Chapter 6: Subgame Perfect Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Sequential Move Game

- Why not just solve them using NE?
  - We would find too many NEs
  - Some NEs would be "sequentially irrational," as we define in this chapter.
- We then need players to behave rationally:
  - Every player maximizes her payoff when called to move (on a node or on an information set)…
    - given her position on the game tree, and
    - given her available information.
- Strategy profiles where every player is sequentially rational will be called Subgame Perfect Equilibria (SPEs).
- We will consider games with:
  - Discrete and continuous strategies.
  - When players observe their rivals' previous actions, and when they do not.
- Before we start, we need to specify some "game tree rules."

Tree Rules

1. **Every node is the successor of the initial node.**
   - Figure 1 illustrates this property.
   - Game tree in the left satisfies it, since there is only one initial node.
   - The game tree in the right panel violates it.
   - If a modeler wanted to indicate that two different players act simultaneously…
     - Then the figure 6.1a at the bottom is the correct way to represent this.
     - Player 1 selects between B or C, then player 2 chooses *c* or *d* without observing player 1's choice (as they are simultaneous).
   - If, instead, the modeler considers the same player in two initial nodes (right panel in Figure 6.1), there must be a mistake because Player 1 cannot choose simultaneously between A and B (on the top of the figure), and between C and D (on the bottom), as if he had multiple personalities!
     - In this case, player 1's decision should happen at the subsequent stage.

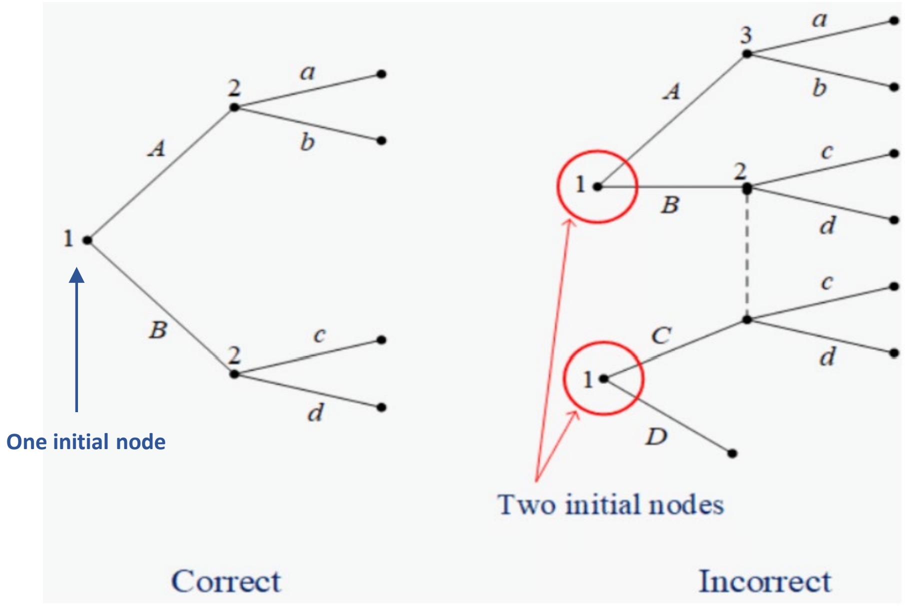

Figure 6.1. Games with one or two initial nodes.

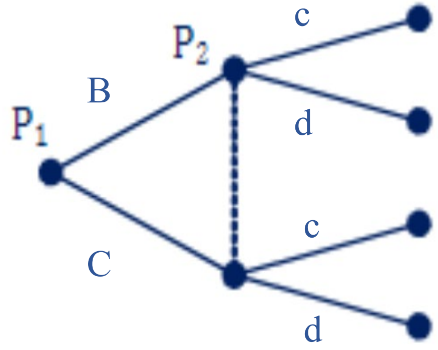

Figure 6.1a: two different players acting simultaneously

Tree Rules

2. **Every node has exactly one immediate predecessor; except the initial node that has no predecessor.**
    - The left panel in Figure 6.1 satisfies this property.
    - But Figure 6.2 depicts a game tree that violates it.
    - If a tree has nodes with more than one predecessor, we could run into misunderstandings:
        - In Figure 6.2, we do not know if player 4 is called to move after player 2 chose b or because player 3 chose c

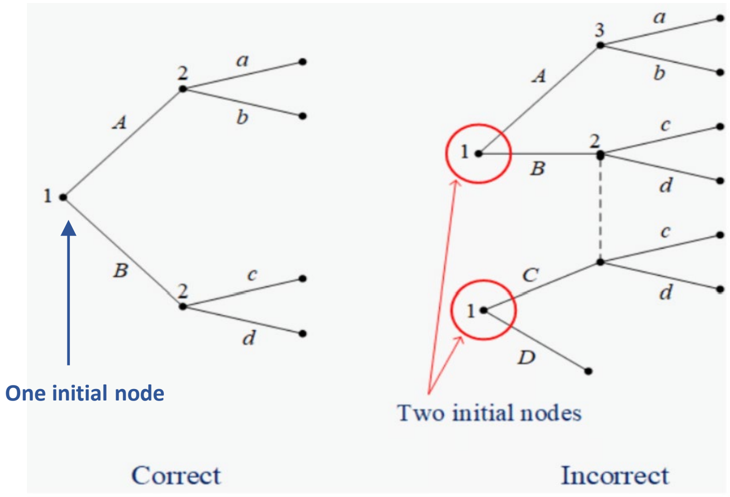

Figure 6.1. Games with one or two initial nodes.

