Contents

These notes introduce the basic ideas of game theory. After a very brief taxonomy and history, we introduce new ideas informally via a simple example of imperfect competition. A more formal presentation follows of games in extensive form (trees) and normal or strategic form (payoff functions or matrices).

The material is covered in every standard game theory text. Students should read these notes in parallel with such a text and with Chapter 7 of Mas-Colell et al.

[[note: for the next edition, consider moving material from the beginning of the next chapter to the end of this chapter: standard notation for bimatrices (and beyond), mixed strategy interpretation, example games.]]

1 Overview of game theory

- Games are situations where two or more decision makers (or "players") interact.
- The interaction of the decisions determines everyone's payoffs, thus there is strategic interdependence.

Brief history:

- 1944. von Neumann and Morgenstern launch the field with *The Theory of Games and Economic Behavior.*
- 1950's. Nash and Shapley
- 1960's-1970's. Aumann, Harsanyi, Selten
- 1980's-1990's. applications, refinements and evolutionary games
- 2000's. coalition theory, networks, bounded rationality, experiments (these started earlier, but exploded in the 1990's to present)

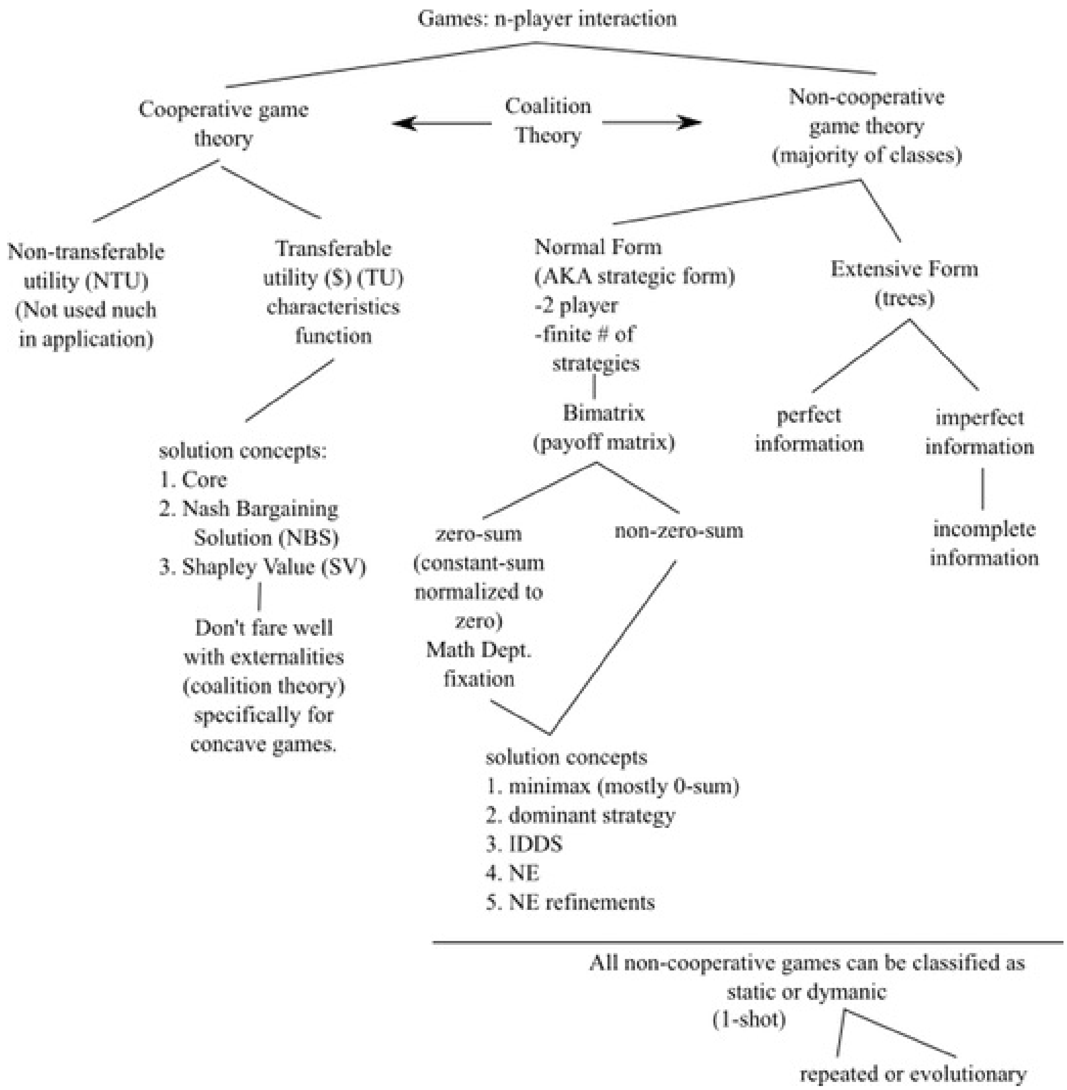

Figure 1: A "family tree" for games. Typos to fix: dymanic

1.1 Different sorts of game trees.

- *Incomplete information* means the game tree is not fully specified.
  - For example, if you don't know someone's utility function then you don't know her payoffs.
- *Complete information* means that everyone knows all there is to know about who chooses when and the consequences to all.

