Perfect Bayesian Equilibrium and Sequential Equilibrium (July 20, July 24)

I. Backwards Induction and Subgame Perfect Equilibrium

- The backwards induction algorithm for games of perfect information – started at penultimate node
- For imperfect information games, this method was not possible because of the following cases:
  - The penultimate node may be in an information set that contained another node.
  - Even if the penultimate node were the only node contained in its information set, there would be a node in an information set that also contained a different node.
- In those cases, when reaching such a node in the induction method, backwards induction could not be applied. Those nodes were ignored in the generalized backwards induction. (This point will be explained using an example in Section II.)
- Today: introduce two equilibrium notions – **(weak) perfect Bayesian equilibrium** and **sequential equilibrium** that can be found using a backward induction method that can be applied to such nodes.
- If time allows: **trembling-hand perfect equilibrium**, a further refinement of sequential equilibrium

II. Motivating Example

- Consider the game below (**Game 1**).

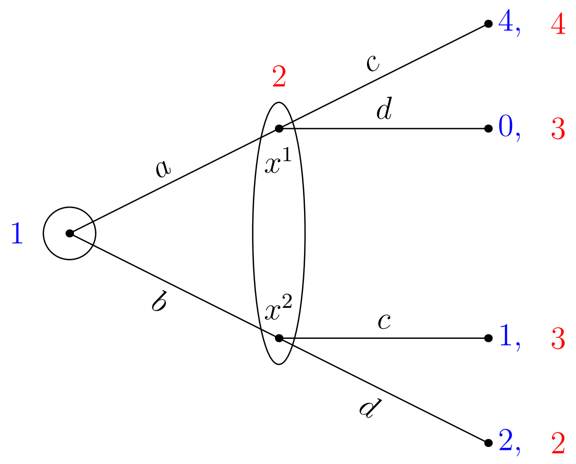

- The game above in matrix form is given by the following.

$1 \setminus 2$	c	d
a	$4, 4$	$0, 3$
b	$1, 3$	$2, 2$

- Two things to note:
  - Strategy d of player 2 is strictly dominated by strategy c. Therefore, when asked to move, player 2 will not choose d.
  - The unique Nash equilibrium is (a, c). Because there is no proper subgame of this game, (a, c) is also the unique subgame-perfect equilibrium.
- Now, modify the game so that player 1 has an "out" strategy (denoted by O) that immediately ends the game. The game tree and matrix are given below (**Game 1'**).

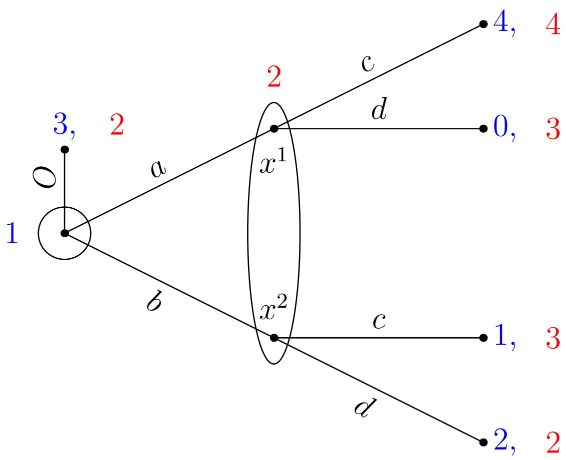

$1 \setminus 2$	c	d
a	$4, 4$	$0, 3$
b	$1, 3$	$2, 2$
O	$3, 2$	$3, 2$

- There are now two Nash equilibria, both of which are subgame-perfect: (a, c) and (O, d).
- However, player 2 when given the chance to move would not want to choose d since this yields a lower payoff than choosing c, **regardless** of whether player 1 had chosen a or b. Moreover, subgame-perfect equilibrium does not rule this out.

- Need a concept similar to backwards induction – (weak) perfect Bayesian equilibrium

III. System of Beliefs and Sequential Rationality

- In the previous example, start with player 2's information set. Player 2, when asked to move, does not whether he/she is at x^1 or x^2 and makes a prediction, or **belief**, as to which one by associating a probability for each event. That is, a player's belief (on an information set) is a probability distribution over the decision nodes in the information set.

- Formally, let H be an information set belonging to player i. A probability distribution over H is given by a function μ, where for each $x \in H$, $\mu(x)$ denotes the probability that player i believes he/she is at decision node x. Because these numbers represent probabilities, $\mu(x) \geq 0$ for all $x \in H$ and $\sum_{x \in H} \mu(x) = 1$.

- Now, consider a collection of such μ for each information set H. Such a collection is called a **system of beliefs**. That is, a function μ is said to be a system of beliefs if for each $H \in \mathcal{I}$,

$$\mu(x) \geq 0 \ \forall x \in H, \ \sum_{x \in H} \mu(x) = 1.$$

- Based on his/her own beliefs, each player chooses the action that maximizes expected payoffs $\to$ sequential rationality (to be formally defined below).

