SF2972: Game theory

Mark Voorneveld, mark.voorneveld@hhs.se

February 5, 2015

Equilibria in extensive form games I:

1. Nash equilibrium
2. Subgame perfect equilibrium
3. Relation between strategies and beliefs: assessments

- We can compute, for each profile of pure strategies, the corresponding (expected) payoffs: every extensive form game has a corresponding strategic/normal-form game.
- Terminology: Jörgen used 'normal-form game', the book of Peters uses 'strategic game'.
- A pure/mixed Nash equilibrium of the extensive form game is then simply a pure/mixed Nash equilibrium of the corresponding strategic game.
- Nash equilibria in behavioral strategies are defined likewise: a profile of behavioral strategies is a Nash equilibrium if no player can achieve a higher expected payoff by unilaterally deviating to a different behavioral strategy.

Example: from extensive to strategic game

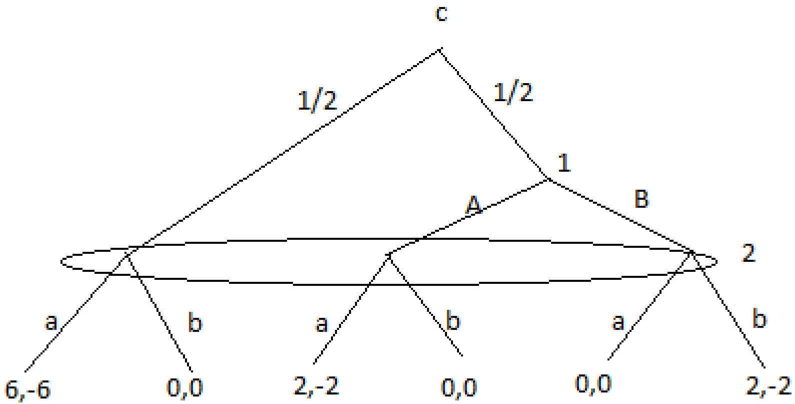

This game (from previous lecture) has strategic form:

	a	b
A	$4, -4$	$0, 0$
B	$3, -3$	$1, -1$

Dominance solvable, unique Nash equilibrium (B, b).

On the definition of strategies

I asked you to think about the following: pure, mixed, and behavioral strategies specify what happens in *all* information sets of a player. Even in those information sets that cannot possibly be reached if those strategies are used. Why do you think that is the case?
Main reason: Nash equilibrium: does each player choose a best reply to the others' strategies?
If a player were to deviate, ending up in a different part of the game tree, we need to know what happens there!

Theorem (Equilibrium existence)

Every finite extensive form game with perfect recall has a Nash equilibrium in mixed/behavioral strategies.

1. For mixed strategies: finite extensive form game gives finite strategic game, which has a Nash equilibrium in mixed strategies.
2. For behavioral strategies: by outcome-equivalence, we can construct a Nash equilibrium in behavioral strategies.

Strategic form analysis of extensive form games

The extensive form game

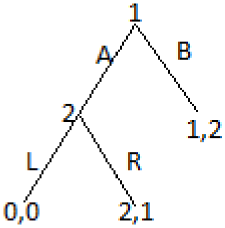

has corresponding strategic form

	L	R
A	$0,0$	$2,1$
B	$1,2$	$1,2$

Pure Nash equilibria: (B, L) and (A, R).

But if pl. 2 is called upon to play, would 2 choose L? This is an implausible choice in the 'subgame' that starts at node A!

To rule out such implausible equilibria, require that an equilibrium is played in each subgame: 'subgame perfect equilibrium'

Subgame perfect equilibrium

- In an extensive form game *with perfect information*, let x be a node of the tree that is not an end node. The part of the game tree consisting of all nodes that can be reached from x is called a *subgame*.
- Each game is a subgame of itself. A subgame on a strictly smaller set of nodes is called a *proper subgame*.
- A *subgame perfect equilibrium* is a strategy profile that induces a Nash equilibrium in each subgame.

In the game on the previous slide, only (A, R) is subgame perfect.

