به نام خدا

An introduction to Game Theory

Lecture 3-4

By: M. Khakestari

University of Economics Sciences

Fall 2013

Game Theory and Its Application In Economics

(Examples : Static Games of Complete Information)

Objectives

- On completion of this presentation you should:
  - understand the place of game theory in Economics
  - be able to represent and solve simple games
  - apply game theory to the issue of collusion
  - model Cournot, Bertrand and von Stackelberg competition
  - be able to take a game-theoretic approach to entry deterrence
  - appreciate the limits of game theory

Cournot and Bertrand Competition

- Cournot competition
  - two firms, identical products, firms choose output levels
  - each firm's profit-maximising output depends on the other firm's output; hence each firm has a reaction function
  - as each firm will operate on its reaction function the point where they cross is the Cournot Nash equilibrium
- Bertrand competition
  - two firms, identical products, firms choose price levels
  - price is forced down to marginal cost

- ***Cournot competition Model***

  Imagine that two firms are trying to decide how much of a specified homogeneous good to produce.

  Let q1 and q2 denote the quantities each firm can produce. Let q denote the sum of q1 and q2

$$p = a - bq$$

$$C_i = cq_i$$

$$q = q_1 + q_2$$

$$u_1 = (a-c)q_1 - bq_1^{\,2} - bq_1 q_2$$

$$u_2 = (a-c)q_2 - bq_2^{\,2} - bq_1 q_2$$

Cournot and Bertrand Competition

$$q_1 = \frac{a-c}{2b} - \frac{1}{2}q_2$$

$$q_2 = \frac{a-c}{2b} - \frac{1}{2}q_2$$

$$q_1^* = q_2^* = \frac{a-c}{3b}$$

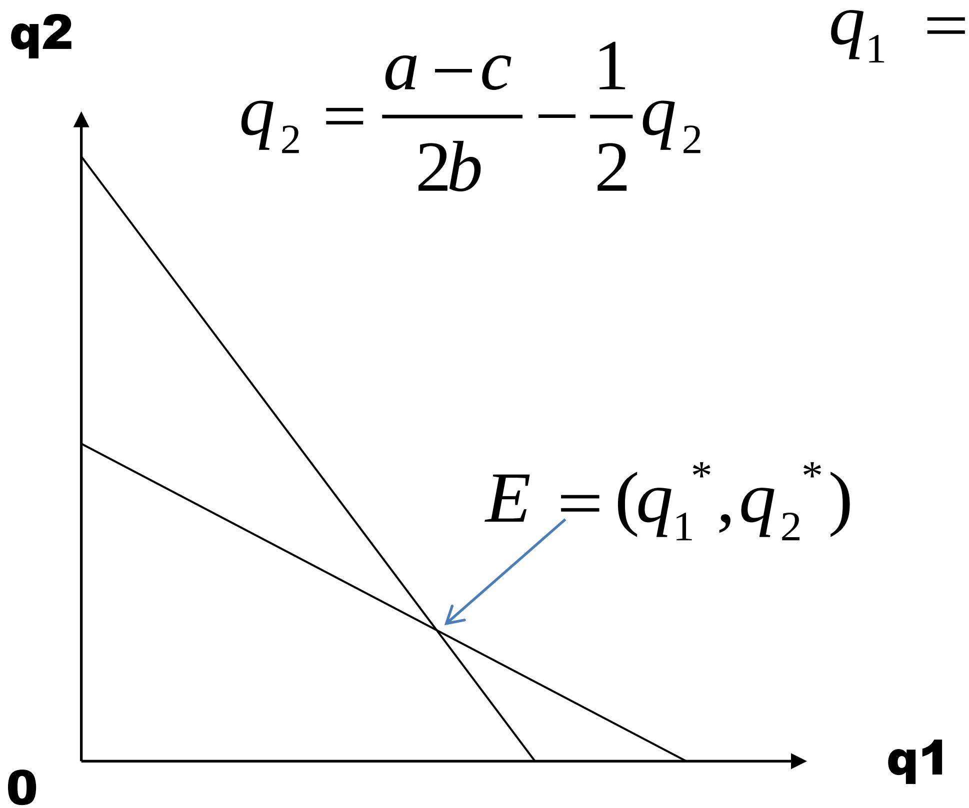

❑ *Bertrand Competition Model*

The assumptions of the model are:

1. 2 firms in the market, $i \in \{1, 2\}$.
2. Goods produced are homogenous, $\Rightarrow$ products are perfect substitutes.
3. Firms set prices simultaneously.
4. Each firm has the same constant marginal cost of c.

Cournot and Bertrand Competition

- What is the equilibrium, or best strategy of each firm?

p1 = p2 = p,

Using logical arguments:

1. Firm's will never price above the monopoly's price
2. In equilibrium, all firm's prices are the same
3. In equilibrium, prices must be at the marginal cost

Cournot and Bertrand Competition

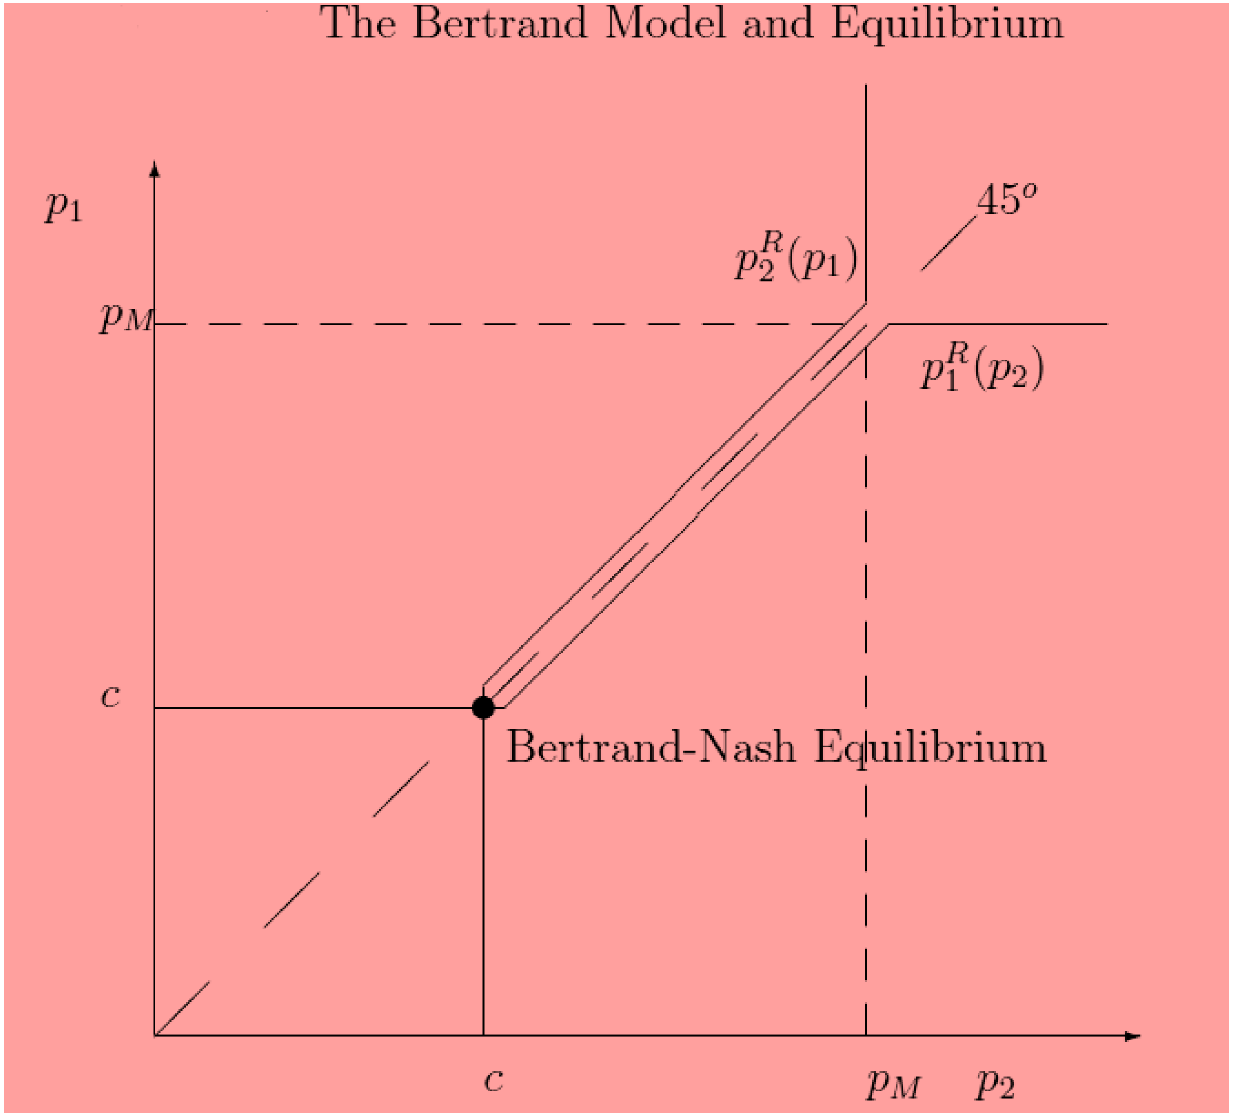

Extensive Form Games

Dynamic Games of Complete and Perfect Information

Dynamic games

- Sequential moves
  - One player moves
  - Second player observes and then moves
- Examples
  - Industrial Organization
  - Chess, Bargaining, Negotiations

Osborne, M. J. and Rubinstein, A. 1994. A Course in Game Theory. MIT Press, Cambridge, MA, England.

Sequential Decisions

- Suppose that decisions can be made sequentially
- We can work backwards to determine how people will behave
  - We will examine the last decision first and then work toward the first decision
- To do this, we will use a decision tree

Extensive Form Games

Extensive form games contain the following:

- A game tree
- A list of players
- The names of players moving at each node
- A set of allowable actions at each node
- Payoffs specified at each node Unlike normal form games, it is easy to depict sequential moves by players in extensive form games
- The decision nodes are partitioned into information sets.

Assumptions in Dynamic Extensive Form Games

- All players are rational.
- Rationality is common knowledge
- Players move sequentially. (Therefore, also called sequential games)
- Players have complete and perfect information
  - Players can see the full game tree including the payoffs
  - Players can observe and recall all previous moves

Example: Prisoners' Dilemma

Prisoner 2

Prisoner 1

	Confess	Don't Confess
Confess	5, 5	15, 0
Don't Confess	0, 15	1, 1

Prisoners' Dilemma in “Extensive” Form

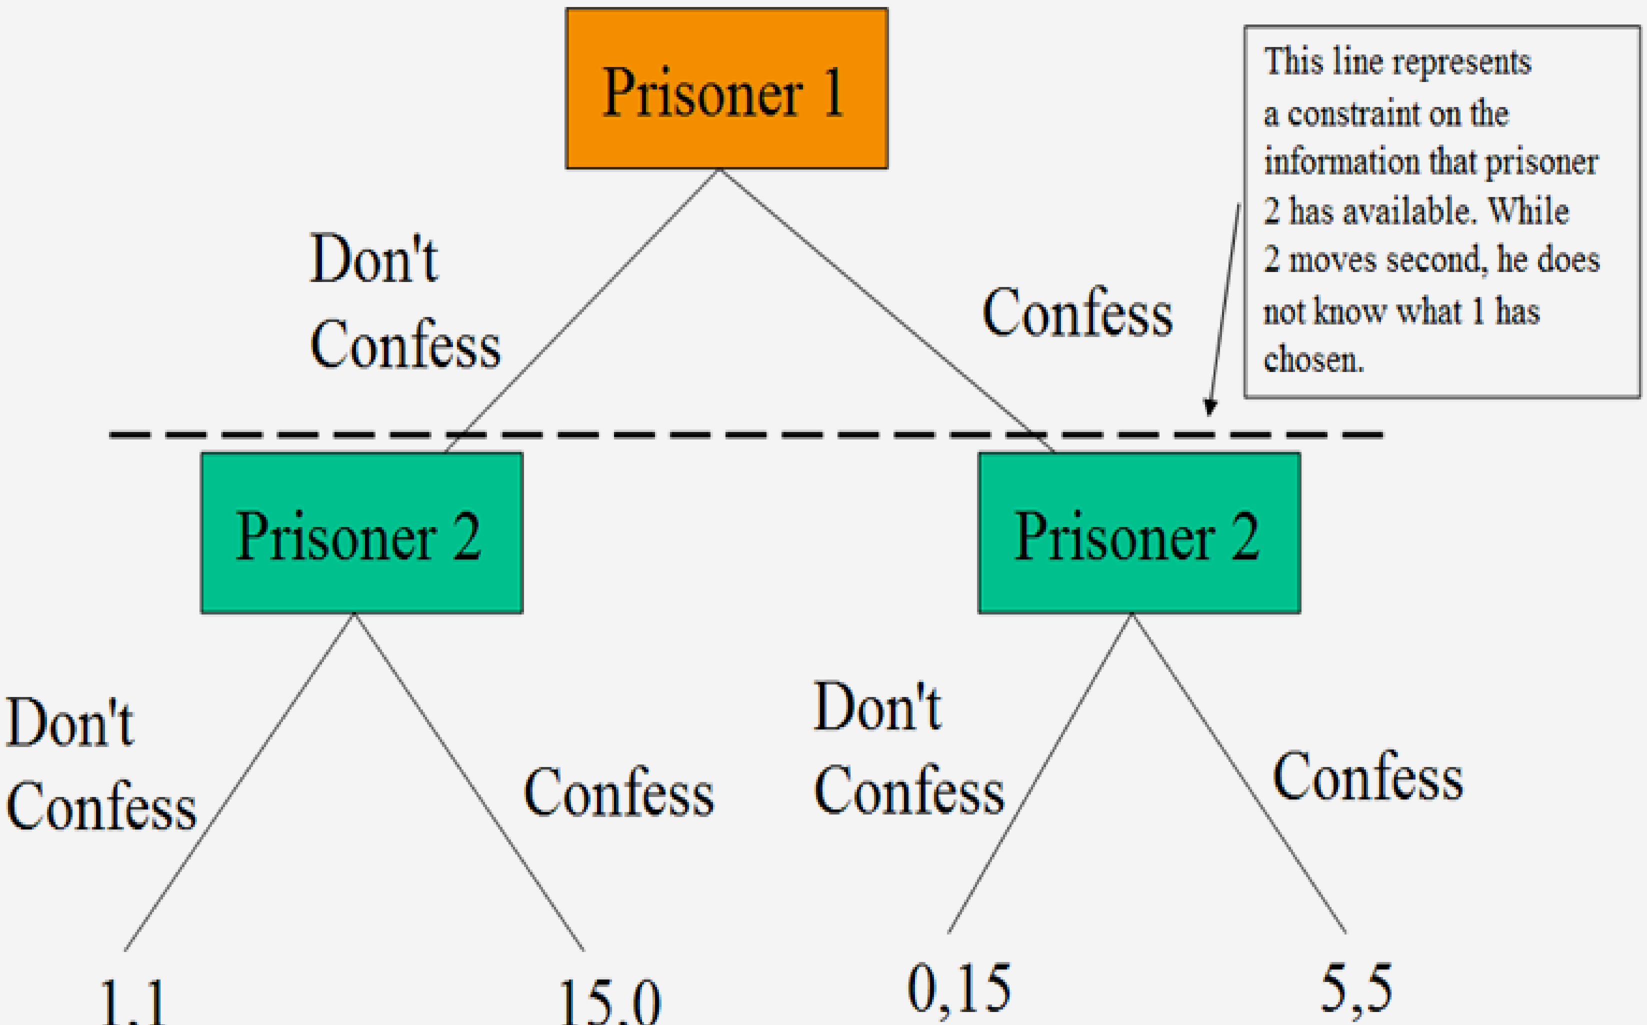

Rules that game trees must satisfy

TREE RULES

1) EVERY NODE IS THE SUCCESOR OF THE INITIAL NODE.