- In Game 1, let $\mu(x^1) = 1/3$ and $\mu(x^2) = 2/3$.
  - If player 2 chooses c, player 2's expected payoff is given by $(1/3)\times 4 + (2/3)\times 3 = 10/3$.
  - If player 2 chooses d, player 2's expected payoff is given by $(1/3)\times 3 + (2/3)\times 2 = 7/3$.
  - Since $10/3 > 7/3$, player 2 choosing c gives a higher expected payoff.

- Some notation before introducing the formal definition:
  - $Eu_i(\sigma_i, \sigma_{-i}|x)$: expected payoff of the behavioral strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ starting from node x. This value represents the payoff of i *as if* the game started at node x (regardless of whether previous play prescribed by σ led to x or not, just as in considering subgames for subgame-perfect equilibrium).

- $Eu_i(\sigma_i, \sigma_{-i}|\mu, H)$: Given beliefs μ, the expected payoff of σ starting from information set H – in the above concept, player i does not know for sure he/she is at node x but does know that he/she is at information set H which contains x. These values were also calculated in Game 1 above.
- Although the notation $Eu_i(\sigma_i, \sigma_{-i}|x)$ and $Eu_i(\sigma_i, \sigma_{-i}|\mu, H)$ may seem to indicate conditional expectation, but they have no such meaning.
- With abuse of notation, let $\Delta(S_i)$ denote the set of behavioral strategies.

Definition. Let μ be a system of beliefs. A behavioral strategy profile σ^* is **sequentially rational with respect to** μ if for each information set H,

$$Eu_i(\sigma_i^*, \sigma_{-i}^*|\mu, H) \geq Eu_i(\sigma_i, \sigma_{-i}^*|\mu, H)\ \forall \sigma_i \in \Delta(S_i)$$

where i is the player who moves at all decision nodes in H. Equivalently, the above inequality can be rewritten by the following:

$$\sum_{x \in H} \mu(x) Eu_i(\sigma_i^*, \sigma_{-i}^*|x) \geq \sum_{x \in H} \mu(x) Eu_i(\sigma_i, \sigma_{-i}^*|x)\ \forall \sigma_i \in \Delta(S_i)$$

- In Game 1, the strategy profile (a, c) is sequentially rational with respect to belief $\mu(x^1) = 1/3$ and $\mu(x^2) = 2/3$.

IV. Consistency of Beliefs and Perfect Bayesian Equilibrium

- Consider once again Game 1. (a, c) is sequentially rational with respect to belief $\mu(x^1) = 1/3$ and $\mu(x^2) = 2/3$, but given this strategy profile (a, c), is it reasonable for player 2 have a belief stated by μ? That is, should player 2 believe that player 1 would choose a with probability 1/3? $\rightarrow$ consistency condition for beliefs.
- In the previous section, beliefs μ were given with no justification to how they were set. Now, we consider the relationship between the strategies and the beliefs.
- The beliefs must be consistent to the actions that precede the information set. That is, the beliefs must match the probabilities induced by the actions specified by the strategies.
- Review of conditional probability.

Let A and B be events. Then, the **conditional probability of A given B** – denoted by $P(A|B)$ – is given by the following equation.

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

(It is assumed that $P(B) > 0$.)

- Given a strategy profile σ^* and an information set H and a decision node $x \in H$, the two events to look at:
  - A: the event that x is reached
  - B: the event that the information set H such that $x \in H$ is reached – that is, the event that at least one decision node in H is reached.
  - From how A and B are defined, note that $A \cap B = A$.
- Some notation: Let a strategy profile σ be given.
  - $P^\sigma(x)$: probability that node x is reached under the strategy profile σ.
  - $P^\sigma(H)$: probability that information set H is reached under the strategy profile σ.
  - Also, note that $P^\sigma(H) = \sum_{x \in H} P^\sigma(x)$.
  - Example of calculation of these probabilities in class.
- Therefore, the probability that a decision node $x \in H$ is reached from strategy profile σ, conditional on the event that H is reached, is given by

$$\frac{P^\sigma(x)}{P^\sigma(H)} = \frac{P^\sigma(x)}{\sum_{x' \in H} P^\sigma(x')}$$

A system of beliefs μ is said to be **consistent with the strategy profile** σ if the following equality is satisfied

$$\frac{P^\sigma(x)}{P^\sigma(H)} = \mu(x)$$

for all information sets H with $P^\sigma(H) > 0$. The above equation is called **Bayes' rule**. In words, μ and σ must satisfy Bayes' rule whenever applicable (that is, whenever $P^\sigma(H) > 0$).