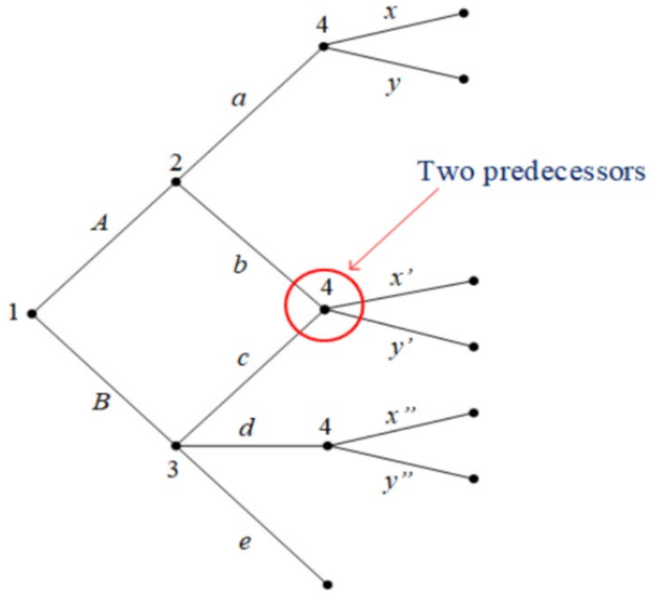

Figure 6.2: Game with two predecessors

Tree Rules

No, you must be referring to a different action. Otherwise collapse everything under the same name.

3. **Multiple branches extending from the same node have different action labels.**
   - Left panel in Figure 6.3 satisfies this property
   - The right panel violates it because label A is on two branches for player 1.
   - If the modeler seeks to represent that:
     - Player 1 chooses A in both the top and middle branches that stem from the initial node, then these two branches should be collapsed into a single branch.
     - If, instead, the modeler tries to represent that player 1 has two different actions in each of these branches, then they should be labelled differently to avoid misunderstandings.

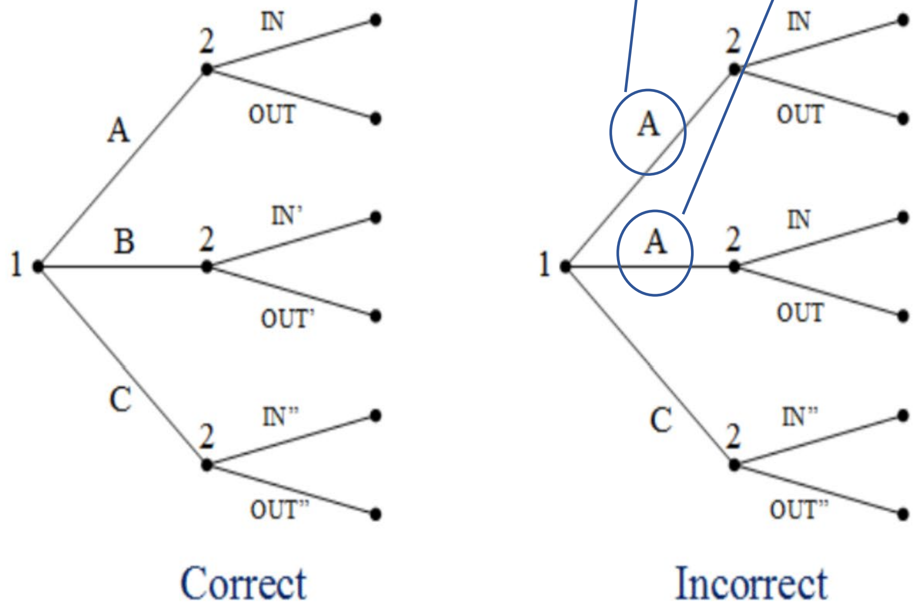

Figure 6.3. Games with same/different action labels.

Tree Rules

4. **Each information set contains nodes for only one player.**
    - The left panel in Figure 6.4 satisfies this property.
    - The right panel violates it.
    - Intuitively, if player 2 is called to move at the top node of the game tree in the right panel, she can infer that player 1 selected A, leaving her no uncertainty about player 1's choices in the previous stage.
        - A similar argument applies to player 3: if she is called to move at the bottom node, she knows that player 1 selected B.
    - Therefore, the information set connecting the top and bottom nodes is incorrect (unnecessary).

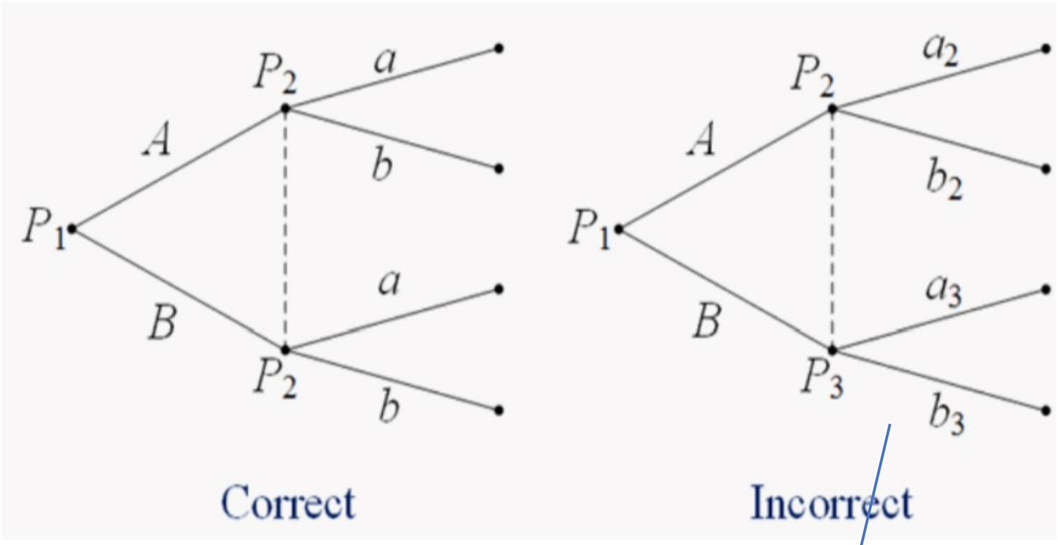

Figure 6.4. Information sets with one or more players.