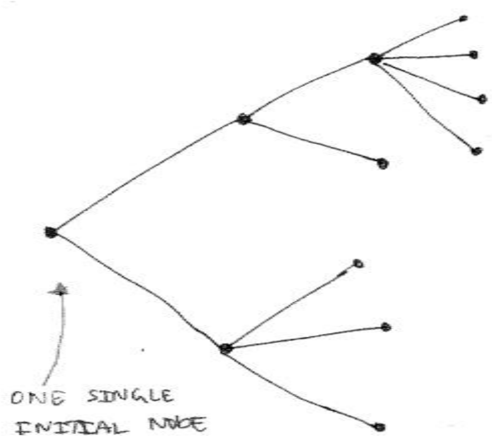

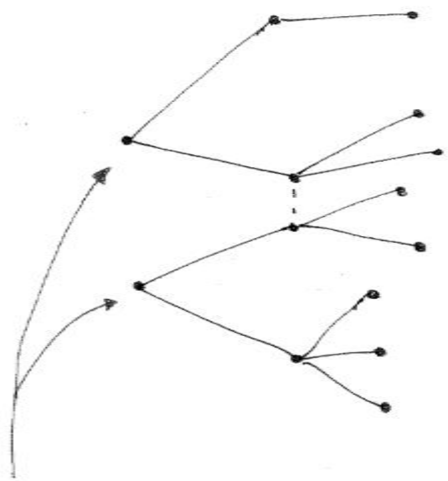

NO! TWO INITIAL NODES.

- IF THEY REFER TO THE SAME PLAYER : TWO SELVES?
- IF THEY REFER TO TWO DIFFERENT PLAYERS ACTING SIMULTANEOUSLY ; WE HAVE A WAY TO REPRESENT SUCH SITUATIONS

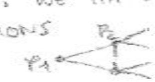

Rules that game trees must satisfy

TREE RULES

2) EVERY NODE, EXCEPT THE INITIAL NODE, HAS EXACTLY ONE IMMEDIATE PREDECESSOR.

- THE INITIAL NODE HAS NO PREDECESSOR.

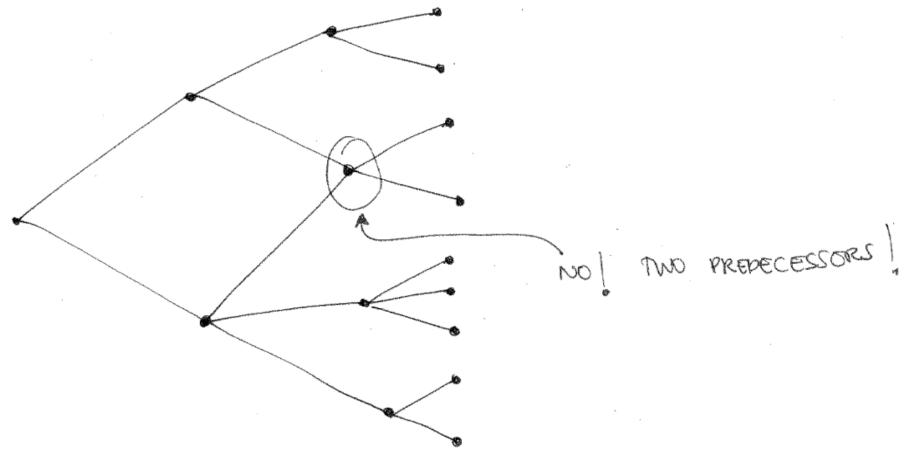

Rules that game trees must satisfy

TREE RULES

3) MULTIPLES BRANCHES, EXTENDING FROM THE SAME NODE, HAVE DIFFERENT ACTION LABELS.

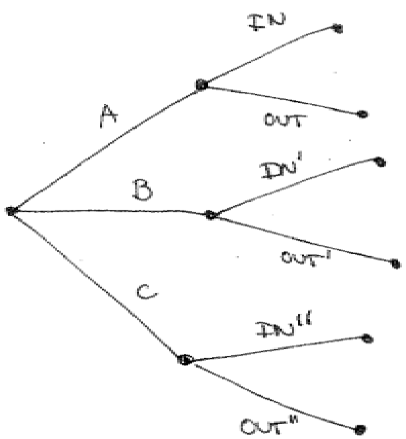

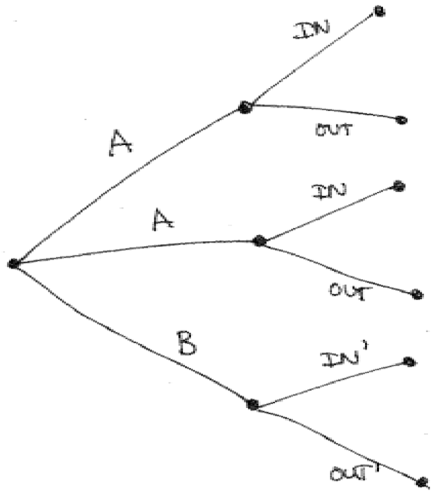

CORRECT

NO, YOU MUST BE REFERING TO A DIFFERENT ACTION... OTHERWISE COLLAPSE EVERYTHING UNDER THE SAME NAME.

Rules that game trees must satisfy

TREE RULES

4) EACH INFORMATION SET CONTAINS DECISION NODES FOR ONLY ONE OF THE PLAYERS

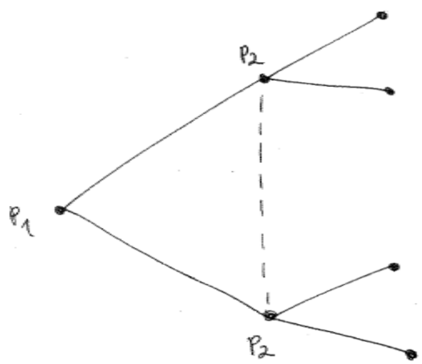

CORRECT

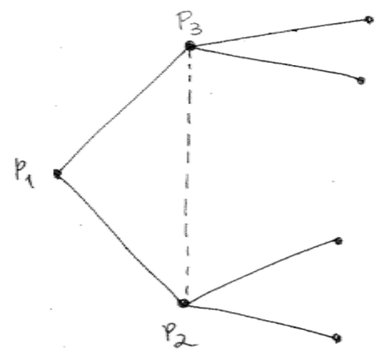

INCORRECT

Rules that game trees must satisfy

TREE RULES

5) ALL NODES IN A GIVEN INFORMATION SET HAVE THE SAME NUMBER OF IMMEDIATE SUCCESSORS, AND THEY MUST HAVE THE SAME NUMBER OF ACTION LABELS LEADING TO THESE SUCCESSORS.

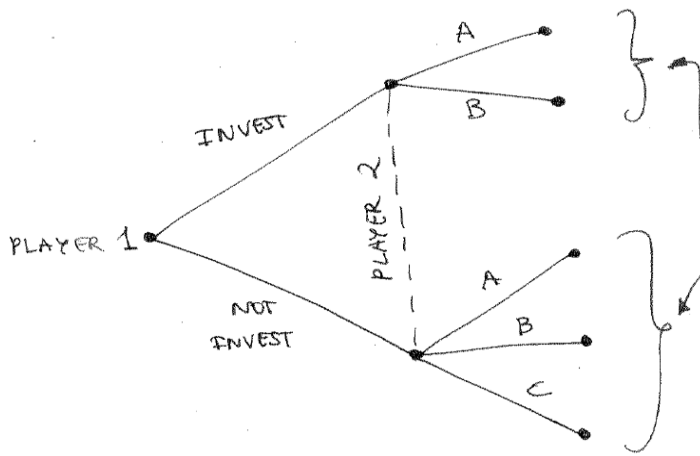

NO!

OTHERWISE PLAYER 2 WOULD KNOW WHERE HE IS (WHAT ACTION PLAYER 1 CHOSE BEFORE HIM), BY JUST OBSERVING THE SET OF AVAILABLE ACTIONS THAT IS OFFERED TO HIM, EITHER {A, B} OR {A, B, C}

Example

Guided exercise (page 19-20 in Watson)

Firm A decides whether to enter firm B's industry. Firm B observes this decision.

- If firm A stays out, firm B alone decides whether to advertise. In this case, firm A obtains zero profits, and firm B obtains \$4 million if it advertises and \$3.5 million if it does not.
- If firm A enters, both firms simultaneously decide whether to advertise, obtaining the following payoffs.
    - If both advertise, both firms earn \$3 million.
    - If both don't advertise, both firms earn \$5 million.
    - If only one firm advertises, then it earns \$6 millino and the other earns \$1 million.

Example

A earns nothing if it does not enter.

Solution: Let E and D denote firm A's initial alternatives of entering and not entering B's industry. Let a and n stand for "advertise" and "not advertise," respectively. Then the following extensive-form diagram represents the strategic setting.

GUIDED EXERCISE, page 19 WATSON

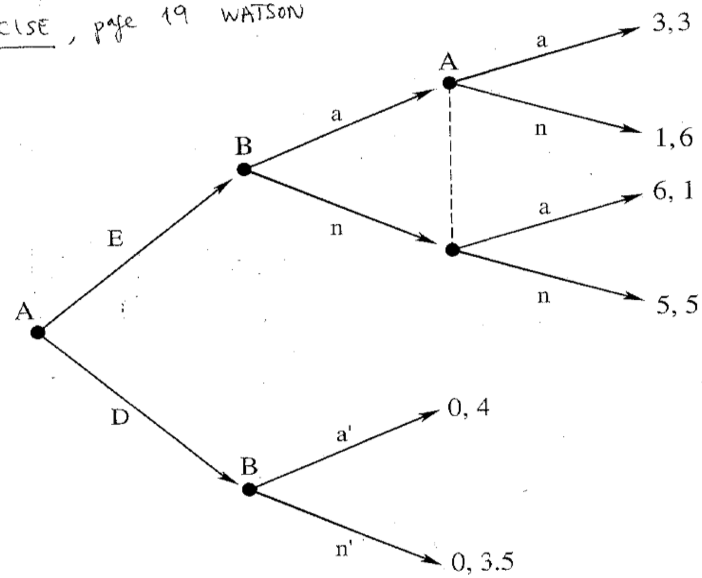

Note that simultaneous advertising decisions are captured by assuming that, at firm A's second information set, firm A does not know whether firm B chose a or n. Also note that primes are used in the action labels at firm B's lower information set to differentiate them from the actions taken at B's top information set.

EXERCISES

Jan
wo
lett
tun
uni

2. Co
the
cid
age
the
If t
shi
by
is 1
shi
get
get
hir
get
the

3. Dr
tor
tw
(H
a h
the
hig
for
wh

4. Th
thr
do

Example

Pepsi and Coke are the largest soft drink manufacturers in the world. They compete in many geographic markets. They also compete in product markets of two types. The first type is in flavors; a form of product differentiation not unlike ready to eat cereal. The second type of market is in packaging. You can buy Pepsi at a fountain, with your Big Mac, or packaged in a can or bottle at a 7-11 convenience store. In this discussion we will focus on entry into a geographic market.

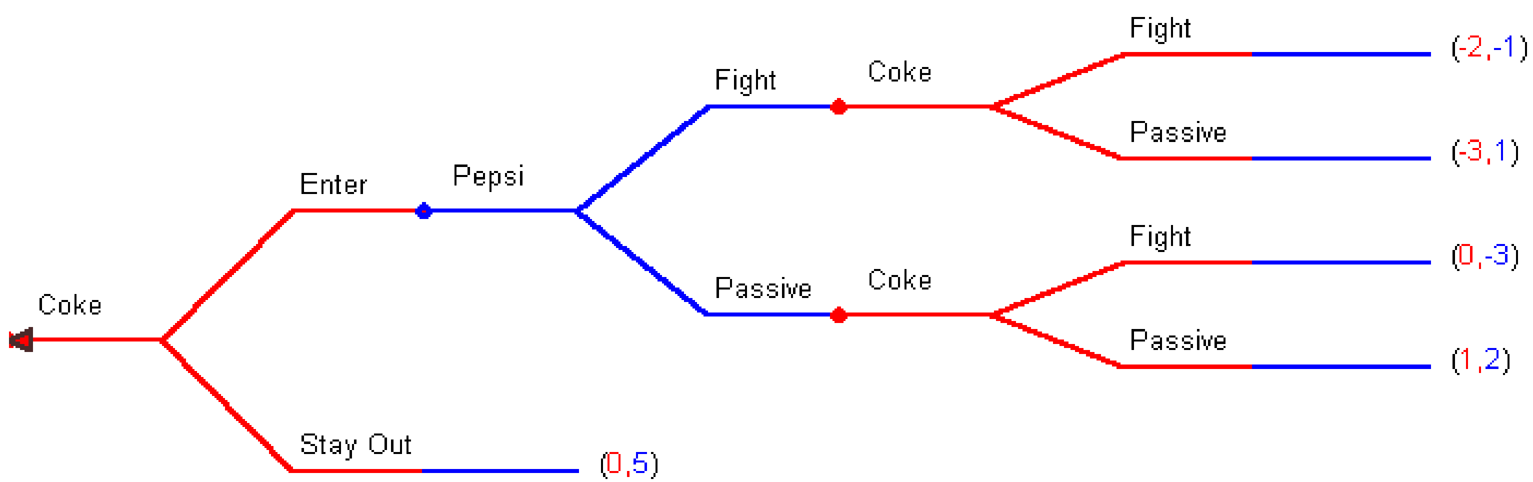

Example

<$\text{Fight}_{\text{Entry}}$, $\text{Fight}_{\text{Out}}$>	Fight when Coke enters, Fight when Coke stays out
<$\text{Fight}_{\text{Entry}}$, $\text{Passive}_{\text{Out}}$>	Fight when Coke enters, Passive when Coke stays out
<$\text{Passive}_{\text{Entry}}$, $\text{Fight}_{\text{Out}}$>	Passive when Coke enters, Fight when Coke stays out
<$\text{Passive}_{\text{Entry}}$, $\text{Passive}_{\text{Out}}$>	Passive when Coke enters, Passive when Coke stays out

Example

Coke will have eight strategic plans, shown in the following table:

<Enter, $\text{Fight}_{\text{Fight}}$, $\text{Fight}_{\text{Passive}}$>	Enter, Fight in response to Pepsi's Fight stance, Fight in response to Pepsi's Passive stance.
<Enter, $\text{Fight}_{\text{Fight}}$, $\text{Passive}_{\text{Passive}}$>	Enter, Fight in response to Pepsi's Fight, Passive in response to Pepsi's Passive
<Enter, $\text{Passive}_{\text{Fight}}$, $\text{Fight}_{\text{Passive}}$>	Enter, Passive in response to Pepsi's Fight, Fight in response to Pepsi's Passive
<Enter, $\text{Passive}_{\text{Fight}}$, $\text{Passive}_{\text{Passive}}$>	Enter, Passive in response to Pepsi's Fight, Passive in response to Pepsi's Passive
<Out, $\text{Fight}_{\text{Fight}}$, $\text{Fight}_{\text{Passive}}$>	Stay Out, Fight in response to Pepsi's Fight stance, Fight in response to Pepsi's Passive stance.
<Out, $\text{Fight}_{\text{Fight}}$, $\text{Passive}_{\text{Passive}}$>	Stay Out, Fight in response to Pepsi's Fight, Passive in response to Pepsi's Passive
<Out, $\text{Passive}_{\text{Fight}}$, $\text{Fight}_{\text{Passive}}$>	Stay Out, Passive in response to Pepsi's Fight, Fight in response to Pepsi's Passive
<Out, $\text{Passive}_{\text{Fight}}$, $\text{Passive}_{\text{Passive}}$>	Stay Out, Passive in response to Pepsi's Fight, Passive in response to Pepsi's Passive

Example

		Pepsi	
		Fight	Passive
Coke	<Enter, Fight$_{Fight}$, Fight$_{Passive}$>	-2, -1	0, -3
	<Enter, Fight$_{Fight}$, Passive$_{Passive}$>	-2, -1	1, 2
	<Enter, Passive$_{Fight}$, Fight$_{Passive}$>	-3, -1	0, -3
	<Enter, Passive$_{Fight}$, Passive$_{Passive}$>	-3, -1	1, 2
	<Out, Fight$_{Fight}$, Fight$_{Passive}$>	0, 5	0, 5
	<Out, Fight$_{Fight}$, Passive$_{Passive}$>	0, 5	0, 5
	<Out, Passive$_{Fight}$, Fight$_{Passive}$>	0, 5	0, 5
	<Out, Passive$_{Fight}$, Passive$_{Passive}$>	0, 5	0, 5

Example

Things to observe about the game as shown:
1. Using IEDS: For Coke, any of the Stay Out strategies weakly dominate EFF and EPF. On seeing this Pepsi will conclude that its Passive strategy dominates Fight. Then Coke sees that it is indifferent between EPP and EFP.
2. There are three Nash equilibria: a. Pepsi plays Fight and Coke chooses to stay out. b. Coke plays EFP and Pepsi plays Passive. c. Coke plays EPP and Pepsi plays Passive. The observed outcomes in 2.b. and 2.c. are the same.
3. The IEDS solution is a Nash equilibrium and is also the solution to the extensive form of the game using backward induction.