- Putting the two ideas together, we have a formal definition of a perfect Bayesian equilibrium.

> Let σ^* be a strategy profile and μ a system of beliefs. (σ^*, μ) is a **perfect Bayesian equilibrium** if the following hold:
>
> - σ^* is sequentially rational with respect to μ
> - μ is consistent with σ^*

V. Some Observations

- In Game 1, $((a, c), \mu)$ is the unique perfect Bayesian equilibrium where $\mu(x^1) = 1$ and $\mu(x^2) = 0$.
- In Game 1', $((a, c), \mu)$ is the unique perfect Bayesian equilibrium where $\mu(x^1) = 1$ and $\mu(x^2) = 0$.
- Consider the following example (**Game 2**).

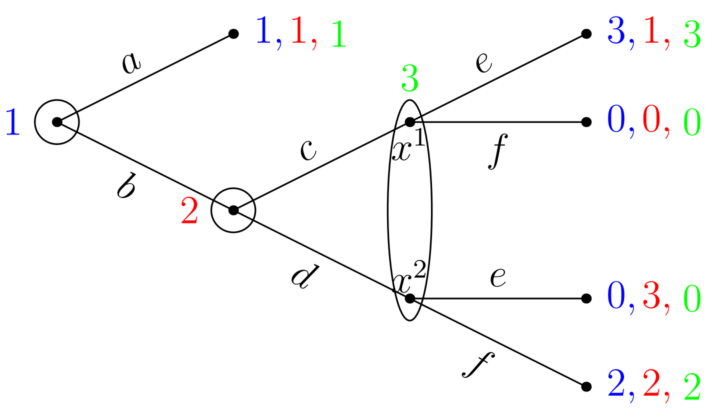

- The game in matrix form is given below, where player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices.

Matrix e	c	d
a	$1, \underline{1}, \underline{1}$	$\underline{1}, \underline{1}, \underline{1}$
b	$\underline{3}, 1, \underline{3}$	$0, \underline{3}, 0$

Matrix f	c	d
a	$\underline{1}, \underline{1}, \underline{1}$	$1, \underline{1}, \underline{1}$
b	$0, 0, 0$	$\underline{2}, \underline{2}, \underline{2}$

- There are three Nash equilibria: (a, d, e), (a, c, f), (b, d, f).

- Only (b, d, f) is a subgame-perfect equilibria. To see this, the subgame starting from player 2's decision node can be expressed in the following matrix form.

$2 \setminus 3$	e	f
c	$3, 1, \underline{3}$	$0, 0, 0$
d	$0, \underline{3}, 0$	$2, \underline{2}, \underline{2}$

- There is only one Nash equilibrium: (d, f). Therefore, (a, d, e) and (a, c, f) are not subgame-perfect equilibria.

- However, (a, d, e) combined with the belief $\mu(x^1) = 1$ is a perfect Bayesian equilibrium. The first equilibrium is shown below.

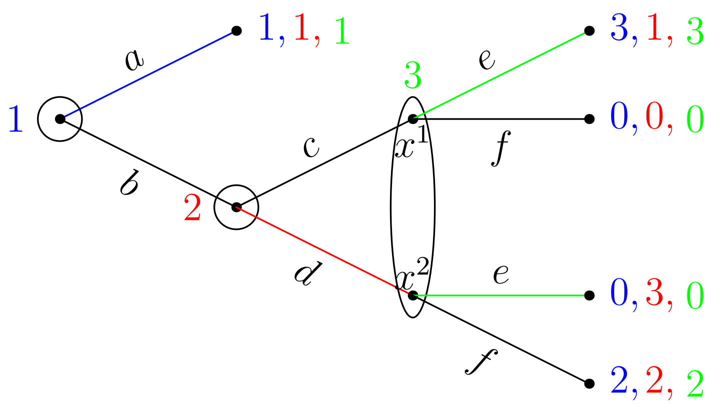

Explanation: Under belief μ, player 3's optimal action is to choose e. Then, when given player 3's action e, player 2's optimal action is d, which yields 3, over c, which yields 1. Given d and e, player 1's optimal action is to choose a. Thus, (a, d, e) is sequentially rational with respect to μ. μ is consistent with (a, d, e) since player 3's information set is not reached under this strategy profile, and Bayes' rule cannot be applied.

VI. Stronger Consistency Requirement and Sequential Equilibrium

- It was shown in the previous example that a weak perfect Bayesian equilibrium may not be a subgame-perfect equilibrium.

- In a perfect Bayesian equilibrium, beliefs can be arbitrary at information sets that are not reached. This results from the phrase "apply Bayes' rule whenever applicable."

- To avoid such unreasonable beliefs for information sets that are not reached in equilibrium, consider the additional requirement that the strategy profiles and beliefs are obtained from convergence of a sequence of completely mixed behavioral strategies and a sequence of system of beliefs that are consistent to these behavioral strategies.