Incorrect, otherwise P3 knows he is called on to move after player 1 chose B.
(no uncertainty)

Tree Rules

5. Same number of branches and labels.

- *1st requirement.* All nodes in a given information set have the same number of branches stemming from them (i.e., the same number of immediate successors)
  - Figure 6.5a depicts a game tree that violates this:
    - Player 2 would be able to infer whether player 1 chose $Invest$ or $Not\ Invest$ by just observing the number of strategies.
  - Drawing the information set is unnecessary in this setting!
- *2nd requirement.* The labels in the branches must coincide across all nodes connected by the same information set.
  - Figure 6.5b depicts a game tree that violates this:
    - Player 2 would be able to infer that:
    - Player 1 must have chosen $Invest$ if she must select between A and B, but…
    - Player 1 must have chosen $Not\ Invest$ is she must select between A and C.

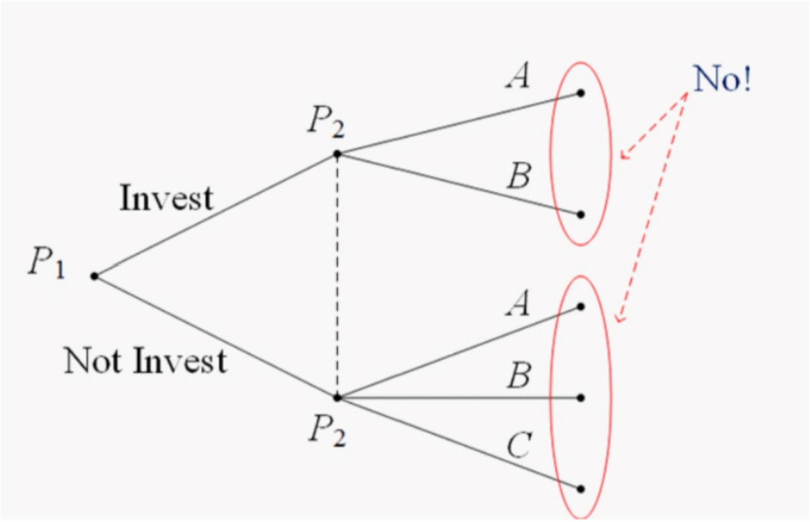

Figure: 6.5a. Different number of branches

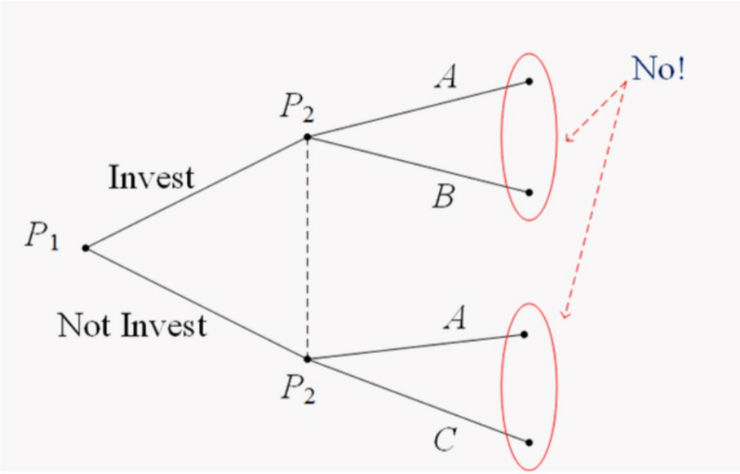

Figure 6.5b. Different immediate successors.

Actions vs. Strategies

- Graphically, an **action** denoted as a_i or $s_i(h_i)$, is just a specific branch that player i chooses when it is her turn to move.
  - Note that the notation is to emphasize one action among those available in information set $h_i, A_i(h_i)$.
- A **strategy** is the list of all the branches that she would choose along the game tree, both in nodes that are reached in equilibrium and those that are not reached.
  - Complete contingent plan.
- Definition. **Pure Strategy.**
  - In a sequential-move game, a pure strategy for player i is a mapping

$$s_i\colon\ H_i \ \to \ A_i$$

  that assigns an action $s_i(h_i) \ \in \ A_i(h_i)$ at information set $h_i \in H_i$, where $A_i(h_i)$ denotes the set of available actions at information set h_i.

Actions and Strategies

- Importantly, the above definition of pure strategy applies to *every* information set, $h_i \in H_i$, describing how player i behaves once she reaches each information set h_i.
- In contrast, an action only describes how player i behaves when reaching a specific information set h_i.
- Example (Figure 6.3):
  - In denotes player 2's action at the top node of the game tree.
  - (In, Out', In'') denotes her strategy, indicating that she responds with:
    - In after player 1 chooses A,
    - Out' after player 1 selects B, and
    - In'' after player1 chooses C).