Example 2: In this version of the model Pepsi and Coke must make their Fight - Passive decision simultaneously. That is, there is imperfect information. However, the payoffs are such that the solution is equivalent to using IEDS in the strategic form of the game.

Example

$$N = \{1, 2\}, \quad \text{Players: } P_1, P_2$$

$$u_1 = 7s_1 - s_1{}^2 - s_1 s_{21}, \quad u_2 = 8s_2 - s_2{}^2 - s_1 s_2$$

$$s_1 = f_1\ (s_2) = \frac{1}{2}(7 - s_2), \quad s_2 = f_2\ (s_1) = \frac{1}{2}(8 - s_1)$$

$$s_1{}^* = 2, \quad s_1{}^* = 4,$$

$$s_2{}^* = 3, \; s_2{}^* = 9.$$

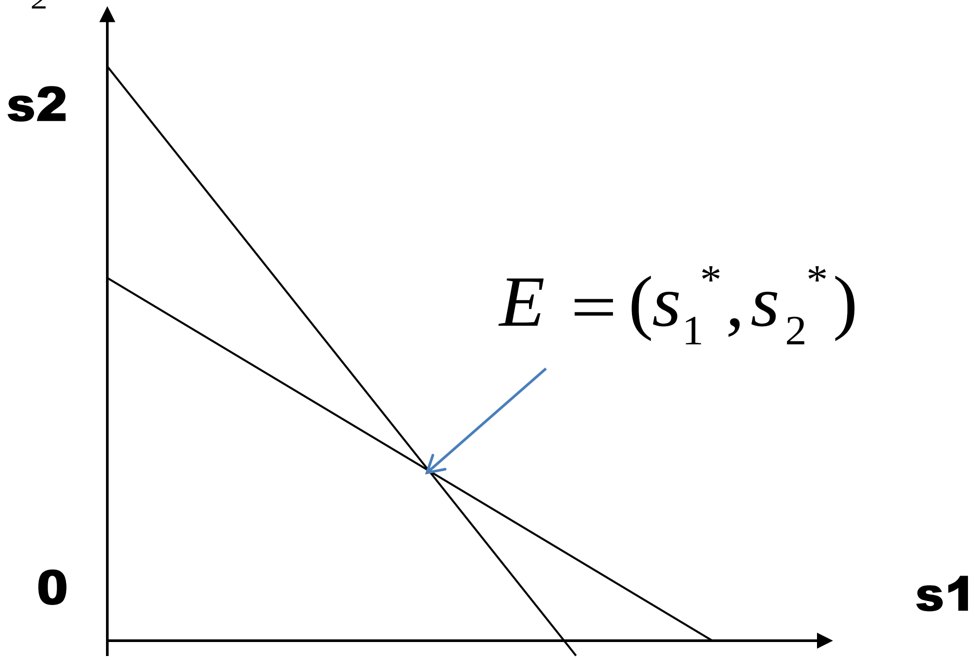

Backward Induction

The conventional method of analyzing perfect-information trees is the backward-induction algorithm. This involves going to the end of the tree and working back towards the beginning.

The first step

In the algorithm is to assign to the last player to move, the choice that maximizes that player's payoff.

The second step

Is then to turn to the second-to-last player and, taking the last player's choice as determined in the first step, to assign to the second-to-last player the choice that maximizes her payoff. And so on.

Backward Induction

Theorem

Backward induction gives the entire set of SPE.

Proof: backward induction makes sure that in the restriction of the strategy profile in question to any subgame is a Nash equilibrium.

- Backward induction is straightforward for games with perfect information and finite horizon.
- For imperfect information games, backward induction proceeds similarly: we identify the subgames starting from the leaves of the game tree and replace it with one of the Nash equilibrium payoffs in the subgame.

Backward Induction

Theorem

Every finite perfect information extensive form game G has a pure strategy SPE.

Proof: Start from the end by backward induction and at each step one strategy is best response.

Theorem

Every finite extensive form game G has a SPE.

Proof: Same argument as the previous theorem, except that some subgames need not have perfect information and may have mixed strategy equilibria.

New example (Slide 57)
Decision tree in a sequential game:
Person 1 chooses first

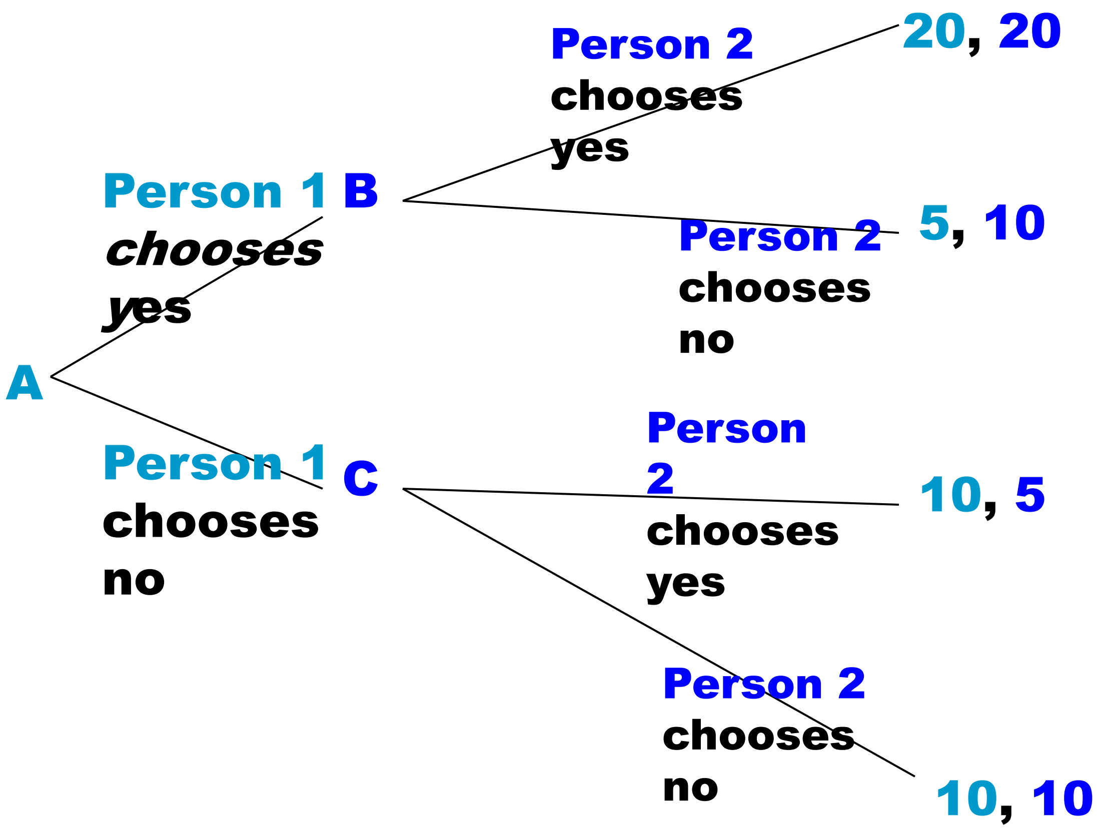

Decision tree in a sequential game: Person 1 chooses first

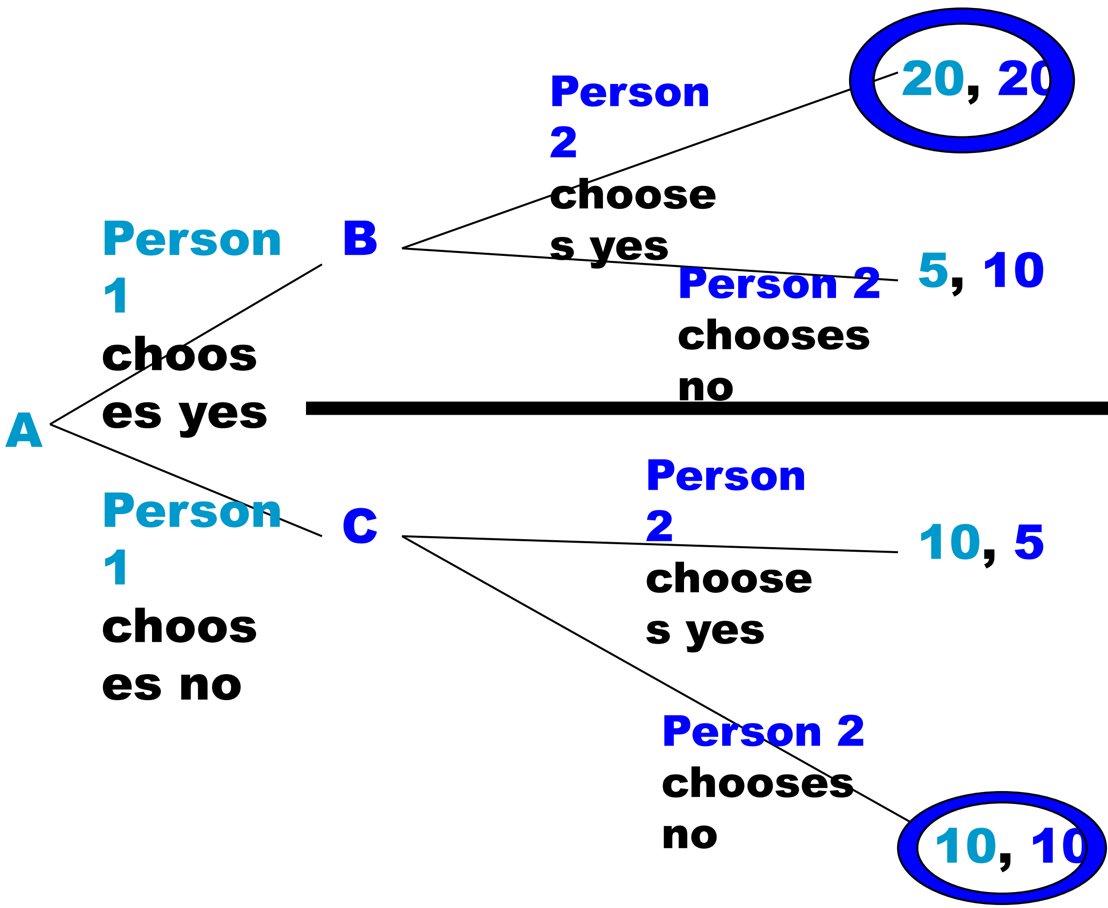

- Given point B, Person 2 will choose yes (20 > 10)
- Given point C, Person 2 will choose no (10 > 5)

Decision tree in a sequential game: Person 1 chooses first

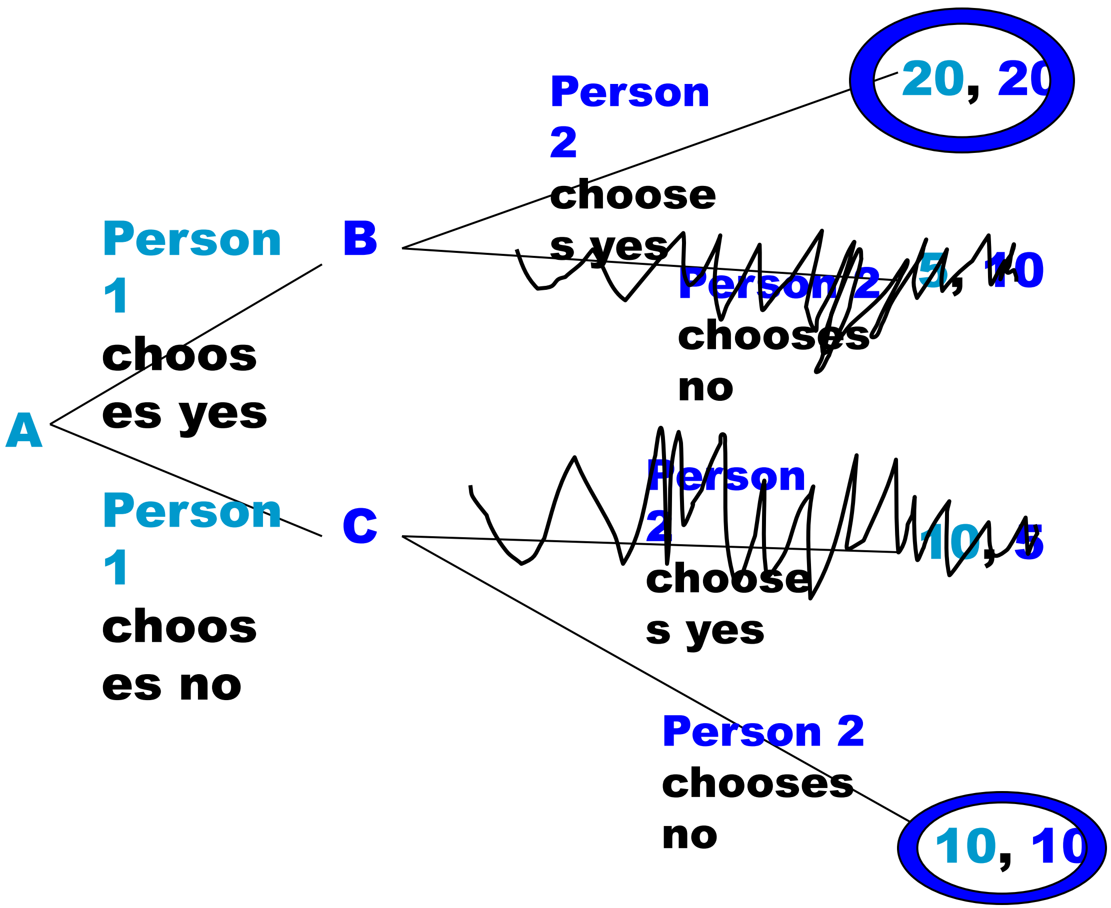

- If Person 1 is rational, she will ignore potential choices that Person 2 will not make
- Example: Person 2 will not choose yes after Person 1 chooses no

Decision tree in a sequential game: Person 1 chooses first

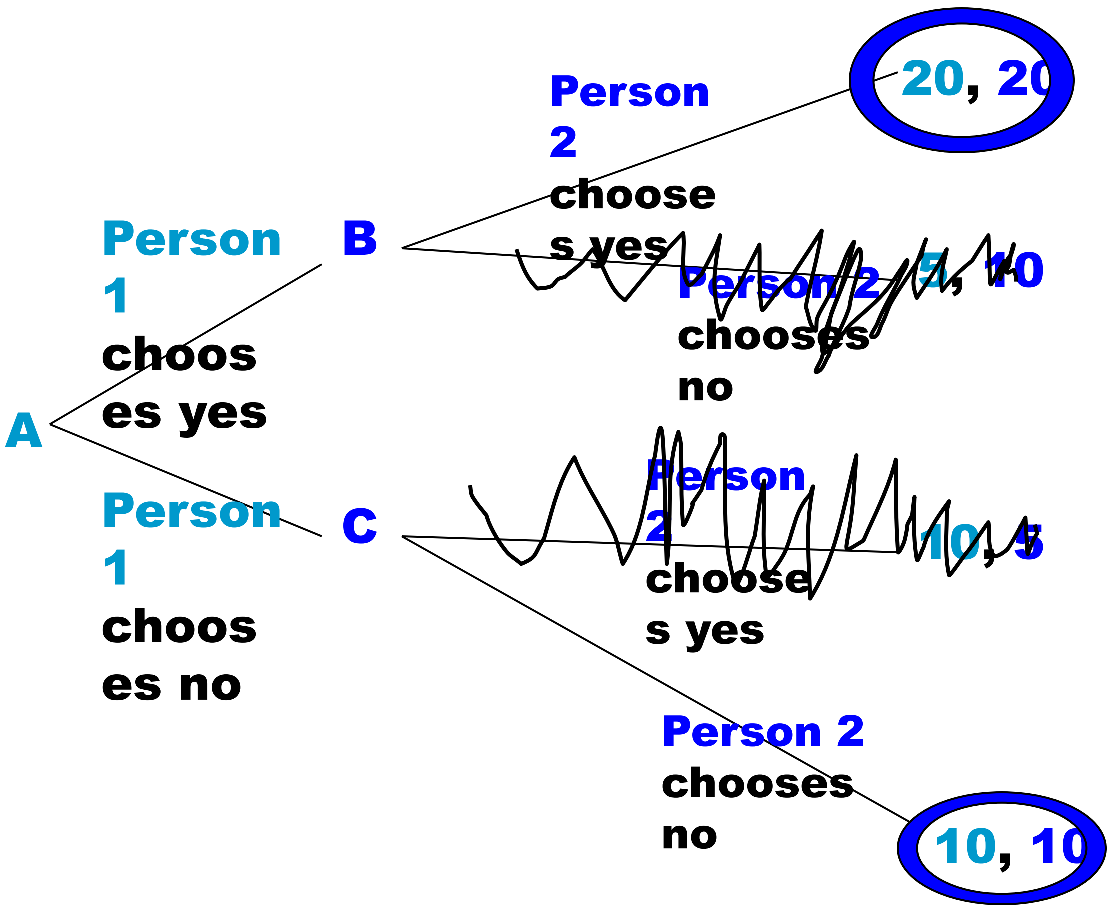

- If Person 1 knows that Person 2 is rational, then she will choose yes, since 20 > 10
- Person 2 makes a decision from point B, and he will choose yes also
- Payout: (20, 20)

Subgame

Subgame: Given an extensive form game, a node x is said to *initiate a subgame* if neither x nor any of its successors are in an information set that contains nodes that are not successors of x.

