Thinking Backwards

MARCO PAGNOZZI
marco.pagnozzi@st-catherines.oxford.ac.uk
(Based on Notes by David Myatt)

Objective:

- The representation of games in extensive form.
- Backward induction.
- Information sets.
- Imperfect Information.
- Subgames.
- Subgame perfect equilibrium.

- In previous lectures, we have considered:
  - Interactive decision-making problems and their representation in *games*.
  - *Strategic-form* or *normal-form* representation of a game.
  - *Strategies* as *complete plans of action*.
  - The deletion of *strictly dominated strategies* by "rational" players, and the notion of *rationalisability*.
  - The *Nash equilibrium* concept.

- We now attempt to refine the Nash equilibrium concept.
- We focus on games that evolve over time and consider:
  - Iterative deletion of *weakly dominated strategies*.
  - The representation of a game in *extensive form*.
  - *Backward induction*.
  - The solution concept of *subgame-perfect Nash equilibrium*.

Deletion of Weakly Dominated Strategies

- **Definition**: Strategy a_i'' is weakly dominated by strategy a_i' if

$$\pi_i(a_i', a_{-i}) \geq \pi_i(a_i'', a_{-i})$$

for all $a_{-i} \in A_{-i}$, with a strict inequality for at least one a_{-i}.

- Strategy B is weakly dominated by Strategy A if:
  - A yields a payoff not lower than that from B, irrespective of the choices made by others.
  - There is at least one strategy profile for which Strategy A yields a strictly higher payoff.

- Consider the following game:

	F	T
S	2 0	2 0
E	-1 -1	1 1

$\Rightarrow$

	T
S	2 0
E	1 1

$\Rightarrow$ The game has two pure-strategy Nash equilibria: {S, F} and {E, T}.

- But removing weakly dominated strategies removes one Nash equilibria, and leaves {E, T}.

- The remaining equilibrium is *trembling hand perfect* (i.e. it is robust to small trembles).
- With tied payoffs, this procedure does not necessarily produce a unique outcome — it can be sensitive to the order in which strategies are deleted.

Entry Deterrence Revisited

- The previous example is the entry deterrence game in strategic form. Representing it in extensive form:

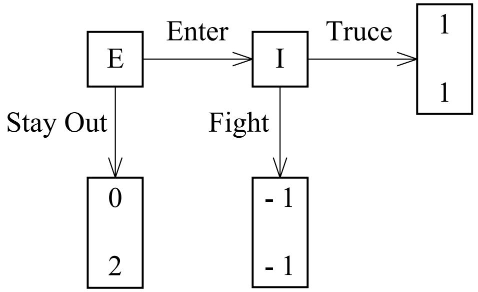

- Suppose there is entry. Then I prefers to accomodate — the threat to fight is not credible. Hence, we obtain a reduced tree:

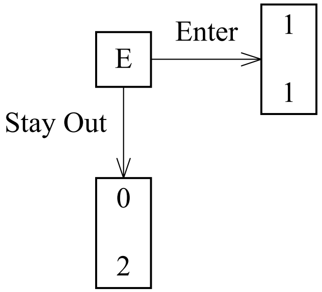

- But then the entrant will choose to enter.

⇒ By *backward induction* we get the equilibrium {E, T}.

- This procedure entailed the deletion of weakly dominated strategies.

Extensive Form Games

- A *finite game* in *extensive form* (with complete information) consists of:

1. A *directed graph* between a set of *nodes* (which forms a *tree*) such that:

(a) Each node maps to either:

(i) A *player*, possibly including *Nature*, or

(ii) A vector of *payoffs*, if it is a *terminal* node.

(b) Each *edge* (link between nodes) maps to an *action* for the player at the starting node.

2. A partition of nodes into *information sets* (i.e. collections of nodes) such that:

(a) The same player acts at each node in the set.

(b) The same actions are available at all nodes in the set.