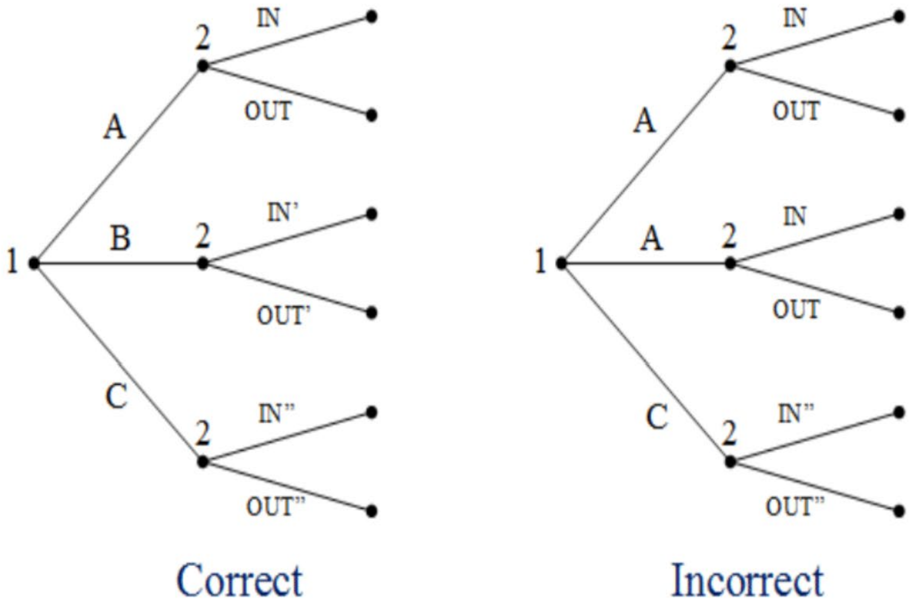

Figure 6.3. Games with same/different action labels.

Why don't we just find the Nash Equilibria of the Game Tree?

- Example. **Entry Game.**
  - Consider the Entry game in Figure 6.6.
  - A firm (potential entrant) chooses whether to enter into an industry, where an incumbent firm operates as a monopolist, or stay out of the market.
  - If the potential entrant stays out (at the bottom of the figure):
    - The incumbent remains a monopolist, earning $10 while the potential entrant earns zero.
  - However, if the entrant joins the industry (top of the figure), the incumbent observes this decision and responds either:
    - Accommodating the entry (e.g., setting moderate prices) which yields a payoff of $4 for each firm, or
    - Fighting it (e.g., starting a price war against the entrant) leading to a payoff of −$2 for each firm.

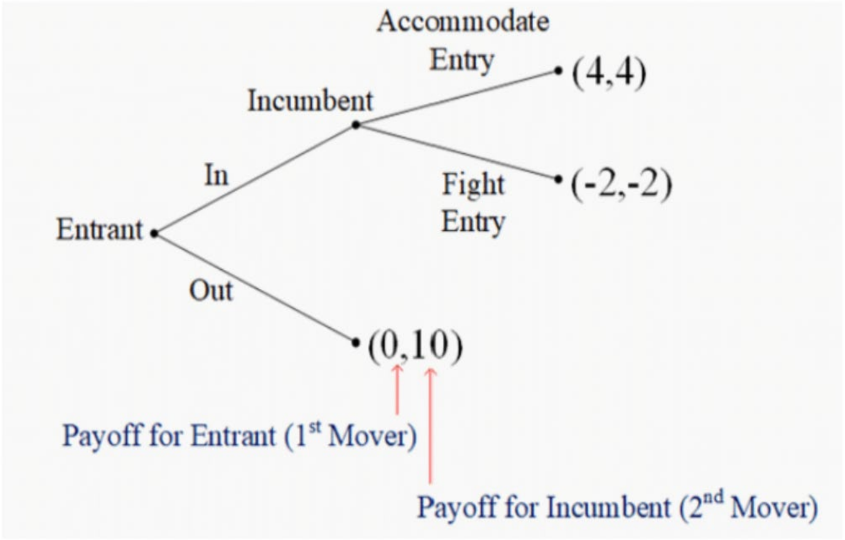

Figure 6.6. The entry game.

Entry Game (contd.)

- To find the NEs of this game tree, we first need to represent the game in matrix form.

		Potential Entrant	
		In	Out
Incumbent	Accommodate	<u>4</u>,<u>4</u>	<u>10</u>,0
	Fight Entry	-2,-2	<u>10</u>,<u>0</u>

Matrix 6.1a. Finding NEs in the Entry Game

- There are two psNE for this game:

$$NE = \{(Accommodate, In) \quad and \quad (Fight, Out)\}$$

Example 6.1. Entry Game (contd.)

- Specifically, for the incumbent, we find that:
  - $BR_{inc}(In) = Acc$ because $4 > -2$
  - $BR_{inc}(Out) = \{Acc, Fight\}$ because both yield a profit of 10
- For the potential entrant,
  - $BR_{ent}(Acc) = In$ because $4 > 0$
  - $BR_{ent}(Fight) = Out$ because $0 > -2$
- Therefore:
  - In the first NE, entry occurs and the incumbent responds accommodating.
  - In the second NE, the entrant does not enter because it believes that the incumbent will start a price war.

Example 6.1. Entry Game (contd.)

- But is this belief credible?
  - NO! The entrant beliefs about the incumbent's decision to $Fight$ after he enters are not rational (in a sequential way):
    - Once the entrant is in, the best thing that the incumbent can do it to accommodate $(4 > -2)$.
    - Hence, the incumbent would never have incentives to start a price war after entry has already occurred.
  - Therefore, among the two NEs:
    - Only $(Accommodate, In)$ is sequentially rational,
    - $(Fight, Out)$ I s *not* sequentially rational.