Subgame perfect equilibria via backward induction

Subgame perfect equilibria are typically found by backward induction:

1. Start with subgames with only one decision left. Determine the optimal actions there.
2. Next, look at subgames with at most two consecutive decisions left. Conditioning on the previous step, the first player to choose (say i) knows what a 'rational' player will do in the subgame that starts after i's choice, so it is easy to find i's optimal action.
3. Continue with subgames of at most 3 consecutive moves, etc.

This is the game-theoretic generalization of the dynamic programming algorithm in optimization theory.

Backward induction: example 1

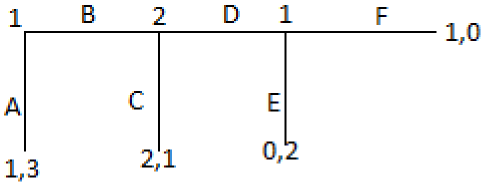

Strategic form:

	C	D
AE	$1,3$	$1,3$
AF	$1,3$	$1,3$
BE	$2,1$	$0,2$
BF	$2,1$	$1,0$

Pure Nash equilibria: (AE, D), (AF, D), and (BF, C).
Subgame perfect equilibrium: (BF, C)

Dividing 2 indivisible objects. Pl. 1 proposes, pl. 2 accepts or rejects.

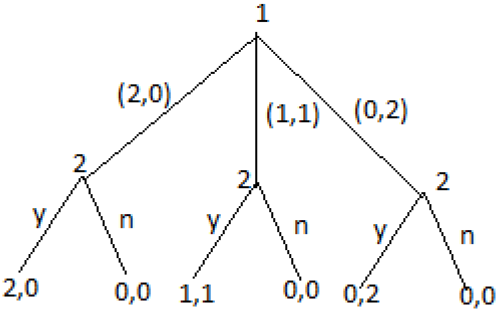

How many pure strategies for player 1? 3
How many pure strategies for player 2? $2^3 = 8$
Subgame perfect equilibria? $((2,0), yyy)$ and $((1,1), nyy)$

Backward induction: example 3 (the 'rotten kid' game)

- A child's action a from some nonempty, finite set A affect both her own payoff $c(a)$ and her parents' payoff $p(a)$; for all $a \in A$ we have $0 \leq c(a) < p(a)$.
- The child is selfish: she cares only about the amount of money she receives.
- Her loving parents care both about how much money they have and how much their child has. Specifically, model the parents as a single player whose utility is the smaller of the amount of money the parents have and the amount the child has. The parents may transfer money to the child (pocket money, trust fund, etc).
- First the child chooses action $a \in A$.
- Then the parents observe the action and decide how much money $x \in [0, p(a)]$ to transfer to the child. The game ends with utility $c(a) + x$ for the child and $\min\{c(a) + x, p(a) - x\}$ to the parents.

Show: in a subgame perfect equilibrium, the child takes an action that maximizes the sum of her private income $c(a)$ and her parents' income $p(a)$. Not so selfish after all!

- In the subgame after action $a \in A$, the parents maximize $\min\{c(a) + x, p(a) - x\}$ over $x \in [0, p(a)]$.
- This is done by choosing x such that $c(a) + x = p(a) - x$, i.e., by $x^*(a) = \frac{1}{2}(p(a) - c(a))$.
- Anticipating this, the child knows that action $a \in A$ leads to transfer $x^*(a)$ and consequently utility $c(a) + x^*(a) = \frac{1}{2}(c(a) + p(a))$. Maximizing this expression is equivalent with maximizing $c(a) + p(a)$.

Finite trees: existence of subgame perfect equilibria

Using backward induction, if there are only finitely many nodes, the first player to move — conditioning on the optimal behavior in the smaller subgames — is optimizing over a finite set: an optimum will always exist. Using this and induction on the 'depth' of the tree, one can show:

Theorem (Existence of subgame perfect equilibria)

In a finite extensive form game with perfect information, there is always a subgame perfect equilibrium in pure strategies.

That's a pretty nice result:

1. no need to consider randomization
2. no implausible behavior in subgames

As an aside: what if there are infinitely many nodes?