- In particular, all players' payoffs are common knowledge.
- This is true of our analysis of decision trees earlier (choice under uncertainty)

- *Perfect information* means that at each move the player knows the full history of the play so far.
  - They know all what moves each player made and the realizations of all random events, so far.
  - This means that they know the exact spot on the game tree which they are at.
  - Games of complete but imperfect information mean that choices are made under uncertainty.

- *Perfect recall* means that players have perfect memory of their previous moves, and all other moves that were previously part of their information set.
  - That is, you don't forget something once you know it.

1.2 Duopoly example

A non-cooperative extensive form game of imperfect information (borrowed from Kreps, 1990).

- Situation: Two firms, A and B, are both are considering introducing a new product.
  - Nature determines if the market for the new product is large or small.
  - Neither A nor B observes the other's choice, but B (who has hired a great market survey team) observes Nature's move.

- Figure 2 is A's decision "tree."

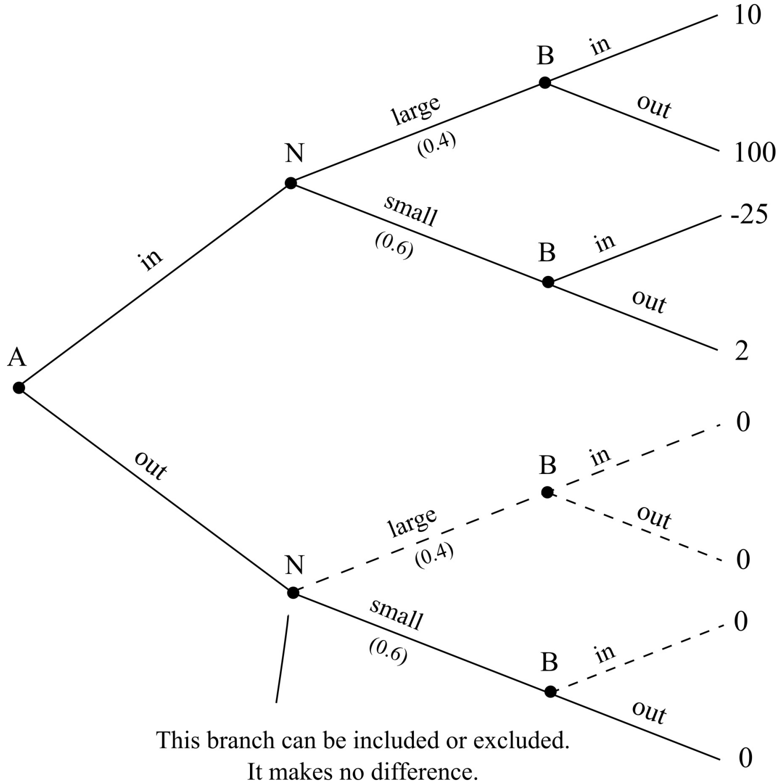

Figure 2: Decision from A's perspective.

  - Note that this not a properly specified tree. Just trying to make the bridge from decision theory.

- Ask: What should A do? (what is a^*?)
  - We don't know expected utility because we can't yet assign probabilities to B's move.
- Similarly, Figure 3 shows B's decision tree.
- Ask: What should B do?
  - Here the best choice is clear irrespective of what A chooses.
  - B has the "dominant" strategy $s = \{in|\text{Large}, out|\text{Small}\}$.