Hence, a subgame is a tree structure defined by such a node x and its successors.

Graphical representation.

Subgame-example

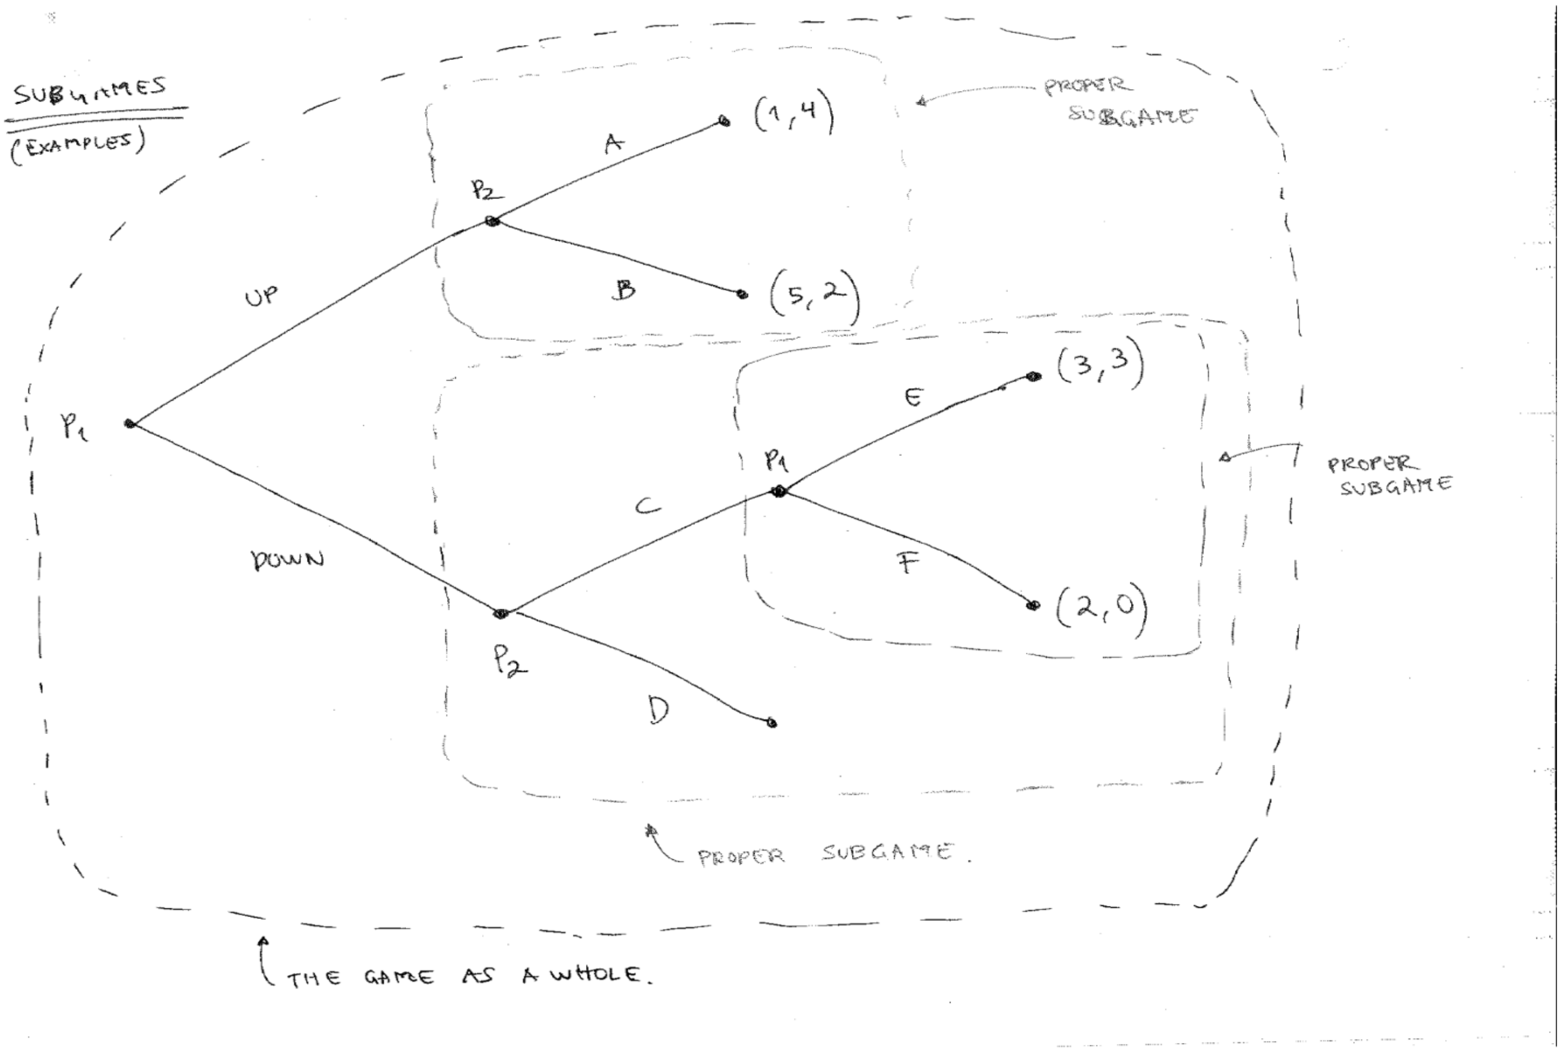

Subgame-example

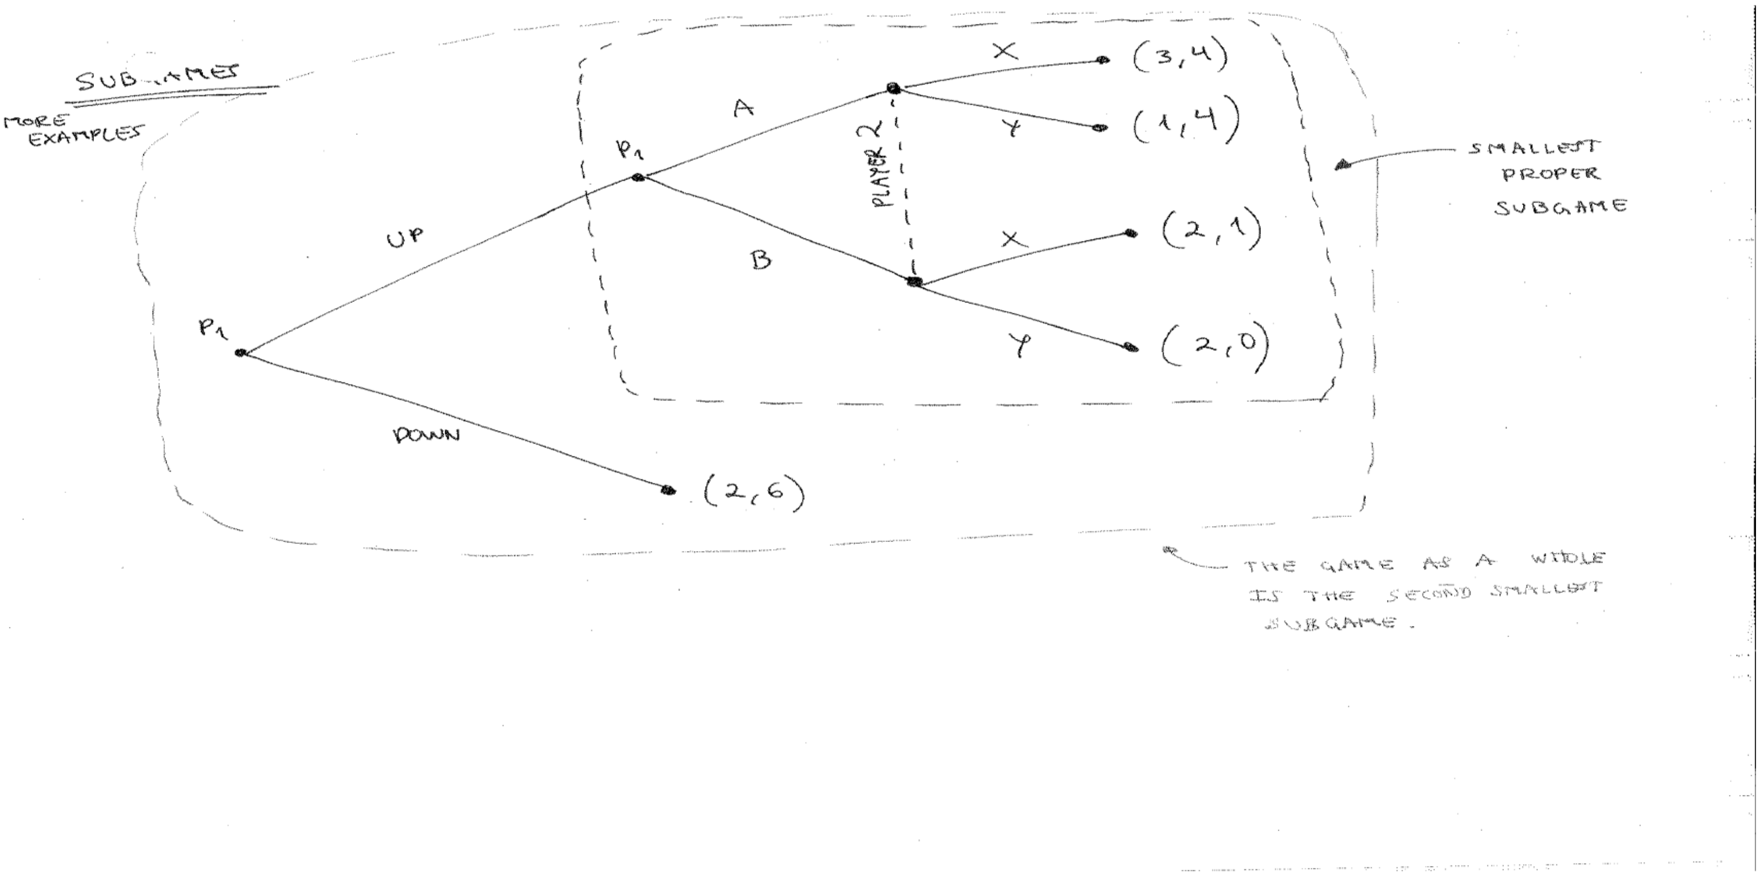

Subgame-example

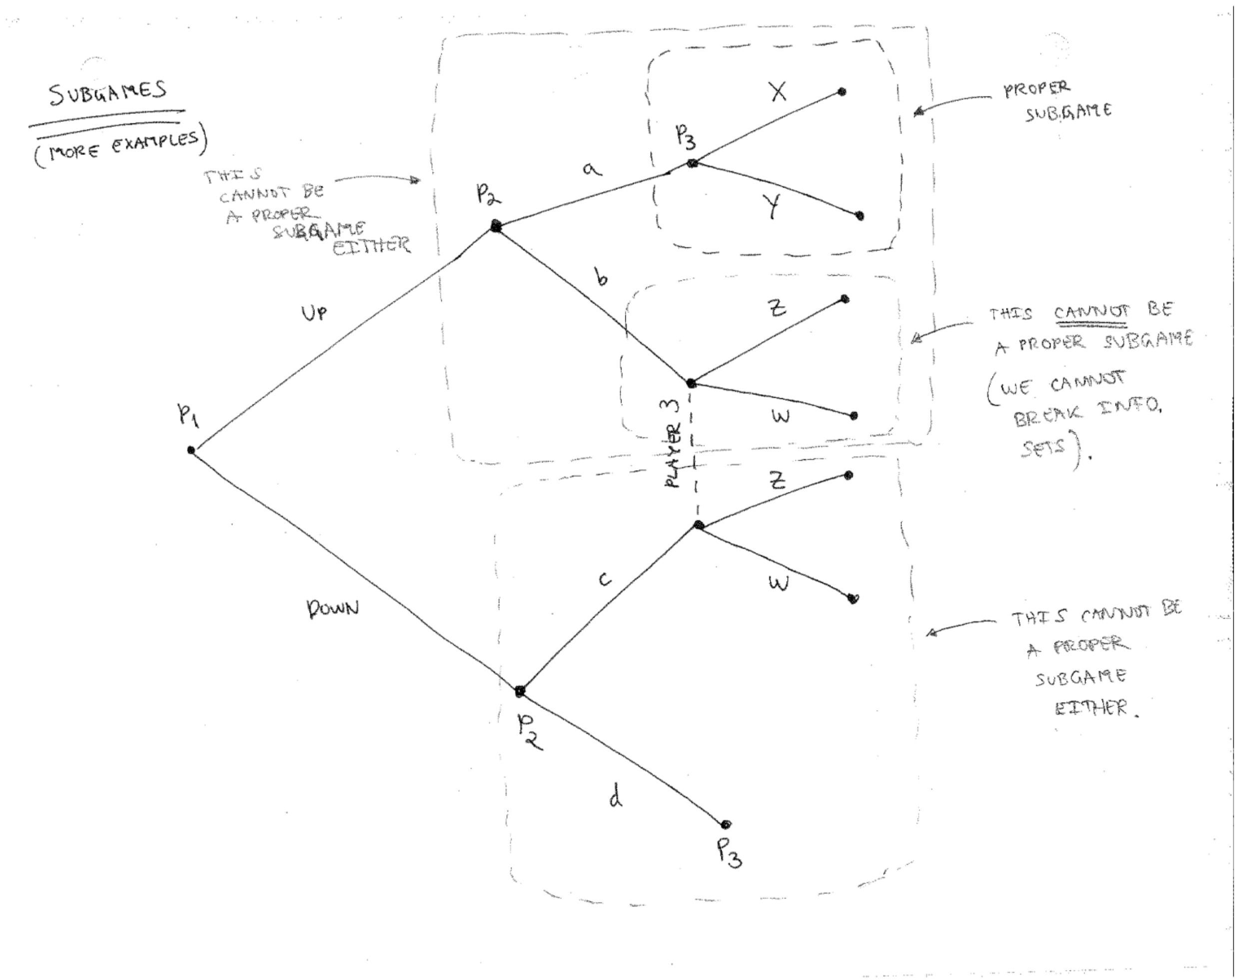

Example

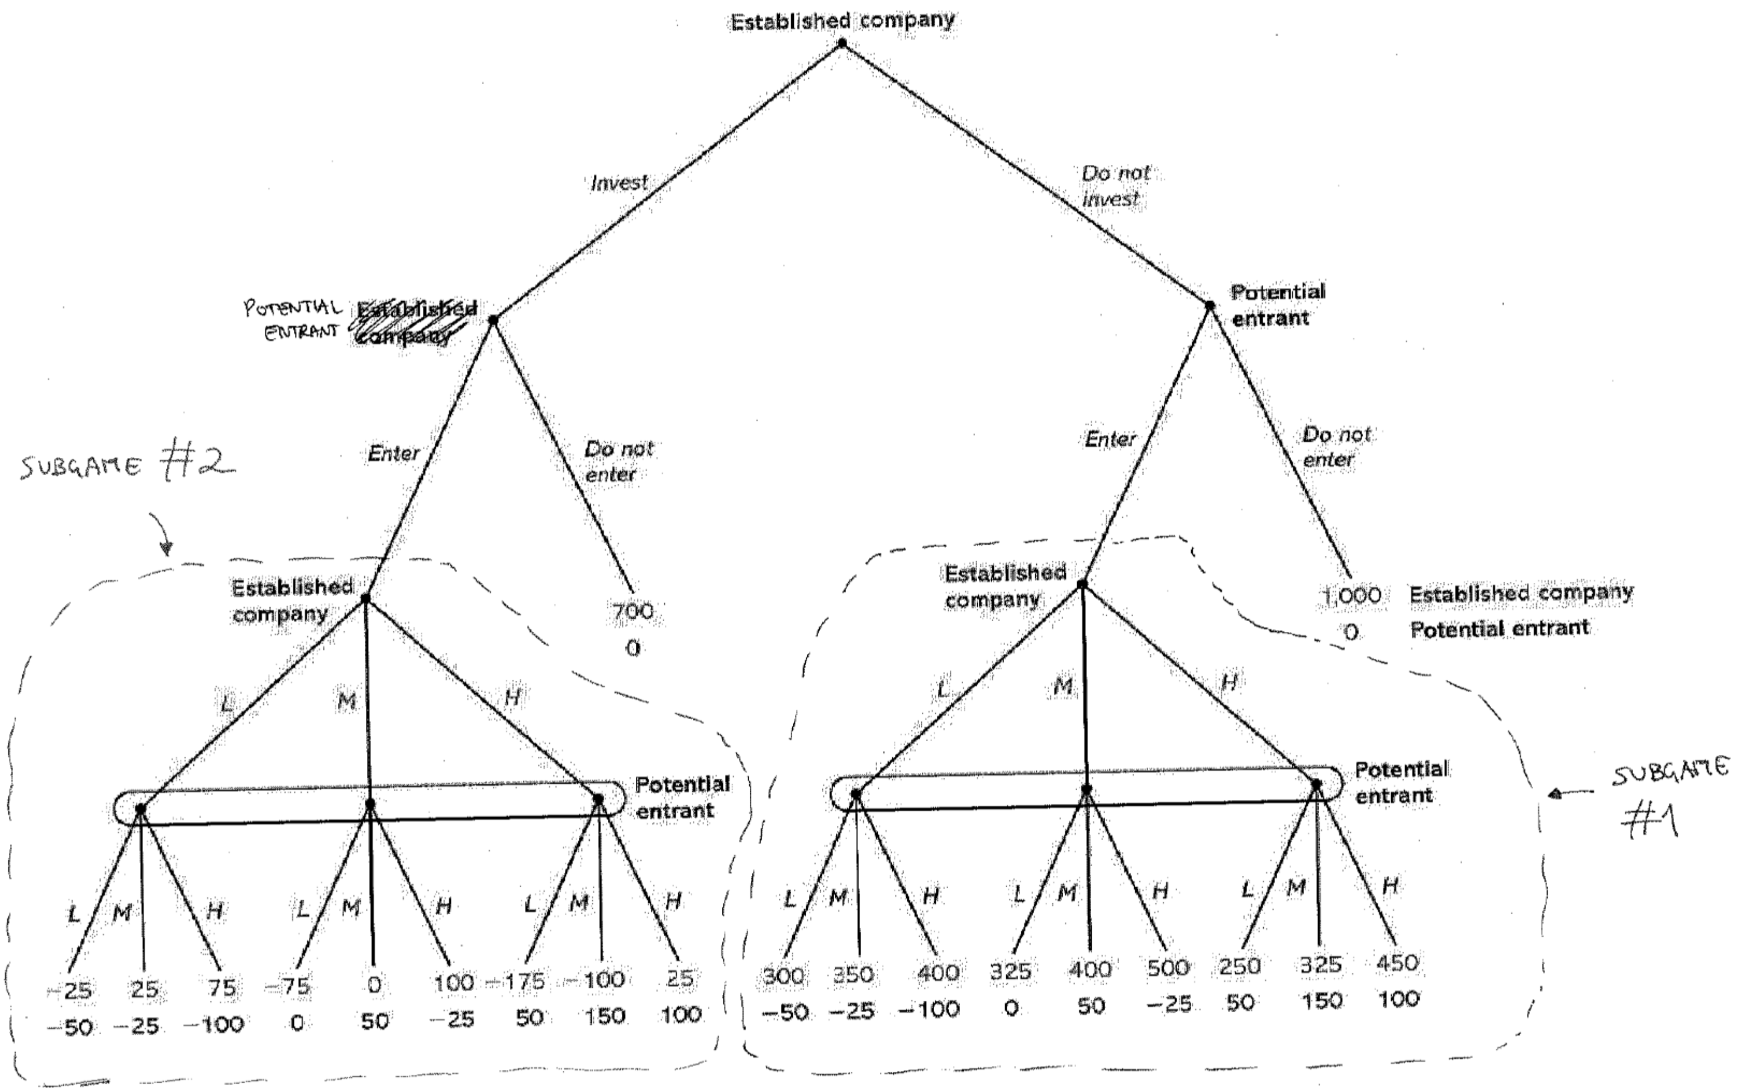

Established company
Invest
Do not invest
POTENTIAL ENTRANT
~~Established company~~
Potential entrant
SUBGAME #2
Enter
Do not enter
Enter
Do not enter
Established company
700
0
Established company
1,000
0
Established company
Potential entrant
L
M