Introducing a new solution concept

- In the Entry Game, we found NEs that were sequentially irrational.
- Can we find only those equilibria that are sequentially rational? Yes!
- We can guarantee sequential rationality by using a new solution concept, **Subgame Perfect Equilibrium (SPE).**

- For this solution concept, we first need to define what we mean by a "subgame."
- Definition. **Subgame:**
- A subgame is a tree structure defined by a node and all its successors.
  - This definition means that:
    - if nodes a and b are connected with an information set (so player *i* does not know whether he is at node a or b), then both nodes must be part of the *same* subgame.
    - Examples in next slide.

Subgames - Example

- Graphically, a subgame can be identified by drawing a rectangle or a circle around a section of the game tree without “breaking” any information set.
- Figure 6.7a depicts the Entry Game, identifying only two subgames:
    1. that initiated when the incumbent is called on to move; and
    2. the game as a whole.

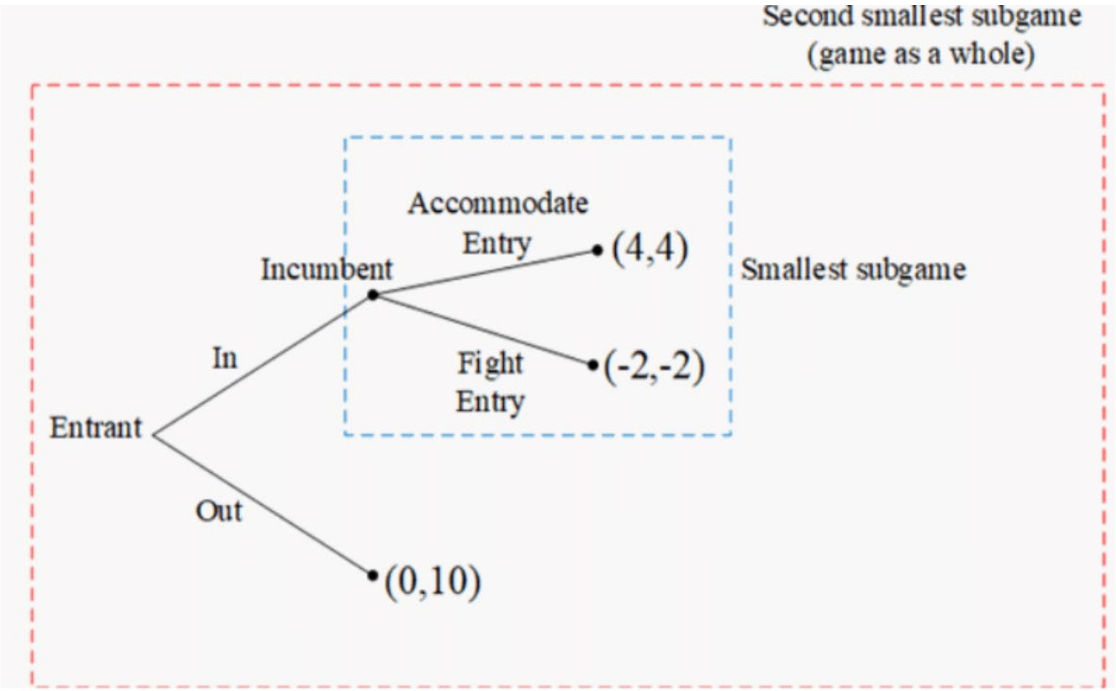

Figure: 6.7a. Subgames in the entry game

Subgames – Example

- There are 4 subgames in Figure 6.7b:
  1. After player 2 is called to move after player 1 chooses Up
  2. After player 2 is called to move after player 1 chooses $Down$
  3. After player 3 is called to move, which happens when player 2 responds to $Down$ with C
  4. The game as a whole