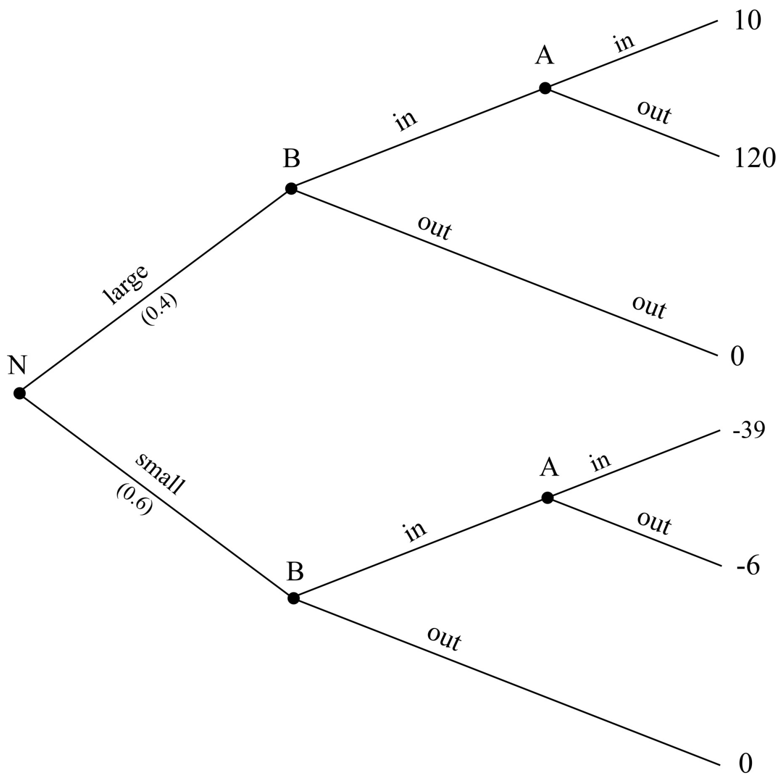

Figure 3: Decision from B's perspective.

- A can figure this out. So now A can solve her tree, using strategy s at B's nodes.
  - $\pi_{A,in}$|Large, B $in = 10$, $\pi_{A,in}$|Small, B $out = 2$ vs
  - $\pi_{A,out} = 0$.
  - $E\pi_{A,in} = .4(10) + .6(2) = 5.2 > 0$.
  - so A chooses 'In'.

1.3 Combined tree

To be more systematic, we need to put the two "decision trees" together. But we need something new to make it work.

An *information set*, denoted by bubbles or dots, is a device to indicate that a player has not observed some previous move.

- Thus, games of imperfect information contain at least one non-singleton information set.

- Now we can combine the improperly specified "decision trees" into one properly specified game tree.

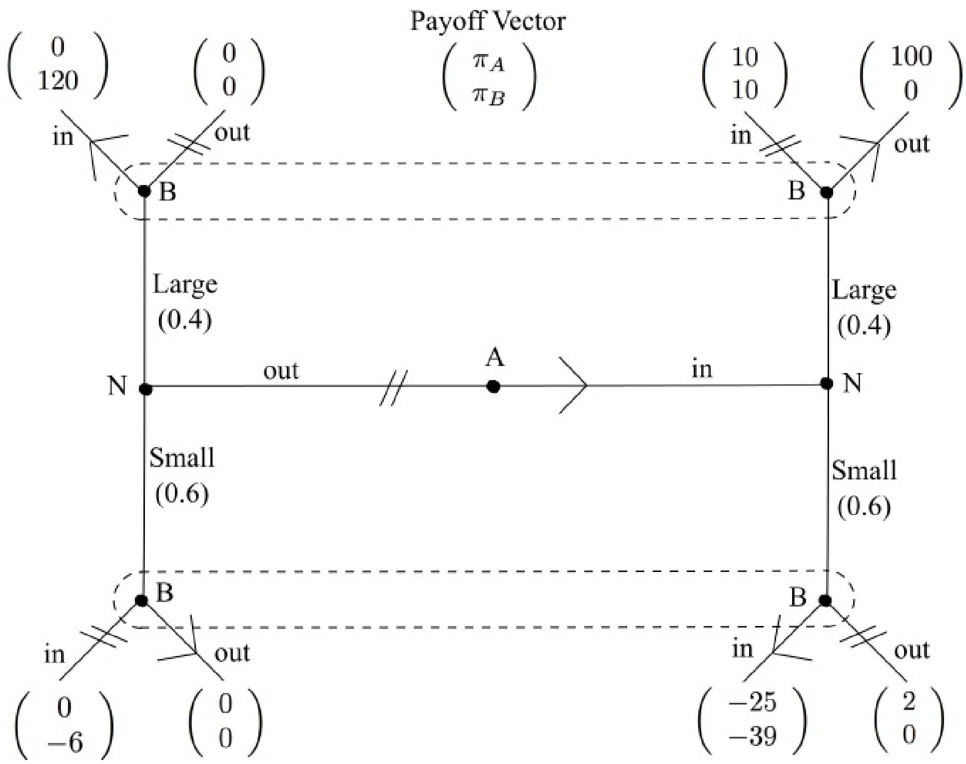

Figure 4: The tree for the duopoly game. Typo in figure on RHS should be in if large and out if small.