H
Potential entrant
Potential entrant
SUBGAME #1
−25 25 75 −75 0 100 −175 −100 25
−50 −25 −100 0 50 −25 50 150 100
300 350 400 325 400 500 250 325 450
−50 −25 −100 0 50 −25 50 150 100

Subgame -Example

2 subgames

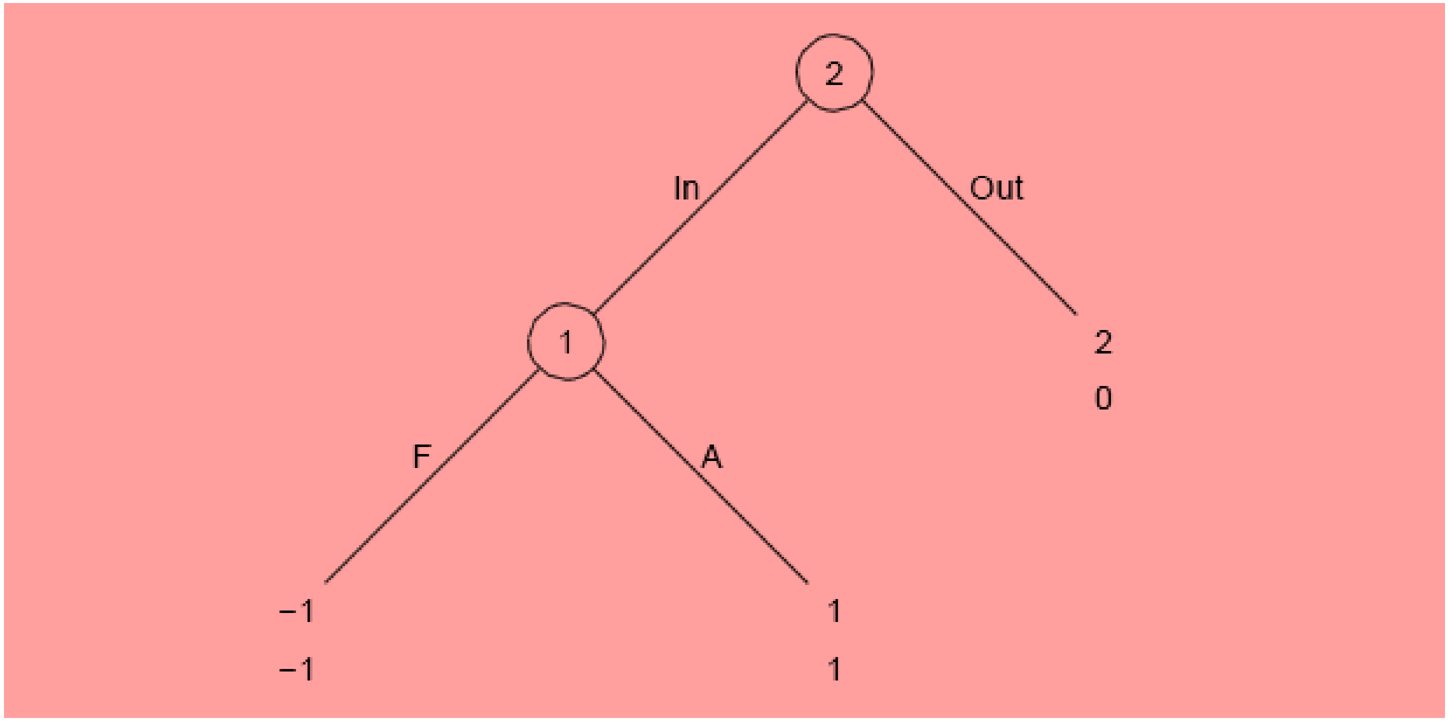

Example: Backward Induction

USING BACKWARD INDUCTION IN THE ENTRY AND PREDATION GAME

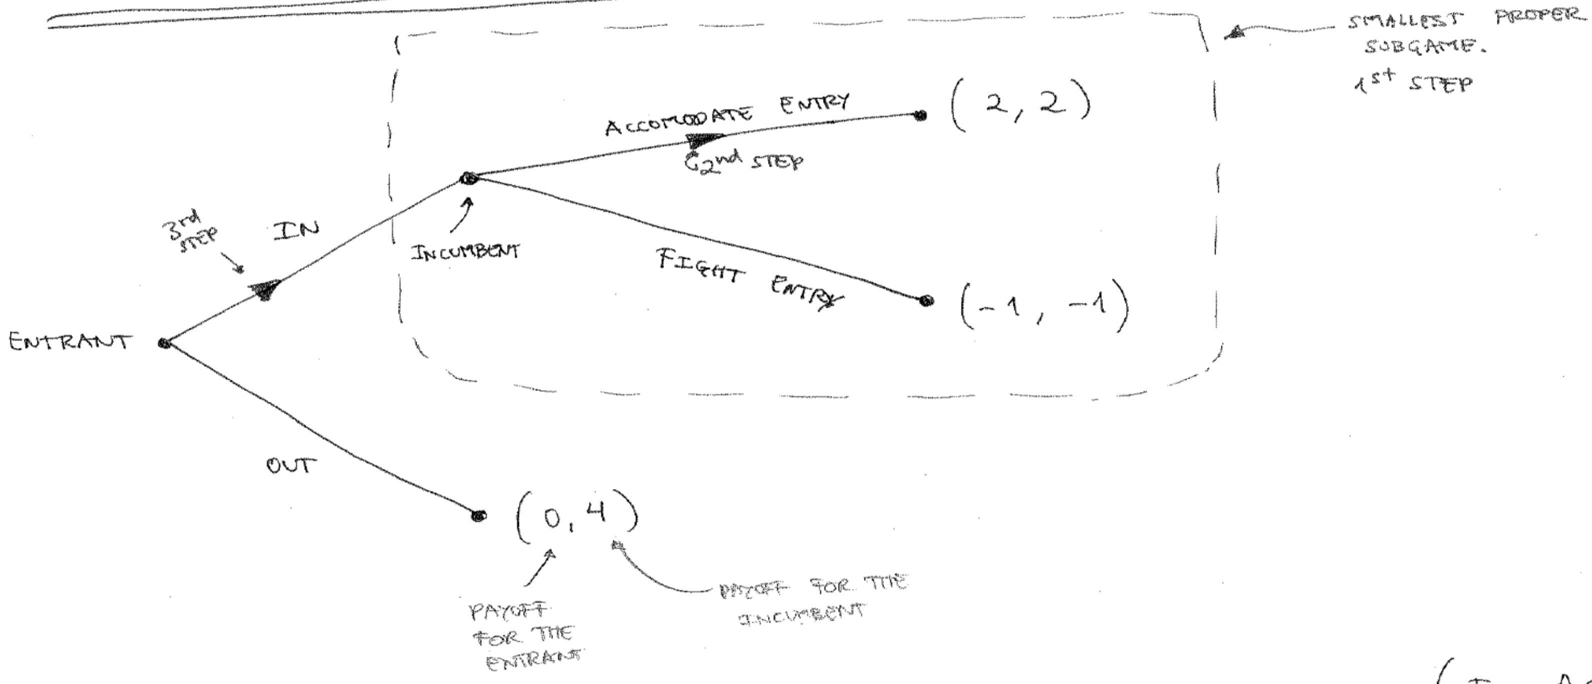

- HENCE, THERE IS ONLY ONE SUBGAME PERFECT EQUILIBRIUM IN THIS GAME: (IN, ACC)
- AMONG THE TWO psNE WE FOUND, i.e., (IN, ACC) AND (OUT, FIGHT), ONLY THE FIRST EQUILIBRIUM IS SEQUENTIALLY RATIONAL.