3. The probability that Nature takes each available action, if she moves.

Rosenthal Centipede's Game

Two players move sequentially. In each stage they decide whether to pass or to terminate the game. Payoffs are:

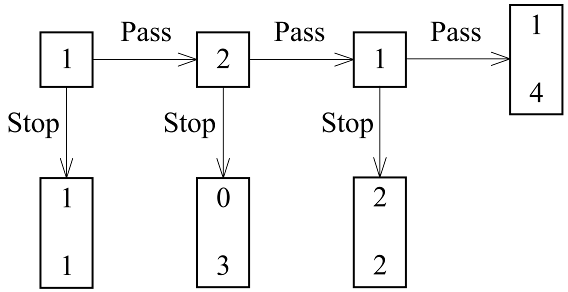

- Let's apply backward induction and start at the end of the game.
- For Player 1, Pass in the last stage is dominated by Stop. So delete it and obtain a reduced game:

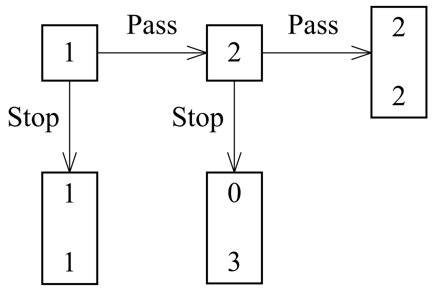

- For Player 2, Pass is dominated so delete that too.

- We are left with:

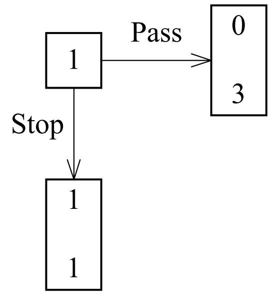

- For Player 1 Pass is dominated, so he chooses Stop.

⇒ The game ends immediately, with payoffs (1, 1).

- This outcome is Pareto-dominated by the outcome obtained by always passing.

- Consider the game in strategic form:

	Stop		Pass	
Stop-Stop	1	1	1	1
Stop-Pass	1	1	1	1
Pass-Stop	0	3	2	2
Pass-Pass	0	3	1	4

- For Player 1, Pass-Pass is weakly dominated. Delete this, and also delete Stop-Pass since it is strategically equivalent to Stop-Stop.

	Stop		Pass	
Stop-Stop	1	1	1	1
Pass-Stop	0	3	2	2

- For Player 2, Pass is now weakly dominated. Delete that as well. This leaves us with:

	Stop	
Stop-Stop	1	1
Pass-Stop	0	3

- Then Player 1 chooses Stop and the game will end immediately.

- Backward induction yields the same outcome in the larger Centipede Game:

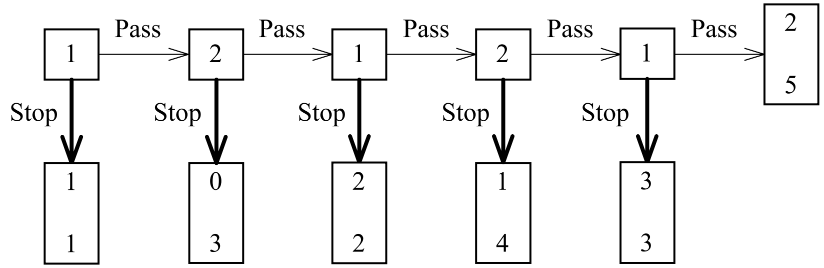

- Notice the features of this game:
  - Payoffs increase the longer the players wait.
  - If a player expects her opponent to stop, she wishes to stop before.

- Zoology provides an example:
  - Consider a group of apes who eat fruits.
  - The longer fruits are on the tree, the riper they get, and the better it is for an ape.
  - But the first ape to take a fruit, can get the best one.
  - Hence, if fruits are harvested, apes prefer to get their first.

Perfect and Imperfect Information

- A player cannot distinguish between nodes in an information set.
- If each information set contains a single decision node, then the game has *perfect* information.
- If there are non-singleton information sets, then the game has *imperfect* information.

- Extensive form games of incomplete information (in which a player is unsure of the opponent's payoffs) can be represented as games with complete but imperfect information.
- To do this, introduce Nature and assume she first chooses players' types.