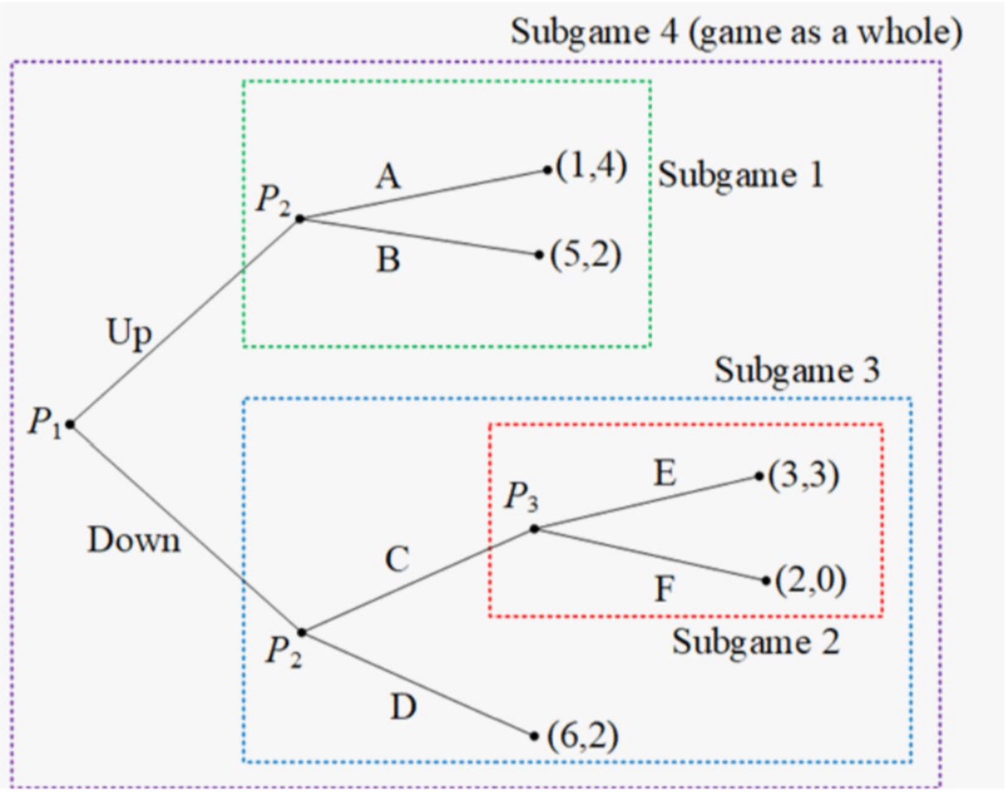

Figure 6.7b. Four subgames in a game tree.

What if the Game Tree has Information Sets?

- Figure 6.8 depicts a game tree with an information set.
  - The smallest subgame must include player 2's information set (otherwise we would be breaking it!)
- The presence of information sets reduces the number of subgames.
- In Figure 6.8, player 2 does not observe whether player 1 selected a or b.
  - This entails that player 2, when choosing whether to respond with c or d, operates "as if" player 1 was selecting a or b at the same time.
  - This is equivalent to a setting where players interact in a *simultaneous-move game*.

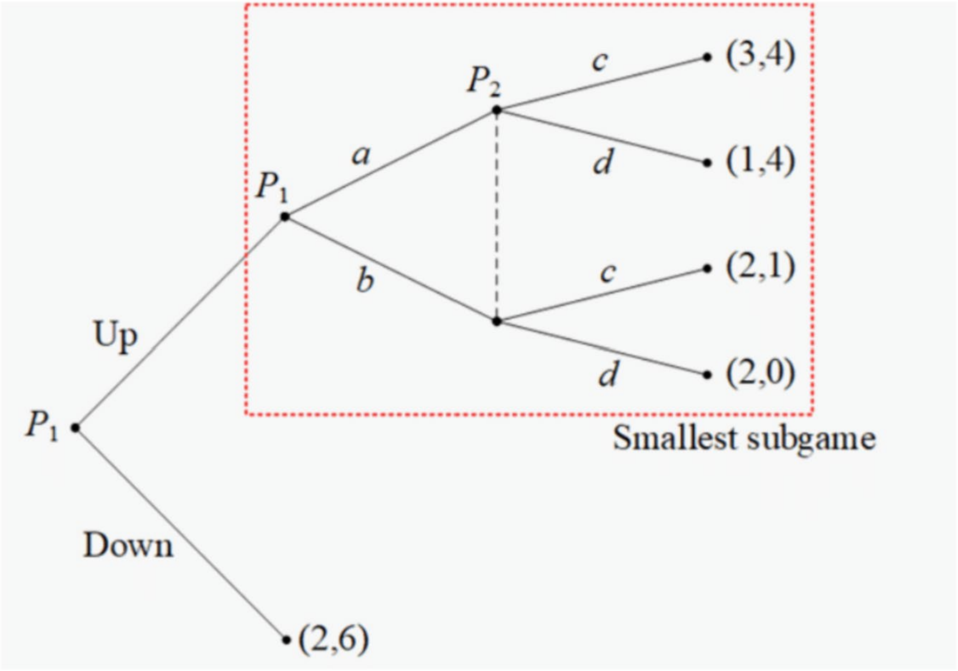

Figure: 6.8. Finding subgames in a game with an information set

Subgame Perfect Equilibrium

- Definition. **Subgame Perfect Equilibrium (SPE).**
- A strategy profile (s_i^*, s_{-i}^*) is a Subgame Perfect Equilibrium if it specifies a NE in each subgame.

- To find SPEs in a sequential-move game, we just need to apply the notion of *backward induction*.

Tool 6.1. Applying backward induction

1. Go to the farthest right side of the game tree (where the game ends), and focus on the last mover.
2. Find the strategy that yields the highest payoff for the last mover.
3. Shade the branch that you found to yield the highest payoff for the last mover.
4. Go to the next-to-last mover and, following the response of the last mover that you found in step 3, find the strategy that maximizes her payoff.
5. Shade the branch that you found to yield the highest payoff for the next-to-last mover.
6. Repeat steps 4-5 for the player acting before the previous-to-the-last mover, and then for each player acting before her, until you reach the first mover at the root of the game.

Example 6.2. Applying Backward Induction – Entry Game

Before we start, recall this game has two subgames.

1st step: Applying backward induction, we first focus on the last subgame (closest to the terminal nodes).

- This corresponds to the last mover, the incumbent.

2nd step: Comparing its payoff from accommodating entry, 4, and starting a price war, -2, we find that its best response to entry is to accommodate.

We shade the corresponding branch in Figure 6.9 to keep in mind the optimal response of the incumbent in this subgame.