Example: Backward Induction

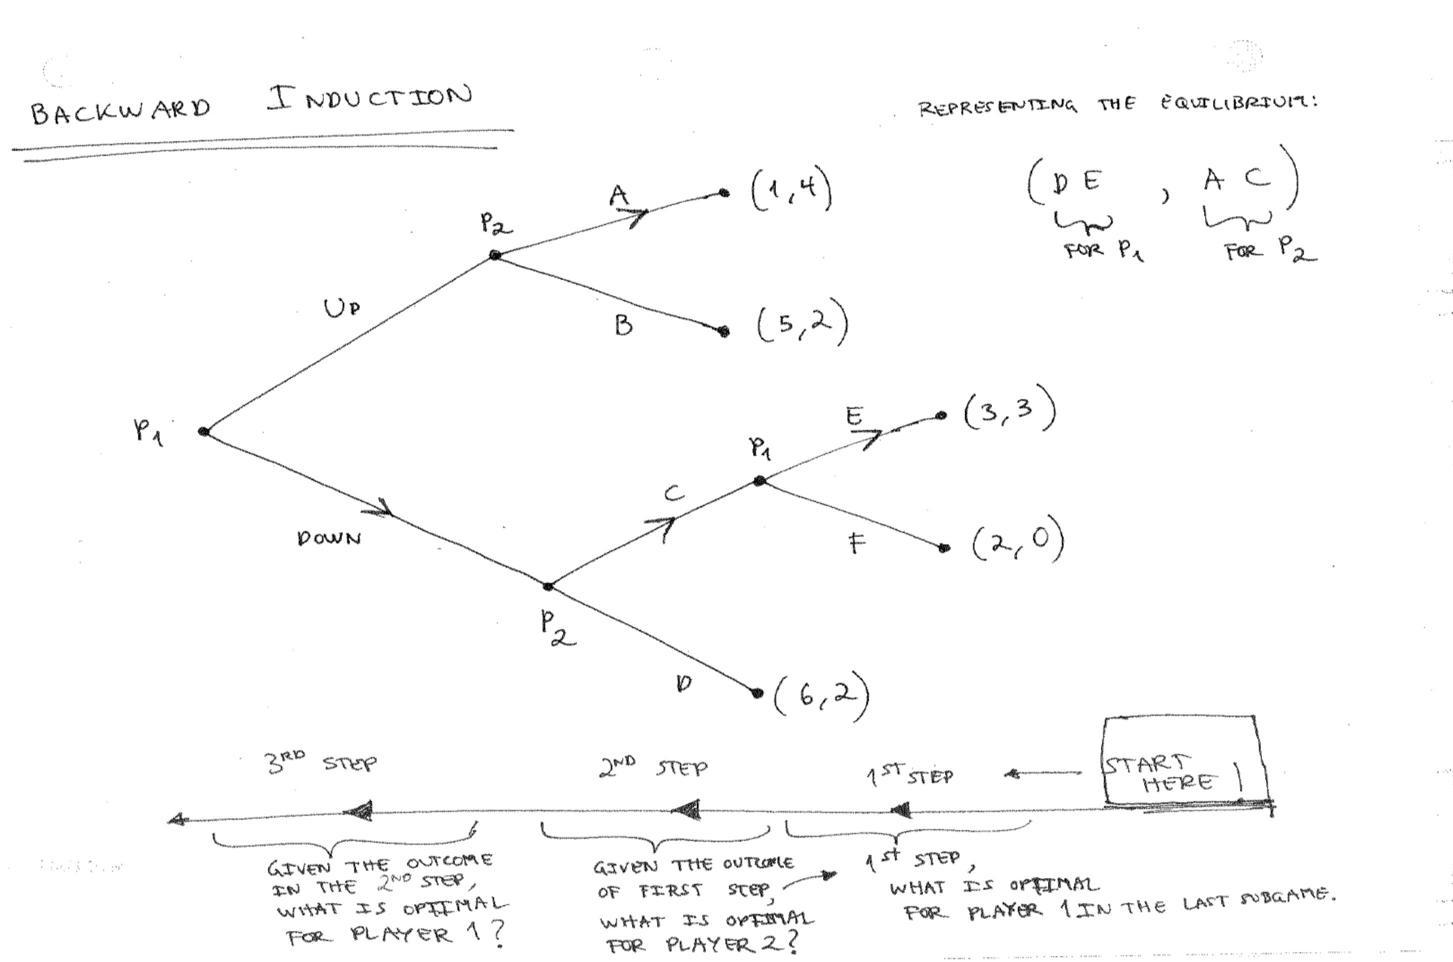

Theorem

Kuhn's Theorem

Also attributed to Zermelo

Every sequential game with perfect information has a Nash equilibrium

The theorem can be proved using backward induction.

Subgame Perfect Equilibrium

A ***subgame*** of an extensive game is a subtree of the game tree that includes all information sets containing a node of the subtree.

subgame perfect equilibrium (SPE) of the game, defined by the property that it induces a Nash equilibrium in every subgame.

Solution of an Extensive Form Game

- Subgame Perfect Equilibrium: For an equilibrium to be subgame perfect, it has to be a NE for all the subgames as well as for the entire game.
  - A subgame is a decision node from the original game along with the decision nodes and end nodes.
  - Backward induction is used to find SPE

Properties of SPE

- The outcome that is selected by the backward induction procedure is always a NE of the game with perfect information.
- SPE is a stronger equilibrium concept than NE
- SPE eliminates NE that involve incredible threats.

- ***Theorem***

Every finite extensive form game has a SPE.

Algorithm (Subgame perfect equilibrium)

Input: An extensive game.

Output: A subgame perfect Nash equilibrium of the game.

Method: Consider, in increasing order of inclusion, each subgame of the game, find a Nash equilibrium of the subgame, and replace the subgame by a new terminal node that has the equilibrium payoffs

An Advertising Example

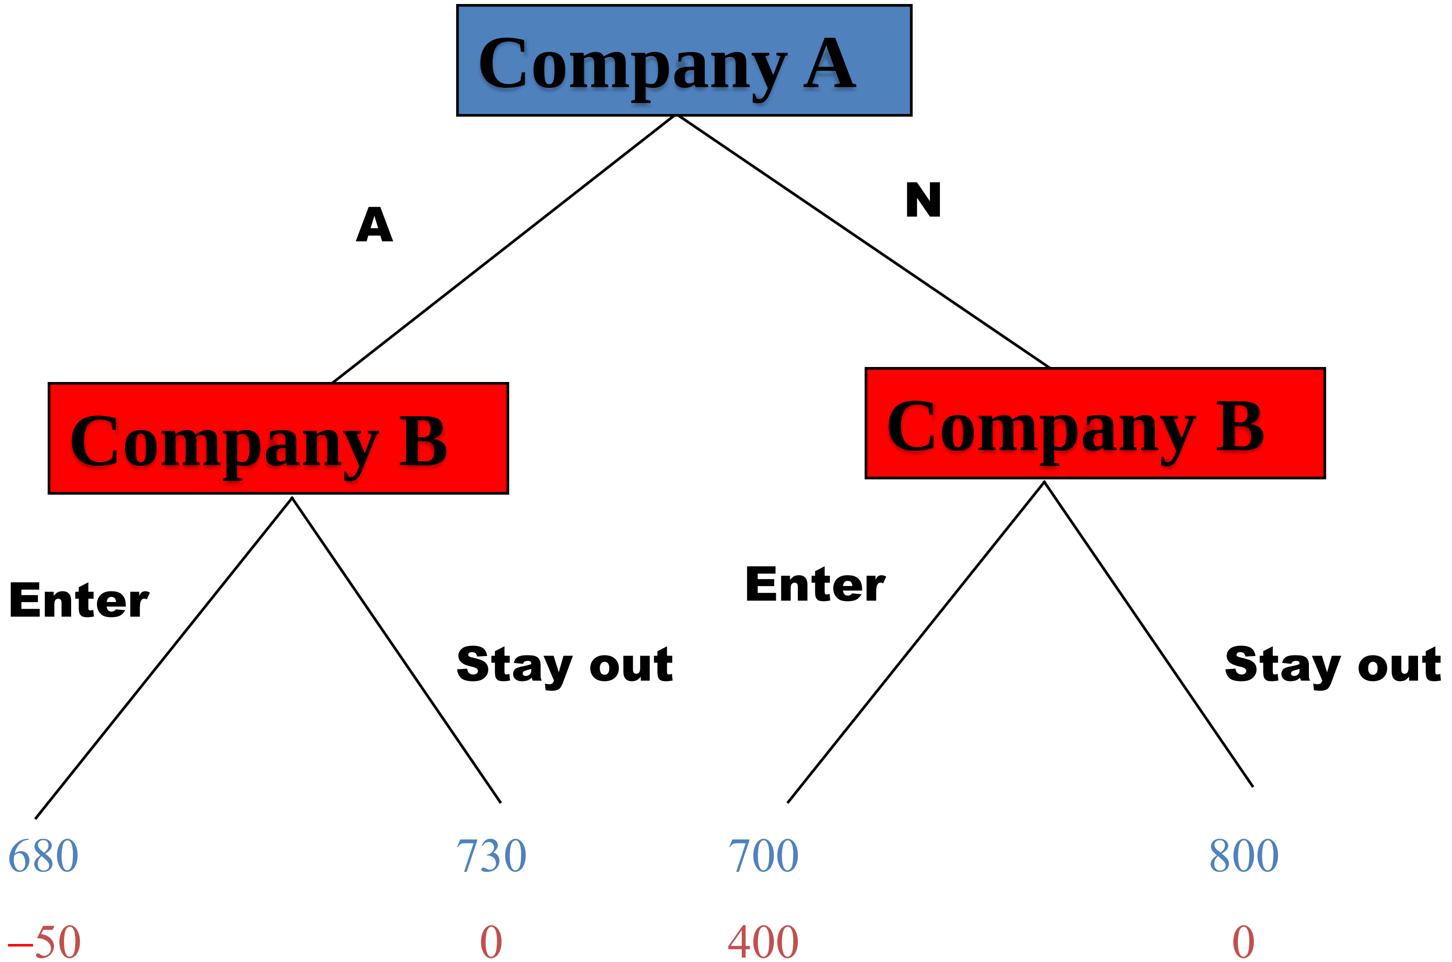

Advertising Example: 3 proper subgames

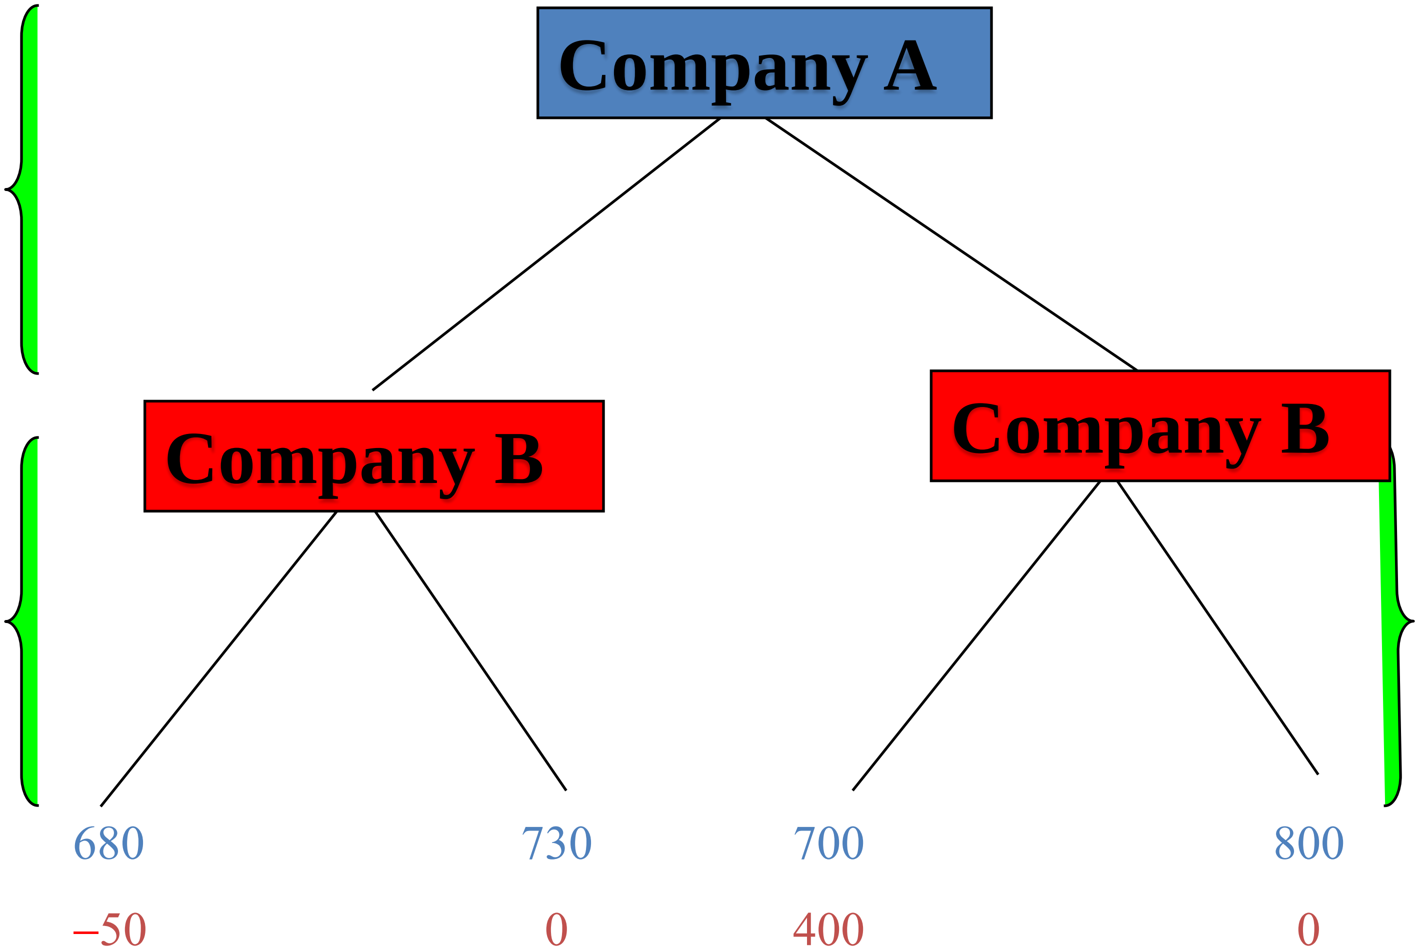

Solution of the Advertising Game

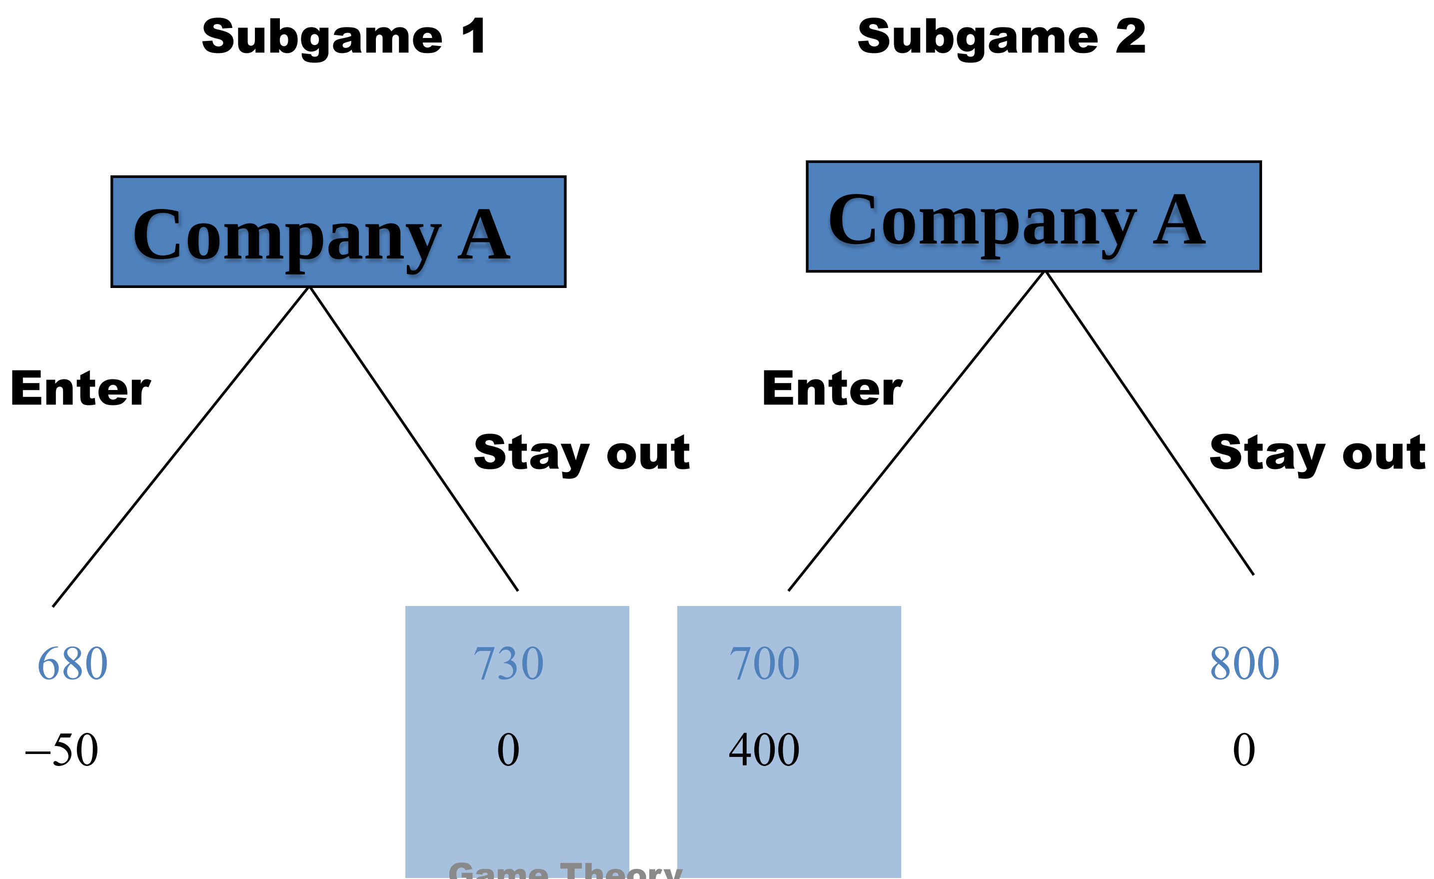

Solution of the Advertising Game (cont.)