Examples

- The Entry Deterrence and Rosenthal's Centipede games are extensive form games with complete and perfect information:

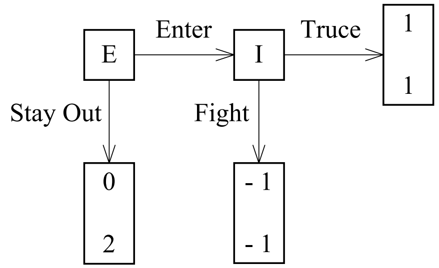

- The Prisoners' Dilemma has complete but imperfect information:

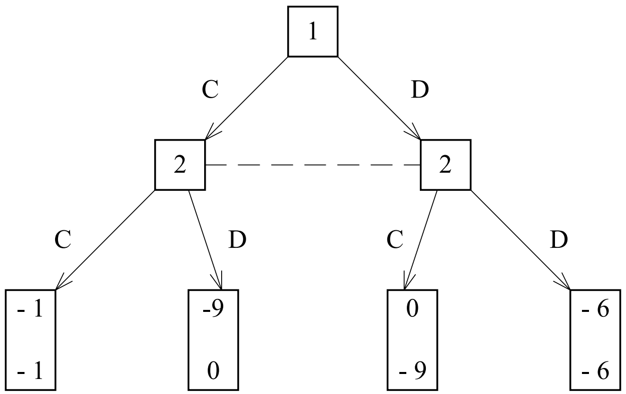

Subgames

- **Definition:** A *subgame* of an extensive form game is a subset of the game such that:
  (a) It begins with a singleton node and contains all of its successors;
  (b) If a node x is in the subgame, then all nodes in the information set containing x are in the subgame too.

- The Entry Deterrence game has two subgames:

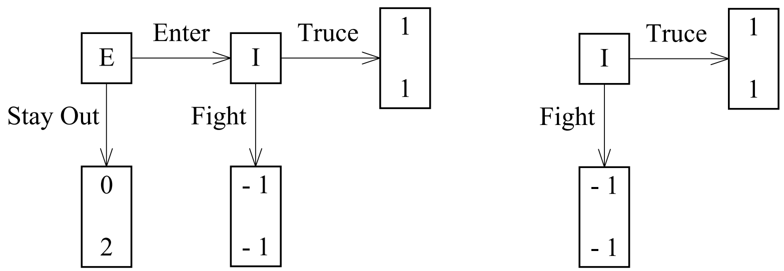

- The Prisoners' Dilemma has no (proper) subgame.

Subgame Perfection

- In an extensive game, a strategy is a *mapping from a player's information sets to actions*.
- **Definition:** A *strategy* for player i is a complete list of actions, one for each decision node in which player i has to choose an action.

- We want to rule out empty threats in dynamic games.
- A Nash equilibrium does not necessarily achieve this, e.g. the entry game.

- **Definition** (Selten, 1965)**:** A *subgame perfect Nash equilibrium* is a profile of strategies that induces a Nash equilibrium in *all* subgames.

⇒ A SPNE specifies optimal behaviour from any point of the game onward — no player has incentive to deviate in any subgame.

⇒ Every SPNE is a NE, but not every NE is a SPNE.

- In finite games of perfect information, the SPNE coincides with the NE that can be derived through *backward induction*.

Subgame Perfection with Perfect Information

- A game of *perfect* information has only *singleton information sets*.

⇒ Every node begins a new subgame and players always know what has happened.

- Every finite game of perfect information can be "solved" by backward induction to yield a subgame perfect equilibrium.
- With no ties in terminal nodes, backward induction yields a *unique* solution — e.g. the Centipede game.

- Each step of backward induction requires an additional level of knowledge of "rationality".
- As the length of the game increases, the assumptions of *common knowledge of rationality* and *complete and common knowledge of payoffs* become more important.

Stackelberg Leadership

Players: Two firms, a *leader* (firm 1) and a *follower* (firm 2).

Order of Play: Players proceed as follows:

(a) The leader chooses (and commits to) quantity q_1.

(b) The follower observes q_1 and then chooses quantity q_2.

Payoffs: Firm i's payoff is given by her profit function:

$$\pi_i(q_i, q_j) = [p(Q) - c_i]\, q_i; \quad i, j = 1, 2, \quad i \neq j.$$

- We assume linear demand: $p(Q) = 1 - (q_1 + q_2)$.
- By backward induction, consider the choice of firm 2 in the last stage after observing firm 1 choosing q_1.
- Firm 2 solves the concave problem:

$$\begin{aligned}\max_{q_2} \pi_2(q_2, q_1) &= [p(Q) - c_2]\, q_2. \\ &= [1 - q_2 - q_1 - c_2]\, q_2.\end{aligned}$$

- From FOC, we obtain 2's reaction function:

$$\frac{d\pi_2}{dq_2} = 0 \quad \Leftrightarrow \quad q_2 = R_2(q_1) = \frac{1 - q_1 - c_2}{2}.$$