- The expected value of *in* for A depends on what B does. Solve this via BI.

- B's decision depends on the observation of Nature's move.

  – If B observes Large, then

    ∗ B's payoff for choosing *in* is: 10 if A is in or 120 if A is out.

    ∗ B's payoff for choosing *out* is: 0 if A is in or out.

        - Thus, if B observes Large, then *in* dominates *out*.
    - If B observes Small, then
        - B's payoff for choosing in is: -39 if A is *in* or -6 if A is *out*.
        - B's payoff for choosing out is: 0 if A is in or out.
        - Thus, if B observes Small, then *out* dominates *in*.
    - Therefore, we have a decision rule for B that is independent of A's action and depends only on the realization of Nature's move.
        - B's decision rule: $\{in|\text{Large}, out|\text{Small}\}$.
        - Note that the decision rule specifies an action for all possible realizations of Nature's move. An action for all possible histories in their information set.

1.4 Strategy

- Strategy $\equiv$ a complete contingency plan. (formal definition later).
- Dominant strategy $\equiv$ A best-response to all possible actions by the other player(s). (formal definition later).
    - Thus, B's decision rule $\{in|\text{Large}, out|\text{Small}\}$ is a dominant strategy.
- Since this is a game of complete information, and B has a dominant strategy, we can solve this game via BI.
    - Strategy sets:
        - A has two possible strategies: $\{in, out\}$.
        - B has 4 possible strategies: $\{(in|\text{Large}, in|\text{Small}), (in|\text{Large}, out|\text{Small}), (out|\text{Large}, in|\text{Small}), (out|\text{Large}, out|\text{Small})\}$.

- B has a dominant strategy and A knows this.
  - Therefore, by EUH A recognizes B chooses: (*in*|Large, *out*|Small), and calculates the expected utility of A's two strategies:

$$\begin{aligned} E(\pi_A|in) &= 0.4(10) + 0.6(2) = 5.2 \\ E(\pi_A|out) &= 0.4(0) + 0.6(0) = 0. \end{aligned}$$

  - Therefore, by EUH, A chooses *in* with expected payoff of 5.2.
    - We can go back to the tree and block off *out* for A.
- The expected payoff for B, given that A chooses *in*, is

$$E(\pi_B|(in|Large, out|Small)) = 0.4(10) + 0.6(0) = 4.$$

Notes:

1. This game can be solved by BI since B has a dominant strategy. This will not always be the case; BI will not always work in EFGs.

2. A game of perfect information is when the information sets are singletons, and thus trivial. Games of perfect information can be solved with BI as with decision trees by taking expectations across nature moves and choosing actions that maximize payoffs at decision nodes.

3. If B did not know Nature's move then the information set would go all the way around the tree and encapsulate all 4 decision nodes for B, as in Figure 5.

4. If B knew A's move, but not Nature's move then there would be two information sets, vertical on the outside of the tree, as in Figure 6.

2 Formalities

There are two main ways to describe a game formally. The first is more detailed.