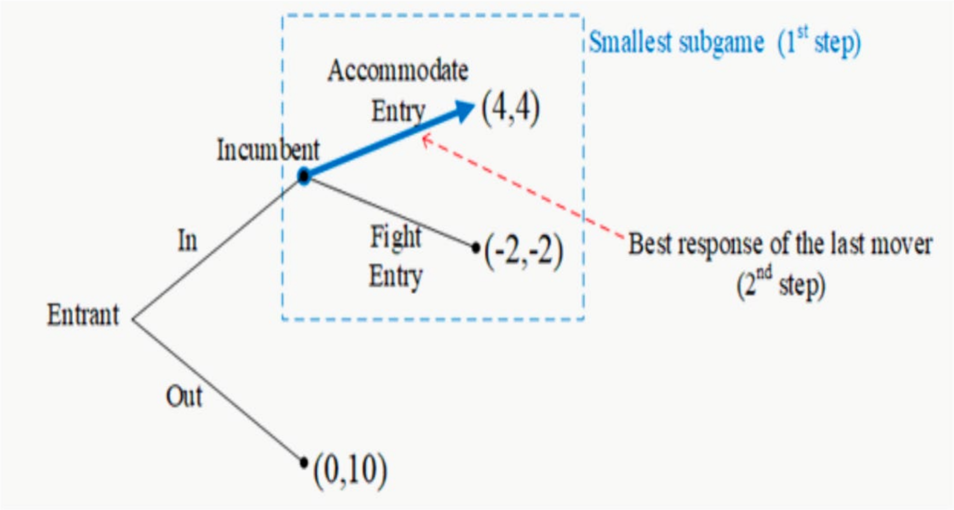

Figure: 6.9. Applying Backward Induction in the entry game –Last Move

Example 6.2. Applying Backward Induction – Entry Game

3rd step: We move to the first mover, the entrant, who anticipates that, if it enters, the incumbent will *accommodate*.

- This means that the entrant expects that, upon entry, the game will proceed through the shaded branch in Figure 6.10 yielding a payoff of 4 from entering.
- If, instead, the entrant stays out, its payoff is only 0. As a consequence, the optimal strategy for the entrant is to enter.

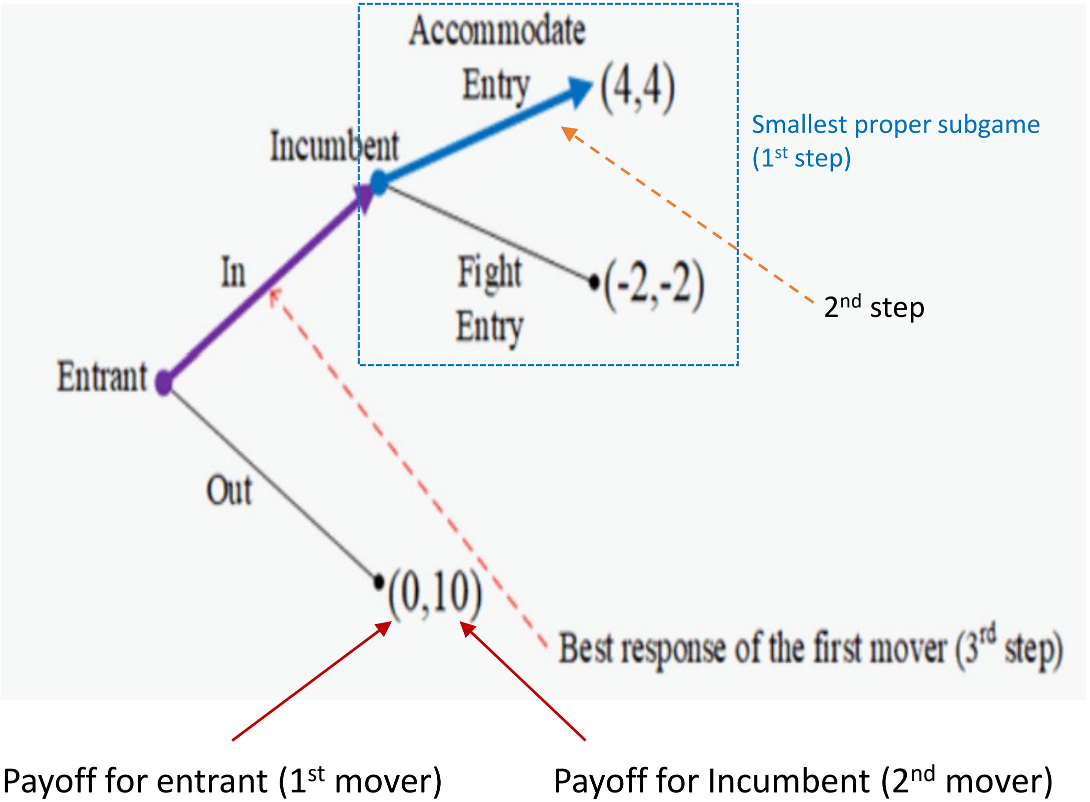

Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying Backward Induction – Entry Game

- Hence:

$$SPE = \{(In, Accommodate)\}$$

- Among the *two psNE* we found, i.e., $(In, Accommodate)$ and $(Out, Fight)$,
  - only the former is sequentially rational.