- Now consider the first stage. Anticipating firm 2's reaction, firm 1 solves the concave problem:

$$\max_{q_1} \pi_1 (q_1, R_2 (q_1))$$

$$= [1 - q_1 - R_2 (q_1) - c_1] \, q_1$$
$$= \left[1 - q_1 - \frac{1 - q_1 - c_2}{2} - c_1\right] q_1$$
$$= \left[\frac{1 - q_1 + c_2 - 2c_1}{2}\right] q_1.$$

- From FOC:

$$\frac{d\pi_1}{dq_1} = \frac{1 - 2q_1 + c_2 - 2c_1}{2} = 0$$

$$\Leftrightarrow \quad q_1^* = \frac{1 + c_2 - 2c_1}{2}.$$

- This is higher than the Cournot quantity.

⇒ The first mover advantage helps firm 1.

- The *commitment* of firm 1 to producing q_1 is critical in this game. Why?

Nuisance Lawsuits

This example is designed to explore the US problem of nuisance lawsuits (Nalebuff, 1987):

Players: A *plaintiff* (P) and *defendant* (D).

Order of Play: Players proceed as follows:

(a) The plaintiff decides whether to pay a cost c and sue the defendant making a settlement offer $s > 0$.

(b) The defendant accepts or rejects the settlement offer.

(c) If the defendant rejects the offer, the plaintiff may go to trial or give up. Going to trial incurs legal costs of p to the plaintiff and d to the defendant.

(d) Going to trial, the plaintiff wins amount x.

Payoffs: We assume that $x < p$, so that the value of winning the lawsuit is relatively small.

- Consider the extensive form of the game:

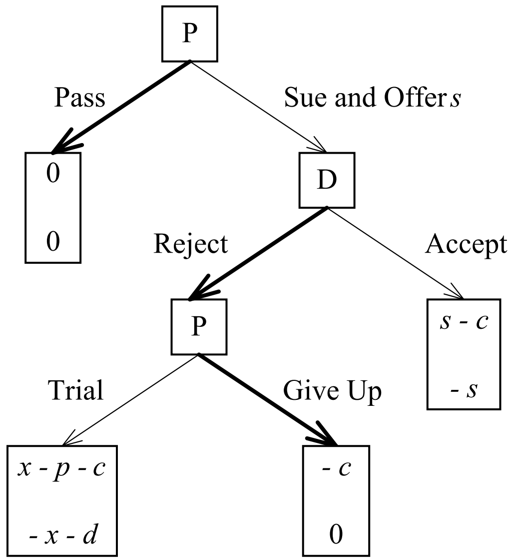

- Use backward induction:
  - Since $x < p$, in the last stage the plaintiff will never go trial,
  - so the defendant will reject his offer in the second stage,
  - so the plaintiff does not sue in the first stage.

⇒ There is no nuisance suit (and player's payoff is 0).

- P's threat of going to trial is not credible.

- Can the plaintiff overcome the credibility problem?
- Suppose the plaintiff pays his lawyer p in advance — a strategic sunk cost. The game becomes:

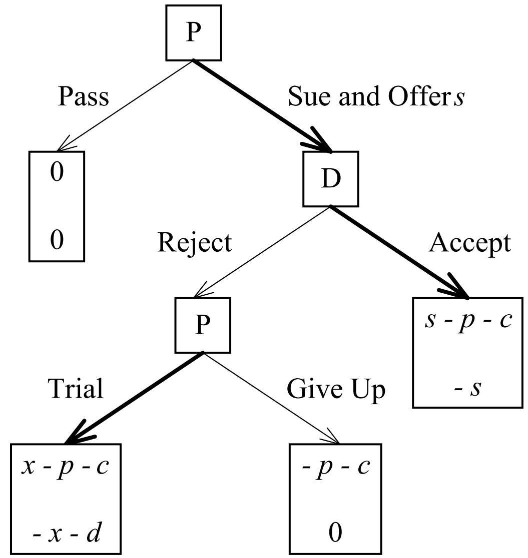

- Use backward induction:
  - Now in the last stage P will go to trial,
  - so in the second stage D will accept if $x + d \geq s$,
  - so in the first stage P sues if $s \geq p + c$

⇒ If $p + c \leq s \leq x + d$, then we have an subgame perfect equilibrium where a lawsuit with settlement occurs.

- Total payoff is $-p - c < 0$!