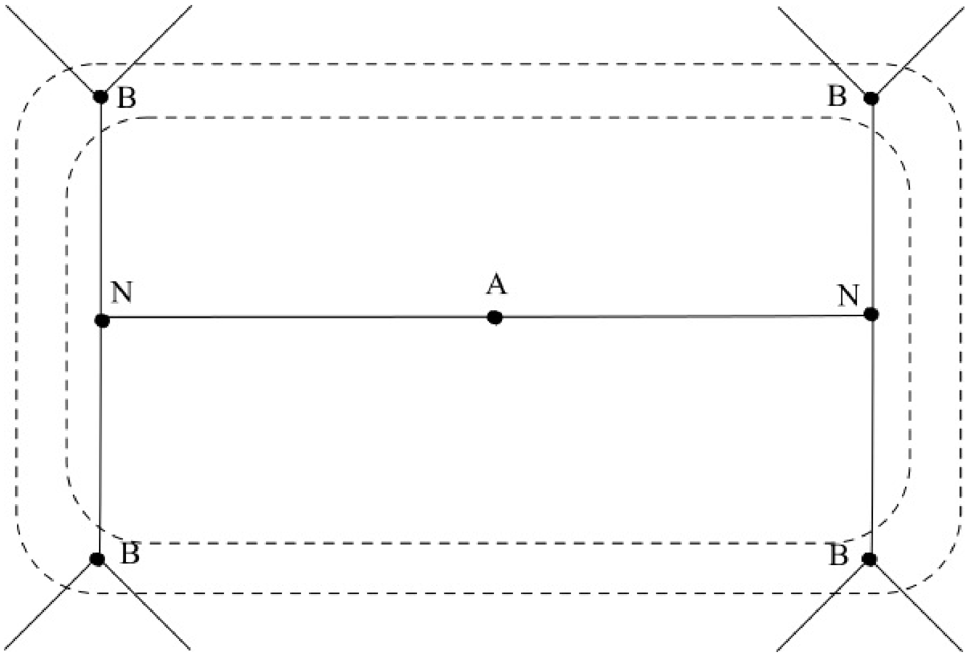

Figure 5: Only B nodes are included in the information set, The N or A nodes are not included.

2.1 Extensive form games (EFG)

In Mas-Colell notation (pg, 227), these are the elements of an EFG.

1. A set of regular players $\{I\}$, with elements $i = 1, \ldots, I$, where $2 \leq I < \infty$, and possibly also nature N as an additional player.

2. A tree, which is a finite set T of nodes with a precedence relation $\prec$,

   (a) $\prec$ is asymmetric, transitive and complete (total) ordering. That is, for all $x, y, z \in T$,

      i. Asymmetric: $x \prec y \Rightarrow \neg[y \prec x]$, read as: if x precedes y then it is not the case that y precedes x). Thus, $x \prec y$ and $y \prec x$ is not allowed.

      ii. Transitive: if $x \prec y$ and $y \prec z$ then $x \prec z$.

      iii. Total: $x, y \prec z \Rightarrow x = y$ or $x \prec y$ or $y \prec x$, i.e. branches don't come together.

