SF2972: Game theory

Mark Voorneveld, mark.voorneveld@hhs.se

February 9, 2015

Equilibria in extensive form games II:

1. Assessments: sequential rationality
2. Sequential equilibrium
3. Perfect Bayesian equilibrium
4. Examples

Reminder (Bayesian) consistency

We defined an *assessment* as a pair (b, β), where

- $b = (b_i)_{i \in N}$ is a profile of behavioral strategies and
- β is a belief system, assigning to each information set h a probability distribution β_h over its nodes.

We considered two possible restrictions on beliefs; informally:

1. Bayesian consistency: in information sets that are reached with positive probability, beliefs are determined by Bayes' law. In information sets reached with zero probability, beliefs are allowed to be arbitrary.
2. Consistency: beliefs are determined as a limit of cases where everything happens with positive probability and — consequently — where Bayes' law can be used.

Our next goal is to assure that players choose optimally given their beliefs and given the strategies of the other players.

Expected payoffs in information sets

Fix assessment (b, β) and an information set h of player i. To formalize the requirement that i plays a best response in info set h, we need to specify i's expected payoff:

1. Conditional on i being in his info set h, belief system β assigns probability $\beta_h(x)$ to being in node $x \in h$.
2. Given such a node x, the probability $\mathbb{P}(e \mid b, x)$ that an end node e is reached, conditional on starting in x and using strategies b is
   - zero if e cannot be reached from x;
   - the product of the probabilities of the corresponding branches from x to e otherwise.
3. In end node e, the payoff to i equals $u_i(e)$.
4. So the expected payoff to agent i in his information set h, given assessment (b, β) is

$$u_i(b_i, b_{-i} \mid h, \beta) = \sum_{x \in h} \beta_h(x) \left(\sum_e \mathbb{P}(e \mid b, x) u_i(e) \right).$$

Sequential rationality

Assessment (b, β) is *sequentially rational* if each player i in each of his information sets h chooses a best response to the belief system β and the strategies of the other players:

$$u_i(b_i, b_{-i} \mid h, \beta) \geq u_i(b_i', b_{-i} \mid h, \beta)$$

for all other behavioral strategies b_i' of player i.

Note:

1. consistency says that beliefs have to make sense given the strategies, without requirements on the strategies;
2. sequential rationality says that strategies have to make sense given the beliefs, without requirements on the beliefs.

Putting the two together, we have:

An assessment (b, β) is a *sequential equilibrium* if it is consistent and sequentially rational.

Theorem (Relations between solution concepts for extensive form games)

(a) *Each finite extensive form game with perfect recall has a sequential equilibrium.*

(b) *If assessment (b, β) is a sequential equilibrium, then b is a subgame perfect equilibrium and (hence) a Nash equilibrium.*

Proof sketch

(a) Via perfect equilibria of an auxiliary 'agent-strategic form game':
- Each player i is split up into agents, one agent for each of i's information sets;
- Agents of i have the same preferences as i;
- A mixed strategy in this agent-strategic form game is a behavioral strategy in the original game;
- Consider a completely mixed seq $b^m \to b$ making b a perfect equilibrium
- For each b^m, Bayes' law gives a belief system β^m.
- Drawing a convergent subsequence if necessary, we can show that $\lim_{m\to\infty}(b^m, \beta^m) = (b, \beta)$ is a sequential equilibrium.

(b) Suppose not. Let i have a profitable deviation b_i' in a subgame starting at some node x. Hence, in this subgame there has to be an information set that is reached with positive probability and where i has a profitable deviation, contradicting sequential rationality and correctness of beliefs.

Perfect Bayesian equilibrium (not on exam)

- Economists sometimes use perfect Bayesian equilibria, a notion that is more restrictive than subgame perfection, but less restrictive than sequential equilibrium.
- The intuition is that assessments are derived from strategies following Bayes' law whenever possible, but the exact definition of 'whenever possible' differs.
- In practice, a common requirement is that beliefs have to be Bayesian consistent with strategies in the game itself, but also in its subgames.
- Therefore, I will not discuss this notion further: if you have a carefully written game theory paper, the authors will make their equilibrium notion precise.