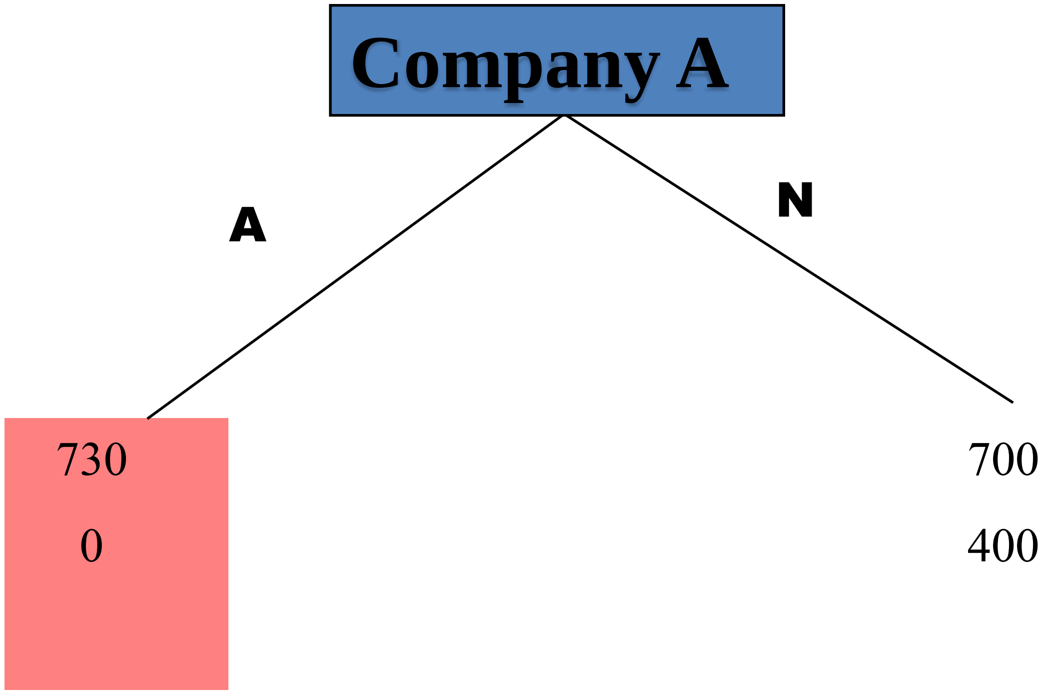

SPE of the game is the strategy profile: {A, (stay out, enter)}

Centipede Game

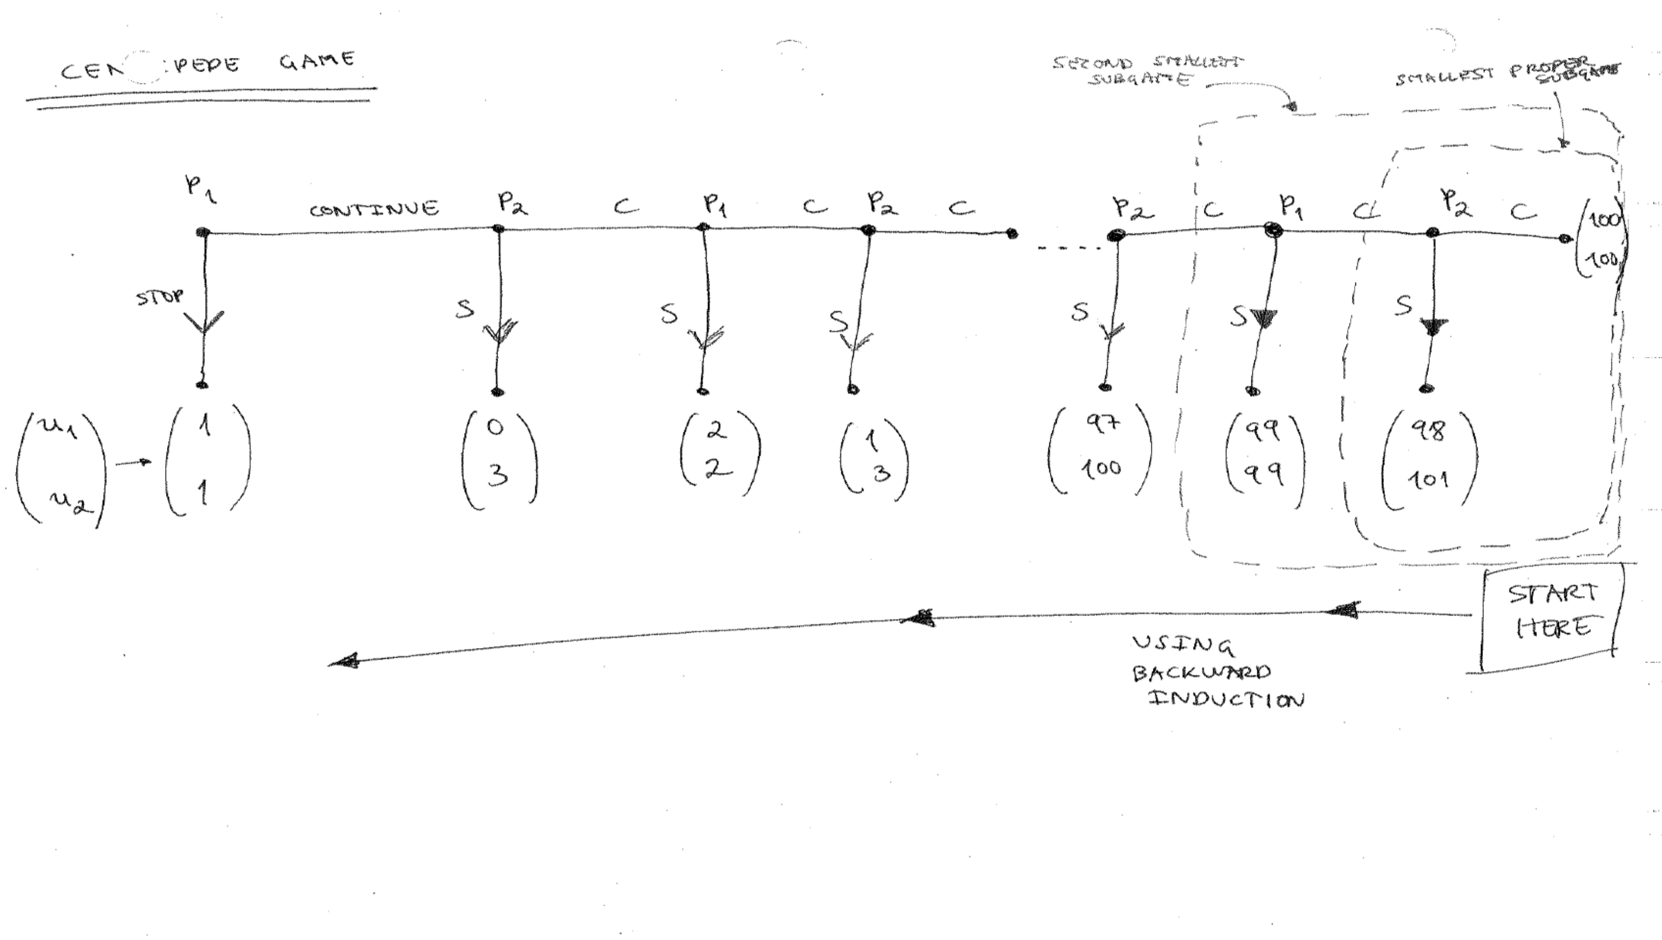

Centipede Game

Centipede game

Let us use backwards induction:

1st) In the last node, P2 is called to move, so he compares

$$u_2(Stop) > u_2(Continue) \text{ since } 101 > 100$$

so he Stops.

2nd) In the previous to the last node, P1 knows that P2 will stop at the last node, then P1 compares

$$u_1(Stop) > u_1(Continue) \text{ since } 99 > 98$$

so he Stops.

.....

In the first node, P1 knows that P2 will stop in the second stage, since P1 stops in the third, etc., so P1 compares

$$u_1(Stop) > u_1(Continue) \text{ since } 1 > 0$$

so P1 Stops.

Centipede Game

- Hence, the unique SPNE of the game is represented as (Stop_t, Stop_t) during every period $t \in T$, and for any finite lenght T of this centipede game.

Centipede Game

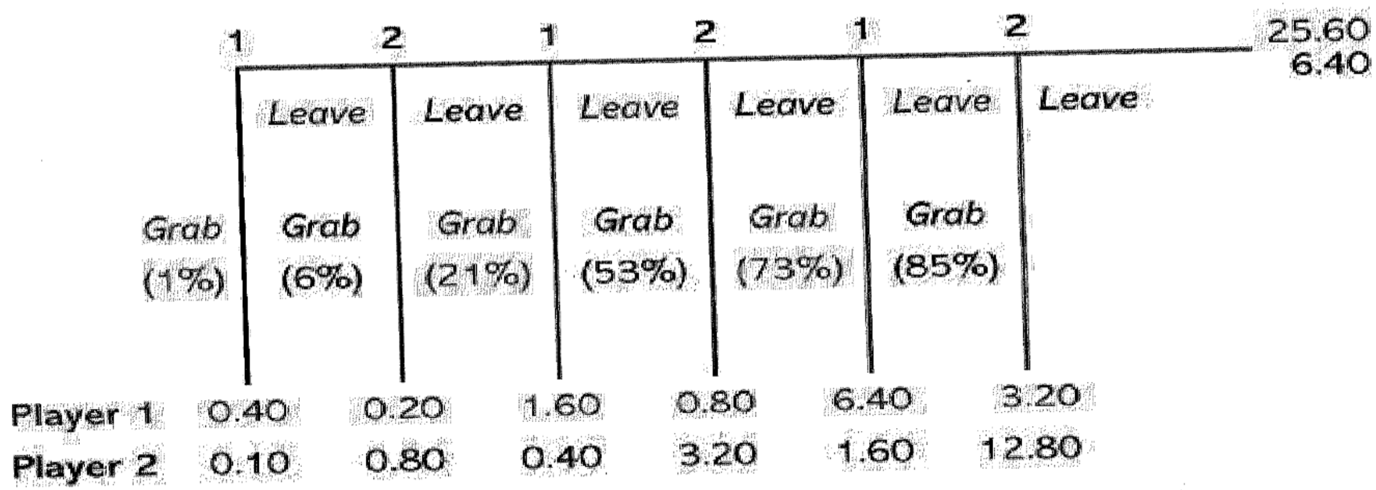

- Difference between the theoretical prediction and individuals' observed behavior in experiments.

1. **Bounded rationality.** People seem to be use backward induction relatively well in the last 1-2 stages of the game, so they can easily anticipate what their opponent will do in just a few of posterior stages.
   - We could summarize this argument as Bounded rationality, since individuals' ability to backward induct is limited, and becomes more hindered as we move further away from the terminal nodes of the game.

Centipede Game

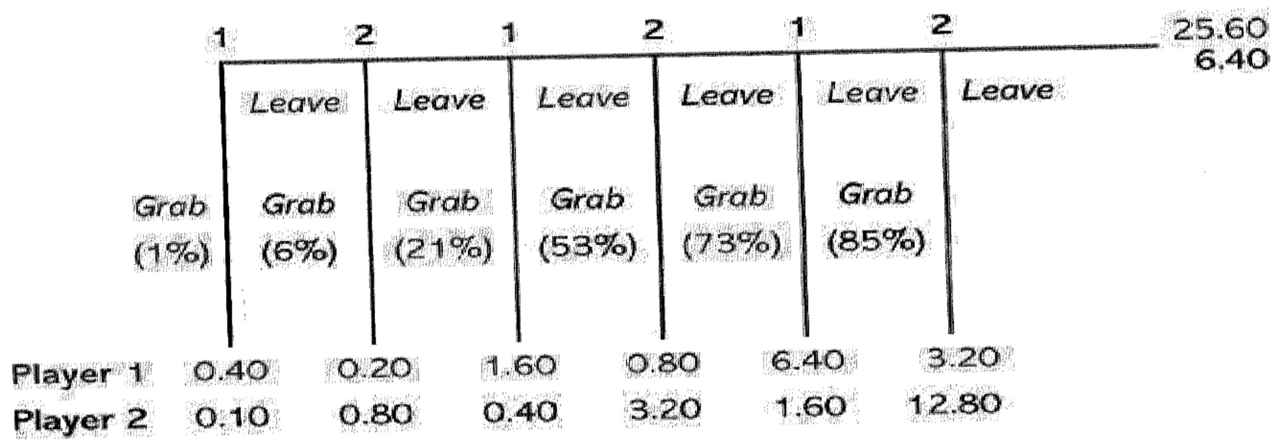

2. **Uncertainty about the presence of altruists in the population.** Another reason for their observed decision to leave money on the table could be their uncertainty about whether their opponent is an altruist.
   - If P2 is an altruist, she values not only her own money, but also the money that P1 receives. Hence, P2 would leave money on the table rather than grab it.
   - If you are in the shoes of P1 and you are uncertain about whether P2 is an altruist, you should then leave money on the table, since P2 will respond leaving it on the table as well, and wait until the last node at which you are called on to move, where you grab all the money.

Von Stackelberg model

STACKELBERG GAME OF SEQUENTIAL QUANTITY COMPETITION

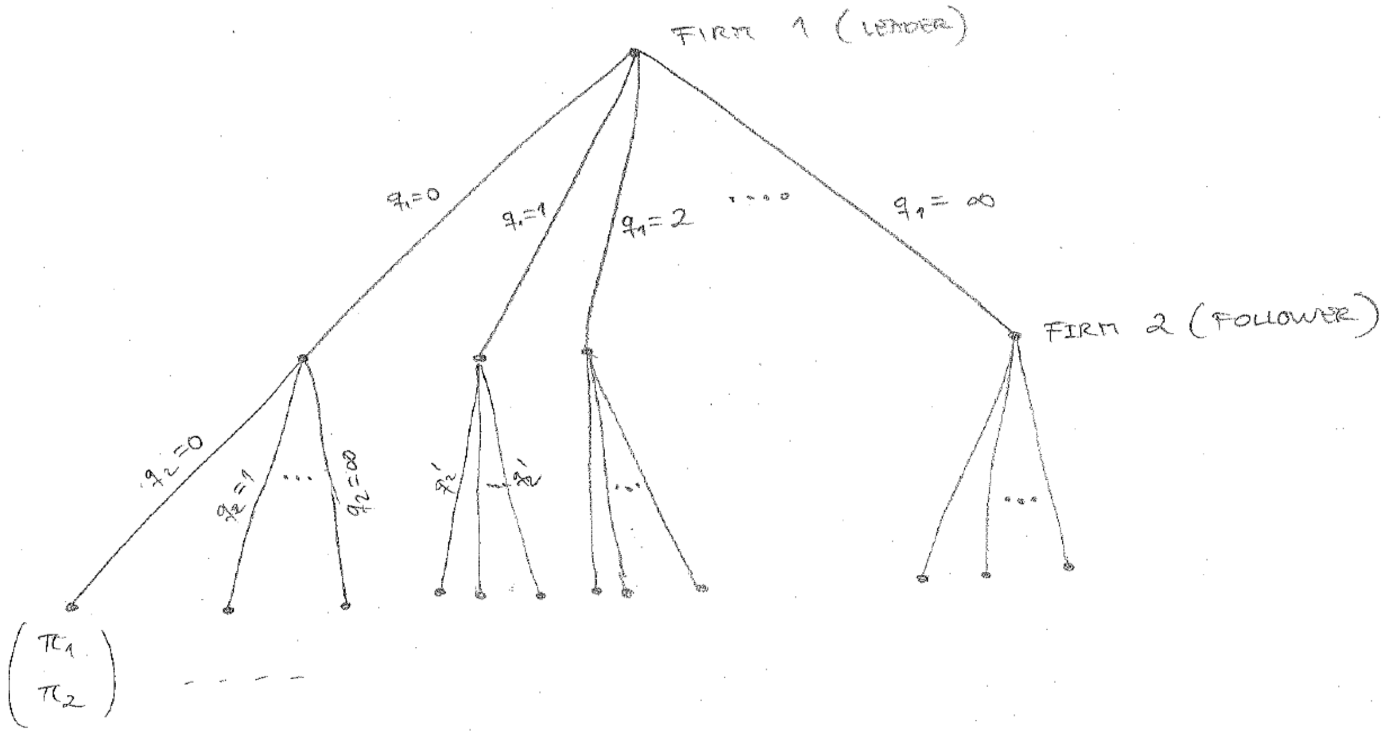

Stackelberg game of sequential quantity competition

Firm 1 is the leader, Firm 2 is the follower. Demand is given by

$$p(q_1, q_2) = 100 - q_1 - q_2$$

and marginal costs are \$10. Operating by backwards induction, we first solve the follower's profit maximization problem

$$\pi_2(q_1, q_2) = [100 - q_1 - q_2]\, q_2 - 10q_2$$

taking FOCs we obtain the BRF2,

$$q_2(q_1) = 45 - \frac{q_1}{2}$$

INTUITIVELY, THIS REPRESENTS THE FOLLOWER'S OPTIMAL ACTION AT THE SMALLEST PROPER SUBGAME (THAT INITIATED AFTER FIRM 1 CHOOSES AN OUTPUT LEVEL, q_1).