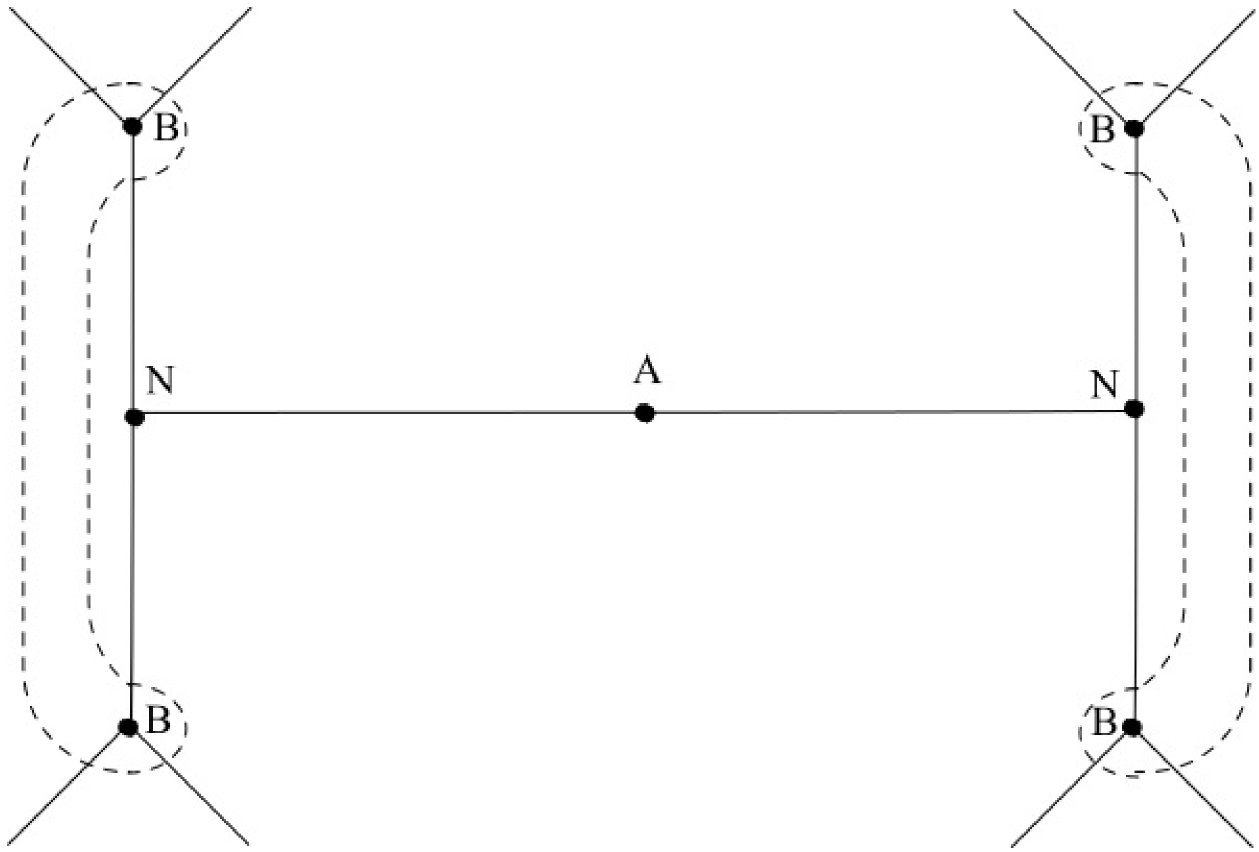

Figure 6:

(b) Therefore, cycles are not allowed.

i. This is one difference between game forms and networks.

ii. There is a unique way to get back to the root from each node.

(c) Three type of nodes.

i. A root, the initial node denoted with an open circle, which has no predecessors.

ii. Branches, regular nodes denoted by a large dot, which have some predecessors and some successors.

iii. Leaves, terminal nodes, which means that there are no successors.

3. An assignment of some specific player i (or nature, denoted by N) to each specific

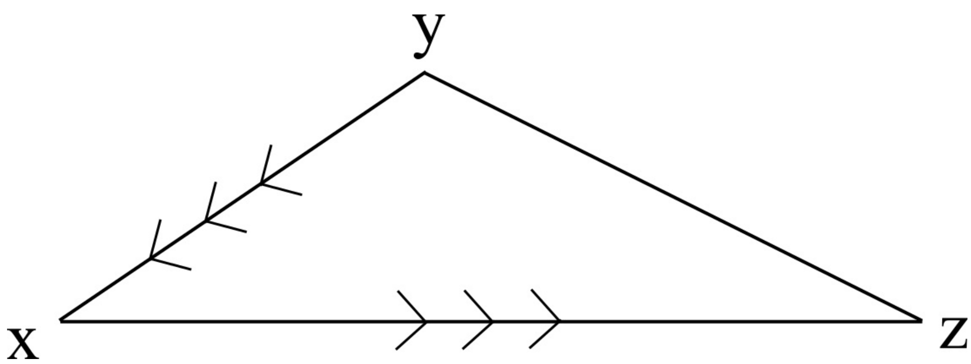

Figure 7: Not a tree because there is a cycle $x \prec z \prec y \prec x$, contrary to the ordering properties asymmetric, transitive and total.

node $t \in T$.

4. A finite set of actions $a_t \in A_t$ at each node t.

5. A partitioning of nodes T into information sets, such that every t is in a given information set. Each info set $t = \{t_1, ..., t_K\}$ must satisfy the following three properties.

   (a) Each node t_k in a given information set belongs to the same player i_ι with the same action set A_ι.

   (b) The nodes in ι do not precede each other.

   (c) Players have perfect recall, which means that they do not forget the past history of the game. (Total Recall, Arnie, 1990)

- (1) through (5) constitute a game form.

- We require two more elements to fully specify an extensive form game:

6. A payoff vector of length I at each terminal node. (All players know these payoffs, since complete EFG, i.e. complete information).

7. A probability vector over A_t at each nature node. Nature has a set of allowable "actions" at each nature node. The convention is that all players know the probabilities of each of these actions from the start.

- Extensive form games are usually presented in the form of a tree.

2.2 Normal form games (NFG)

A normal form game consists of the following elements.

1. A set of regular players $i = 1, \ldots, I$ where $2 \leq I \leq \infty$. Note: no Nature.

2. A strategy set S_i for each player $i \in I$..

3. A payoff function $f_i : S_1 \times ..... \times S_I \to R$ for each player $i \in I$

   (a) for each player, it maps every strategy vector ("profile") to a real number.

   (b) Commonly denoted: $f_i(s_1 \ldots, s_{i-1,} s_i, s_{i+1}, \ldots s_I) = f_i(s_i, s_{-i})$.