- The convergence of a sequence $\approx$ robustness of beliefs to small changes in the strategy profiles. (Equilibria defined on principles of "robustness" are typically defined in this way.)

- Formal definition is given below.

> **Definition.** Let σ^* be a strategy profile and μ a system of beliefs. (σ^*, μ) is a **sequential equilibrium** if there exists a sequence of strategies $(\sigma^k)_{k=1}^{\infty}$ such that for all k, $\sigma_i^k(a) > 0$ for every action a available to player i (σ_i^k is said to be a **completely mixed behavioral strategy**) and a sequence of systems of beliefs $(\mu^k)_{k=1}^{\infty}$ such that
>
> - $\sigma^k \to \sigma^*$ and $\mu^k \to \mu$,
> - μ^k is consistent with σ^k for each k,
> - σ^* is sequentially rational with respect to μ – that is, for each information set H,
>
> $$Eu_i(\sigma_i^*, \sigma_{-i}^*|\mu, H) \geq Eu_i(\sigma_i, \sigma_{-i}^*|\mu, H) \ \forall \sigma_i \in \Delta(S_i)$$
>
> where i is the player who moves at all decision nodes in H.

- It is known that the strategy profile in every sequential equilibrium is a subgame-perfect equilibrium.

- Thus, in Game 2, the strategy combination of (a, d, e) cannot be part of a sequential equilibrium.

VII. Trembling-hand Perfect Equilibrium

- There is a further refinement of sequential equilibrium called trembling-hand perfect equilibrium (Selten (1975)).

Definition. Let σ^* be a strategy profile and μ a system of beliefs. (σ^*, μ) is a **trembling-hand perfect equilibrium** if there exists a sequence of completely mixed behavioral strategies $(\sigma^k)_{k=1}^{\infty}$ and a sequence of systems of beliefs $(\mu^k)_{k=1}^{\infty}$ such that

- $\sigma^k \to \sigma^*$ and $\mu^k \to \mu$,
- μ^k is consistent with σ^k for each k,
- For each $k = 1, 2, \cdots$ the following property holds: for each information set H,

$$Eu_i(\sigma_i^*, \sigma_{-i}^k | \mu^k, H) \geq Eu_i(\sigma_i, \sigma_{-i}^k | \mu^k, H) \ \forall \sigma_i \in \Delta(S_i)$$

  where i is the player who moves at all decision nodes in H.

- Main difference: third condition has to hold for all $k = 1, 2, \cdots$ for trembling-hand perfect, while this condition need to be held only at limit for sequential.
- Example below (**Game 3**):

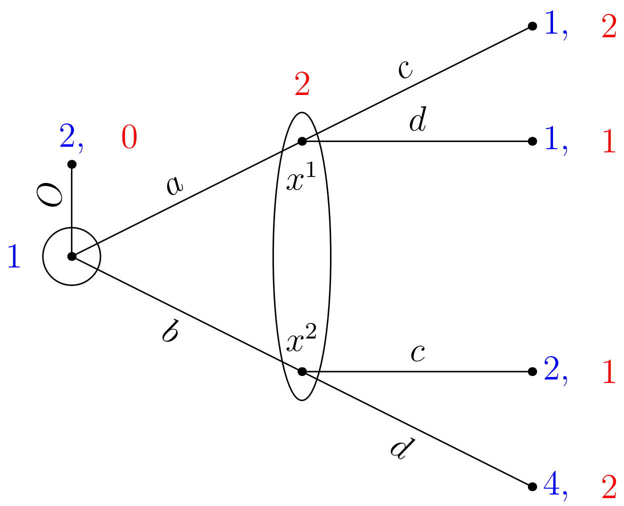

- $((O, c), \mu)$ wth $\mu(x^1) \geq 1/2$ is a sequential equilibrium but not a trembling-hand perfect equilibrium.
- Reason: completely mixed $\to$ positive probability that player 2 plays d ("small error"), but in that case, O is not a best response to player 2's completely mixed behavioral strategy and violates the third condition.

VIII. Further Topics and Notes on the Literature

- Mas-Colell, Whinston, and Green (1995) Section 9.C
- The original paper on sequential equilibrium: Kreps and Wilson (1982)
- Relationship between perfect Bayesian and sequential equilibrium and an equilibrium concept between them (Fudenberg and Tirole (1991))

References

Fudenberg, D. and J. Tirole (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory 53*, 236–260.

Kreps, D. and R. Wilson (1982). Sequential equilibria. *Econometrica 50*, 863–894.

Mas-Colell, A., M. Whinston, and J. Green (1995). *Microeconomic Theory*. New York: Oxford University Press.

Selten, R. (1975). Reexamination for perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory 4*, 25–55.