Von Stackelberg model

- Now, the leader inserts firm 2's BRF into her own profit function, since she knows how firm 2 will react to firm 1's production decision during the first stage of the game. Hence,

LEADER'S PROFITS

$$\pi_1(q_1, q_2) = \left[100 - q_1 - \underbrace{\left(45 - \frac{q_1}{2}\right)}_{q_2(q_1)}\right] q_1 - 10q_1$$

$$= \frac{1}{2}(90 - q_1)q_1 = \frac{1}{2}(90q_1 - q_1^2)$$

- Taking FOCs with respect to q_1, we obtain

$$\frac{90}{2} - \frac{2q_1}{2} = 0 \iff 90 = 2q_1 \iff q_1^* = 45$$

- Plugging this result into the follower's BRF (BRF2), we find

$$q_2(\underset{45}{\cancel{22.5}}) = 45 - \frac{\overset{45}{\cancel{22.5}}}{2} = 22.5$$

SPNE

Von Stackelberg model

SPNE OF STACKELBERG GAME IS, HOWEVER, MORE GENERAL:

- FIRM 1 CHOOSES OUTPUT $q_1^* = 45$
- FIRM 2 RESPONDS TO q_1 OUTPUT FROM FIRM 1 BY PRODUCING:

$$q_2(q_1) = \underbrace{45 - \frac{q_1}{2}}_{\text{MORE GENERAL THAN } q_2 = 22.5} \qquad (BRF_2)$$

- GRAPHICALLY, BRF_2 REPRESENTS FIRM 2'S BEST RESPONSE TO ANY PRODUCTION OF FIRM 1, q_1, THAT INITIATES ANY SUBGAME (IN WHICH FIRM 2 CHOOSES OUTPUT).

Von Stackelberg model

- For practice, you can check that this same exercise played simultaneously (a la Cournot), leads to

$$q_1^* = q_2^* = 30$$

Summary

- Game theory
  - Simultaneous decisions → NE
  - Sequential decisions → Some NE may not occur if people are rational

- ***Theorem***

Every finite extensive form game has a SPE.

- ***Kuhn's Theorem***

Every sequential game with perfect information has a Nash equilibrium

Outline
1 Repeated Games
2 The Iterated Prisoners' Dilemma

Why we need a new model of repeated games?

- Consider the prisoners' dilemma, in reality, many crooks do not squeal, how do we explain this?
- Consider the Cournot duopoly, we showed that cartels were unstable, but in real-life, many countries need to make (or enforce) anti-collusion laws. How do we explain this?
- In real-life, decisions may not be made once only, but we make decisions based on what we perceive about the future.
- In the prisoners' dilemma, crooks will not squeal because they are afraid of future retaliation. For cartels, they sustain the collusion by making promises (or threats) about the future.
- Inspired by these observations, we consider situations in which players interact repeatedly.

Stage game

- If a player only needs to make a *single* decision, he is playing an stage game.
- After the stage game is played, the players again find themselves facing the same situation, i.e., the stage game is repeated.
- Taken one stage at a time, the only sensible strategy is to use the Nash equilibrium strategy for each stage game.
- However, if the game is viewed as a whole, the strategy set becomes much richer:
  - players may condition their behavior on the past actions of their opponents, or
  - make threats about what they will do in the future, or
  - collusion.

Repeated Games

– Repeated Games
– Recall Prisoner's Dilemma
– Unique NE, both Confess

Example: Iterated Prisoners' Dilemma

- Consider the following prisoners' dilemma game with cooperation (C) and defection (D):

	C	D
C	3,3	0,5
D	5,0	1,1

- Let say the game is repeated just once so there are two stages. We solve this like any dynamic game by backward induction.
- In the final stage, there is no future interaction, so the payoff to be gained is at this final stage. We choose the best response of playing D. So (D, D) is the NE of this subgame.
- Consider the first stage (the subgame is the whole game). Since payoff is fixed for the final stage, the payoff for the entire game is:

	C	D
C	4,4	1,6
D	6,1	2,2

Example: Iterated Prisoners' Dilemma

- Note that the pure-strategy set for each player in the entire game is $\boldsymbol{S} = \{CC, CD, DC, DD\}$.
- But because we are only interested in a subgame perfect NE, we only consider two strategies: $\{CD, DD\}$ (since the last stage is fixed).
- Analyzing the above game (previous payoff table), the NE of the entire game is (DD, DD). So the subgame perfect NE for the whole game is to play D in both stages.
- Note that the player cannot induce cooperation:
  - in the first stage by promising to cooperate in the 2nd stage (since they won't);
  - in the first stage by threatening to defect in the 2nd stage since this is what happens anyway.

Infinite Iterated Prisoners' Dilemma

If the length of the game is infinite, we need the following strategy:

A stationary strategy is one in which the rule of choosing an action is the same in every stage. Note that this does not imply that the action chosen in each stage will be the same.

Examples of stationary strategy are:

- Play C in every stage.
- Play D in every stage.
- Play C if the other player has never played D and play D otherwise.

Infinite Iterated Prisoners' Dilemma

A strategy is called a trigger strategy when a change of behavior is triggered by a single defection.

Example of trigger strategy

- Consider a trigger strategy $s_G =$"Start by cooperating and continue to cooperate until the other player defects, then defect forever after".
- If both players adopt s_G, $\pi_i(s_G, s_G) = \sum_{t=0}^{\infty} 3\delta^t = \frac{3}{1-\delta}$.
- But is (s_G, s_G) a Nash equilibrium?

Infinite Iterated Prisoners' Dilemma

- The payoff for a stationary strategy is the "*infinite sum*" of the payoffs achieved at each stage. Let $r_i(t)$ be the payoff for player i in stage t. The total payoff is $\sum_{t=0}^{\infty} r_i(t)$.
- Unfortunately there is a problem. If both players choose $s_C =$"Play C in every stage", then: $\pi_i(s_C, s_C) = \sum_{t=0}^{\infty} 3 = \infty$.
- If one chooses $s_D =$"Play D in every stage" and other chooses s_C, then: $\pi_1(s_D, s_C) = \pi_2(s_C, s_D) = \sum_{t=0}^{\infty} 5 = \infty$.
- Introduce a discount factor δ $(0 < \delta < 1)$ so the total payoff is: $\sum_{t=0}^{\infty} \delta^t r_i(t)$.
- One can use δ to represent (a) inflation; (b) uncertainty of whether the game will continue, or (c) combination of these.
- Applying, $\pi_i(s_C, s_C) = \sum_{t=0}^{\infty} 3\delta^t = \frac{3}{1-\delta}$.
  $\pi_1(s_D, s_C) = \pi_2(s_C, s_D) = \sum_{t=0}^{\infty} 5\delta^t = \frac{5}{1-\delta}$.

Infinite Iterated Prisoners' Dilemma

Is (s_G, s_G) a Nash Equilibrium?

- Let's do an informal analysis (formal analysis follows).
- Assume both players are restricted to a pure-strategy set $\boldsymbol{S} = \{s_G, s_C, s_D\}$.
- Suppose player 1 decides to use s_C instead, payoff is: $\pi_1(s_C, s_G) = \pi_2(s_C, s_G) = \frac{3}{1-\delta}$. Same result applies if player 2 adopts s_C, so this will not be better off than (s_G, s_G).
- Assume player 1 adopts s_D, the sequence is:

	$t =$	0	1	2	3	4	5	$\cdots$
player 1	s_D	D	D	D	D	D	D	$\cdots$
player 2	s_G	C	D	D	D	D	D	$\cdots$

  For player 1: $\pi_1(s_D, s_G) = 5 + \delta + \delta^2 + \cdots = 5 + \frac{\delta}{1-\delta}$.
- Player 1 cannot do better by switching from s_G to s_D if $\frac{3}{1-\delta} \geq 5 + \frac{\delta}{1-\delta}$. The inequality is satisfied if $\delta \geq 1/2$. So (s_G, s_G) is a NE if $\delta \geq 1/2$.

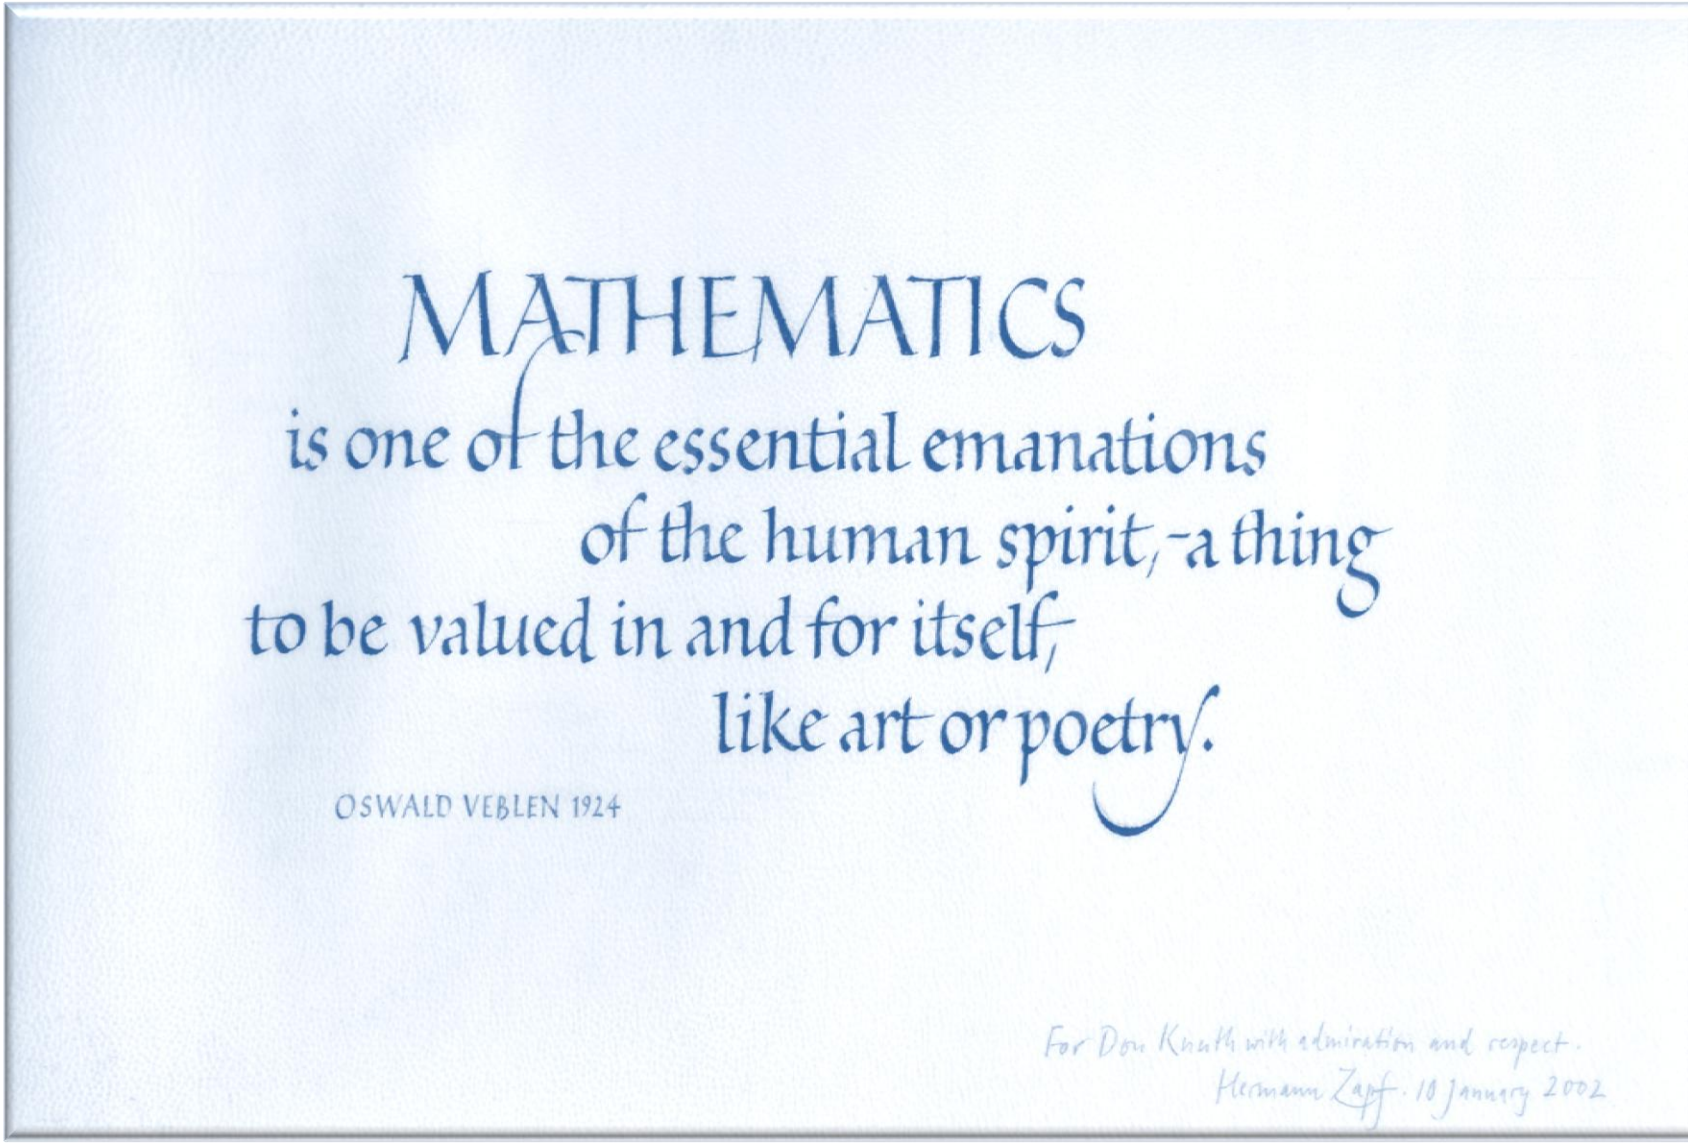
MATHEMATICS
is one of the essential emanations
of the human spirit,–a thing
to be valued in and for itself,
like art or poetry.
OSWALD VEBLEN 1924
For Don Knuth with admiration and respect.
Hermann Zapf. 10 January 2002