude that the worker's type is θ_1^L.