Subgame perfection and backward induction coincide

Theorem

In a finite extensive form game with perfect information, subgame perfect equilibria and those found by backward induction are identical.

Difficult! Main step is the 'one-deviation property': a strategy profile is subgame perfect if and only if for each subgame the first player to move cannot obtain a better outcome by changing only the *initial* action.

Centipede games

Although subgame perfect equilibria were introduced to rule out implausible behavior in subgames, there are examples where such equilibria lead to outcomes that some people find counterintuitive. This is sometimes corroborated with experimental support. One well-known example consists of Rosenthal's centipede games, characterized by the following properties:

- Players 1 and 2 take turns during at most $2T$ rounds ($T \in \mathbb{N}$).
- At each decision node, the player can choose to (S)top or (C)ontinue.
- The game ends *(i)* if one of the players decides to stop, or *(ii)* if no player has chosen stop after $2T$ periods.
- For each player, the outcome when he stops the game in period t is:
  - *better* than the outcome if the other player stops in period $t+1$ (or the game ends),
  - *worse* than any outcome that is reached if in period $t+1$ the other player continues.

Centipede games

Here is an example of a centipede game with 6 periods:

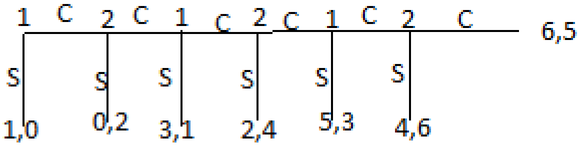

It is tempting to continue the game if you can be sure that the other player does so as well: the longer the game goes on, the higher the payoffs.

But in the unique subgame perfect equilibrium, players choose (S)top in each node. In particular, the game ends immediately in the initial node.

Reason: in the final node, player 2's best reply is to (S)top. Given that 2 (S)tops in the final round, 1's best reply is to stop one period earlier, etc.

There are other Nash equilibria, but they all lead to the same outcome: player 1 ends the game immediately.

Subgame perfect equilibrium in games with imperfect information

- Subgame perfect equilibria in games with perfect information require each player to play a best reply to other players' strategies in each subgame — regardless of whether that subgame is reached or not.
- It is possible to extend the notion of subgame perfect equilibria to games with imperfect information. But the definition of subgames is trickier: information sets must lie entirely outside the subgame or entirely inside the subgame.
- Formally, let x be a (non-end) node and let V^x be the nodes of the tree that can be reached from x. A well-defined subgame starts at x if and only if each information set h of the original game is a subset of V^x or is a subset of its complement.
- Since extensive form games with imperfect information need not have proper subgames, the notion of subgame perfection typically has little 'bite'.

Homework exercise 2

In the game of homework exercise 1:

(a) Find the corresponding strategic game.

(b) Find all pure-strategy Nash equilibria.

(c) What is the outcome of iterated elimination of weakly dominated (pure) strategies?

(d) Find all subgame perfect equilibria (in behavioral strategies).

Question: Can we find a suitable equilibrium refinement for imperfect information games that

1. makes sense even if there are no subgames and
2. still insists that players choose 'rationally' even in information sets that are reached with zero probability?

First attempt: require best responses in each information set.
Problem: the best response depends on where in the information set the player believes to be!

The following game has no proper subgames:

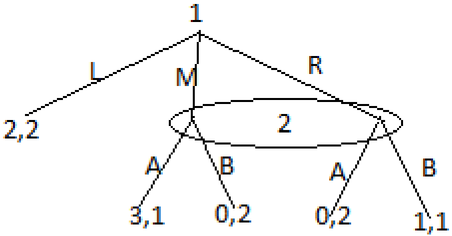

Beliefs affect optimal strategies: consider pl 2 in info set $\{M, R\}$. A is a best response if and only if the player assigns at most prob $1/2$ to being in node M.

Strategies affect reasonable beliefs: If pl 1 assigns to actions (L, M, R) probabilities $\left(\frac{1}{10}, \frac{3}{10}, \frac{6}{10}\right)$, pl 2 is twice as likely to end up in node R than in node M. Bayes' law requires that he assigns conditional prob $1/3$ to M and $2/3$ to R.