Compute the sequential equilibria of the game below:

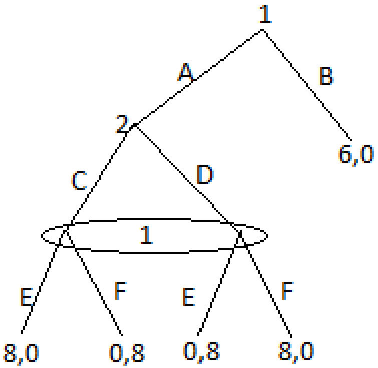

Intuition: What should it be? Player 1 chooses between sure payoffs $(6, 0)$ or the strategic game

	C	D
E	$8, 0$	$0, 8$
F	$0, 8$	$8, 0$

Behavioral strategies $b = (b_1, b_2)$ can be summarized by three probabilities:

1. p, the prob that 1 chooses A in the initial node;
2. q, the prob that 2 chooses C in his information set $\{A\}$;
3. r, the prob that 1 chooses E in information set $\{(A, C), (A, D)\}$.

Belief system β can be summarized by one probability α, the prob assigned to the left node (A, C) in the information set $\{(A, C), (A, D)\}$. Consistency: completely mixed beh. str. have $p, q, r \in (0, 1)$. Bayes' law then gives

$$\alpha = \frac{pq}{pq + p(1 - q)} = q,$$

So for each consistent assessment (b, β), it follows that $\alpha = q$. Which of these assessments also satisfies sequential rationality?

Distinguish 3 cases:

1. If $q = 0$, then $\alpha = 0$, so $r = 0$ is 1's unique best reply in the final info set. But if $r = 0$, then $q = 0$ is not a best reply in 2's info set. Contradiction.
2. If $q = 1$, then $\alpha = 1$, so $r = 1$ is 1's unique best reply in the final info set. But if $r = 1$, then $q = 1$ is not a best reply in 2's info set. Contradiction.
3. If $q \in (0, 1)$, rationality in 2's info set $\{A\}$ dictates that both C and D must be optimal. C gives $8(1 - r)$, D gives $8r$, so $r = 1/2$.
   In the info set $\{(A, C), (A, D)\}$ of pl. 1, his expected payoff is
   $$\alpha[8r] + (1 - \alpha)[8(1 - r)] \underbrace{=}_{\alpha = q} 8 - 8q + 8r(2q - 1).$$
   Choosing $r = 1/2$ is rational only if $q = 1/2$.
   Finally, in the initial node, A gives expected payoff 4 and B gives expected payoff 6, so $p = 0$.

Conclude: there is a unique sequential equilibrium with $p = 0, q = r = \alpha = 1/2$.

Homework exercise 3

Find the sequential equilibria (b, β) of the game in homework exercise 1.

Reading guide

1. Assessment: sequential rationality: slides 1–3, book 207
2. Sequential equilibrium: slides 4–5, 7–9, book 207–209
3. Perfect Bayesian equilibrium: slides 6, book §4.4, 207

- Peters, p. 207: "There is hardly any general method available to compute sequential equilibria: it depends very much on the game at hand what the best way is."
- During the remainder of the lecture, I will solve some more examples to illustrate different methods of finding sequential equilibria:
  1. Method 1: first find all consistent assessments, then find which of these are sequentially rational.
  2. Method 2: first find all sequentially rational assessments, then find which of these are consistent.
  3. Method 3: by the previous theorem: if (b, β) is a sequential eq, then b is a (subgame perfect) NE. So first find all (subgame perfect) NE. This is easier (no belief system) and often rules out many candidates. Then verify which can be turned into sequential equilibria.