   (c) The payoff to player i is a function of the strategy profile: the strategy s_i played by player i and vector of strategies s_{-i} played by all the other players.

Given an EFG, a NFG can be specified by writing down all the possible strategies for each player and the expected payoff of each possible strategy profile.

Warning: a strategy in the NFG is a *complete* contingency plan for the player in the EFG. For example, if player B in the EFG owns 2 different nodes plus an info set consisting of 4 nodes, then

- there are 3 contingencies (2 trivial info sets plus the larger info set).

- A single strategy for player B then specifies the action to be taken at all 3 contingencies.

- If each action set allowed only two alternatives, then player B would have 2^3 distinct strategies.

- Students often just write down a fragment of a strategy (e.g., what the player does at the first or second node) and forget that, to write out even one strategy, they must specify an action at *every* contingency.

A **reduced NFG** comes from eliminating redundant rows and columns.

Solution procedure (preliminary steps):

1. Draw extensive form game and try BI. (Later we will see when BI does or does not help.)

2. Specify all strategies and write them down, then form NFG with expected payoffs.

3. Eliminate redundant rows and/or columns to obtain reduced NFG.

2.3 Duopoly example NFG

- From last time we had 2 strategies for A and 4 for B.

  - A's possible strategies: $\{in, out\}$.

    * Label $in = a_1$ and $out = a_2$.

  - B's possible strategies: $\{(in|\text{L}, in|\text{S}), (in|\text{L}, out|\text{S}), (out|\text{L}, in|\text{S}), (out|\text{L}, out|\text{S})\}$,

    * where L denotes large market realization and
    * S denotes small.
    * Label these b_1, b_2, b_3, b_4.

- The normal form is then a 2 by 4 matrix in strategies.

- To find the cell entries, we need to find the expected payoff vectors from the EFG. Cell entries are: (π_A, π_B).

		B			
		b_1	b_2	b_3	b_4
		(in|L, in|S)	(in|L, out|S)	(out|L, in|S)	(out|L, out|S)
A	$a_1 = in$	$-11, -19.4$	$\boxed{\underline{5.2},\underline{4}}$	$\underline{25}, -23.4$	$\underline{41.2}, 0$
	$a_2 = out$	$\underline{0}, 44.4$	$0,\underline{48}$	$0, -3.6$	$0, 0$

  - For example, cell (1,1) is from: $\{A$ in, B strategy 1$\}$

$$0.4\begin{pmatrix}10\\10\end{pmatrix} + 0.6\begin{pmatrix}-25\\-39\end{pmatrix} = \begin{pmatrix}4-15\\4-23.4\end{pmatrix} = \begin{pmatrix}-11\\-19.4\end{pmatrix}$$

  - Cell (2,1) is from: $\{A$ out, B strategy 1$\}$

$$0.4\begin{pmatrix}0\\120\end{pmatrix} + 0.6\begin{pmatrix}0\\-6\end{pmatrix} = \begin{pmatrix}0\\48-3.6\end{pmatrix} = \begin{pmatrix}0\\44.4\end{pmatrix}$$

  - Cell (1,2) is from: $\{A$ in, B strategy 2$\}$

$$0.4\begin{pmatrix}10\\10\end{pmatrix} + 0.6\begin{pmatrix}2\\0\end{pmatrix} = \begin{pmatrix}5.2\\4\end{pmatrix} \quad \text{NE expected payoffs from EFG}$$

  - A has no dominant strategy.

  - Note: b_2 is B's dominant strategy and a_1 is A's best response.

- We can use the standard method of underlining the best-response to one of the other player's strategies.

  - This yields the BI solution, A in, B choosing strategy 2 (in|L, out|S).

3 Another example

The main goal for the applied economists at this juncture is to become comfortable with

1. Given a story about strategic interaction, to write down an EFG that captures the essential features.

2. Given an EFG, write down the NFG.

3. The next few chapters will focus on solving NFGs and EFGs. But steps 1 and 2 come first!

To practice, consider the following story.

You are captain of a canoe. Entering a narrow channel from the opposite direction is an oil tanker ship. Both you (player C) and the oil tanker skipper (player T) have to decide whether to turn (L)eft or (R)ight. If one ship turns right and the other left, there will be a collision, in which case the canoe is destroyed, giving you a payoff of -5 and the oil tanker is undamaged, but the tanker skipper has to fill out an annoying accident investigation report form, giving her a payoff of -1. An (L, L) outcome is best for you since it positions you a little closer to the dock on the other end of the channel; the payoffs here are 2 for you and 1 for T. An (R, R) outcome helps T make an upcoming turn so here the payoff is 1 for C and 2 for T.