Question: What are reasonable beliefs if 1 chooses L with prob 1?

We consider two requirements on beliefs that give different answers to the final question:

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.

In particular, in both of these notions, we need to define two things: strategies and beliefs over the nodes in the information sets. The difference will lie in the constraints that are imposed.

Formally, consider a finite extensive form game with perfect recall. An *assessment* is 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.

Two belief requirements

Given node x and behavioral strategies b, let $\mathbb{P}_b(x)$ be the probability that node x is reached using b: it is the product of the probabilities assigned to the branches leading to x. Similarly, if h is an information set, it is reached with probability $\mathbb{P}_b(h) = \sum_{x \in h} \mathbb{P}_b(x)$.
Assessment (b, β) is:

- *Bayesian consistent* if beliefs in information sets reached with positive probability are determined by Bayes' law:

$$\beta_h(x) = \mathbb{P}_b(x)/\mathbb{P}_b(h)$$

  for every info set h with $\mathbb{P}_b(h) > 0$ and every node $x \in h$.
- *consistent* if there is a sequence of Bayesian consistent assessments $(b^m, \beta^m)_{m \in \mathbb{N}}$ with each b^m completely mixed (all actions in all info sets have positive prob) and $\lim_{m \to \infty}(b^m, \beta^m) = (b, \beta)$.

Note: (b, β) consistent $\Rightarrow$ (b, β) Bayesian consistent.

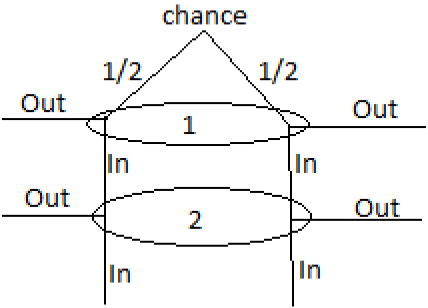

In the game above, where payoffs are omitted since they are irrelevant to the question:

(a) Find all Bayesian consistent assessments (b, β).

(b) Find all consistent assessments (b, β).

Summarize an assessment (b, β) by a 4-tuple $(p, q, \alpha_1, \alpha_2) \in [0,1]^4$, where

- p is the probability that 1 chooses *In*,
- q is the probability that 2 chooses *In*,
- α_1 is the probability that the belief system assigns to the left node in 1's info set,
- α_2 is the probability that the belief system assigns to the left node in 2's info set.

(a) Distinguish two cases:

1. If $p \in (0,1]$, 2's information set is reached with positive probability. In that case, Bayes' Law dictates that $\alpha_1 = \alpha_2 = \frac{1}{2}$. Conclude: all $(p, q, \alpha_1, \alpha_2) \in (0,1] \times [0,1] \times \left\{\frac{1}{2}\right\} \times \left\{\frac{1}{2}\right\}$ are Bayesian consistent.
2. If $p = 0$, 2's information set is reached with zero probability and 2 is allowed any belief $\alpha_2 \in [0,1]$ over the nodes in the information set. Bayes' Law only dictates that $\alpha_1 = \frac{1}{2}$. Conclude: all $(p, q, \alpha_1, \alpha_2) \in \{0\} \times [0,1] \times \left\{\frac{1}{2}\right\} \times [0,1]$ are Bayesian consistent.

(b) Every completely mixed profile of behavioral strategies leads to $\alpha_1 = \alpha_2 = \frac{1}{2}$.
Indeed, in 2's information set, both nodes are reached with equal probability $\frac{1}{2}p$.
Conclude: consistent are all
$(p, q, \alpha_1, \alpha_2) \in [0, 1] \times [0, 1] \times \left\{\frac{1}{2}\right\} \times \left\{\frac{1}{2}\right\}$.

Reading guide

1. Nash equilibrium: slides 1–4, book 47–48, 202–203
2. Subgame perfect equilibrium and backward induction: slides 5–16, book §4.3, 204
3. Assessments: consistency of beliefs: slides 18–24 , book §4.4 (partly), §14.3.2 (204–206)