The oil tanker is bigger and heavier, therefore once the skipper makes a decision between L and R she cannot change that decision later. Once the oil tanker decides, you (player C) sees how the oil tanker moved, and decide which way to go.

Ok, how do we write out the extensive form of this game?

- Pretty simple. Player 1 (or T) moves first, choosing between L and R at the initial node.

- At each of the two successor nodes, Player 2 (or C) choses l or r.

- The four terminal nodes each has its own payoff vector (π_T, π_C).

Figure 8 illustrates.

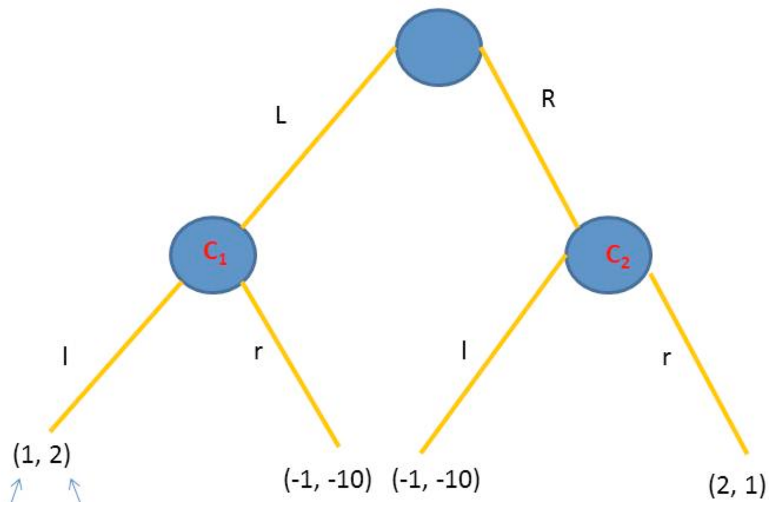

Figure 8: EFG for the first Canoe-Tanker story.

What is the NFG that comes from this EFG? Let's break that down. How many strategies does player 1 have? What are they? (Easy question)
How many strategies does player 2 have? What are they? (Be a little careful; use the notation l_L for example to denote "play l at player 2's L branch node".)

Hint: here is a *wrong* answer to the last question: $\{l_L, r_L, l_R, r_R\}$. Why? These are all strategy *fragments*, not complete contingency plans! Here a complete contingency plan can be written in the form $x_L y_R$, meaning "play x at the L branch node and play y at the R branch node."

There are 4 possible combinations of x and y, so Player 2 has 4 pure strategies. The payoff matrix has 2 rows and 4 columns, with each cell giving its own (π_T, π_C).

Write out this payoff matrix. Are there any identical rows or columns? If so, reduce the matrix by collapsing the redundancies.

Now consider a slightly different game. Now suppose that you (player 2, or C) are also operating a large boat. Therefore both you and player 1 (or T) have poor maneuverability and thus must SIMULTANEOUSLY choose between L and R. Assume that the payoffs for the various outcomes – (L,L),(L,R), (R,L), and (R,R) – are the same as before.

What is the EFG now?

Hint: player 2 now has an information set!

What are the strategy sets of each player? What is the NFG?

Finally, if you are so inclined, complicate the story even more by including a Nature move, say whether or not it is so foggy that the boats can't see each other.

The point is to practice writing out EFGs and converting them to NFGs.

4 Behavioral Considerations

The most interesting behavioral considerations arise once we have predictions of what players will do in a given strategic situation. We'll develop solution concepts in the next few chapters, and the sometimes discuss their predictive accuracy.

For now, we might just say something about payoffs.

- In principle, payoffs are utilities, given some Bernoulli function for each player.
- So positive linear rescaling (separately for each player) should have no effect on the predicted behavior.
- But, empirically speaking, utilities are difficult or impossible to observe.
- The simplest way to make predictions about game outcomes is to assume that payoffs are monetary, and players are risk neutral (or approximately so over the relevant range).
- We'll later see that this assumption is reasonable in some contexts.
- But in other cases, as we will see, players may

1. exhibit apparent risk aversion, or
2. care about others' monetary payoffs as well as their own.
3. In such cases, the predictions need modification.

- We will return to these matters in later chapters.

5 Readings

TBAdded