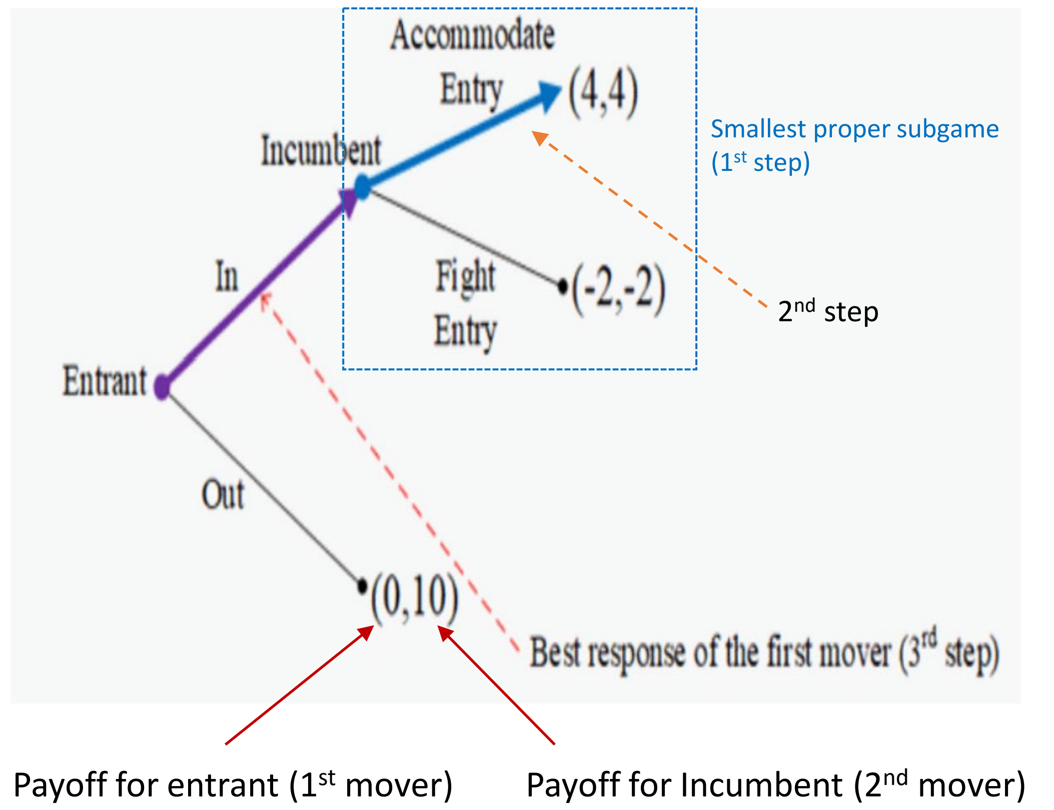

Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying Backward Induction – Entry Game

- Figure 6.10 shades the branches that players choose in equilibrium, known as the *"equilibrium path".*
- The equilibrium path of play is a visual tool to understand how players behave in equilibrium:
    - From the initial node to one of the terminal nodes in the tree,
    - but does not coincide with the SPE in more involved games.

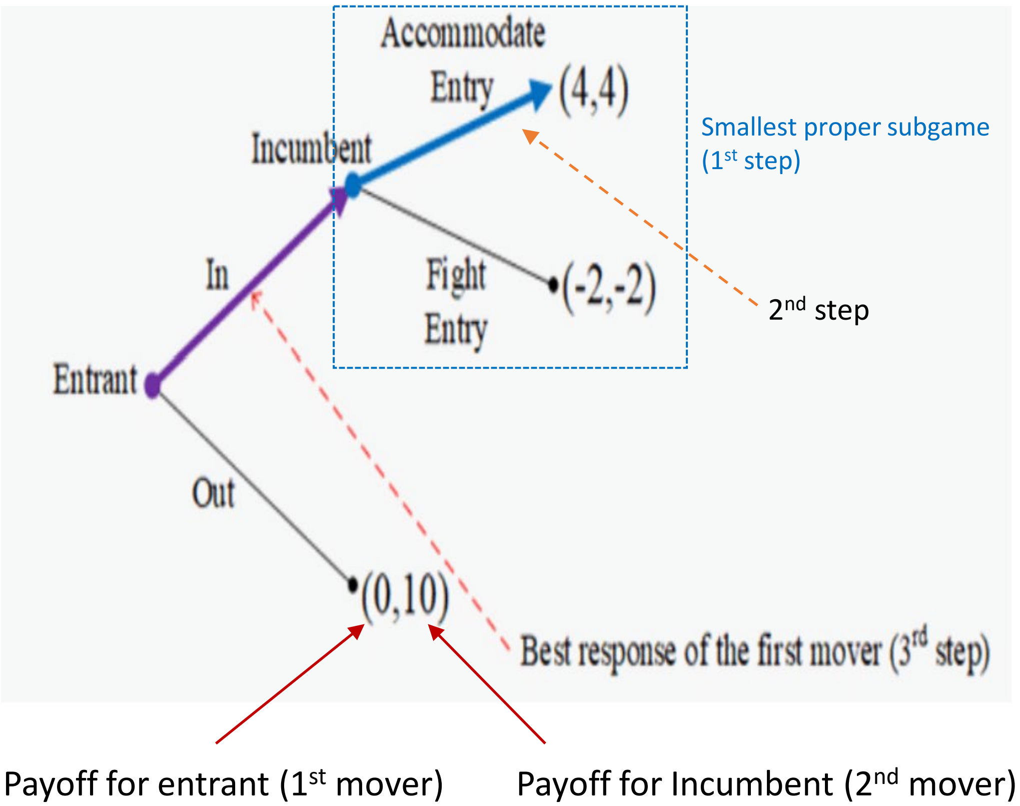

Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying backward induction – Entry Game

- Recall that $(In, Accommodate)$ was also one of the NEs in the Entry game, as shown in Example 6.1.
- Then Examples 6.1 and 6.2 illustrate that:
  - Every SPE must be a NE, since $(In, Accommodate)$ is both a SPE and one of the NEs in the Entry game, but…
  - The converse ins not necessarily true.
  - More formally, for a strategy profile s^*,

$$s^* \text{ is a SPE} \Rightarrow s^* \text{ is a NE}$$
$$\nLeftarrow$$

- Alternatively, the set of strategy profiles that can be supported as SPEs of a game is a subset of those strategy profiles that can be sustained as NE, as depicted in Figure 6.11.

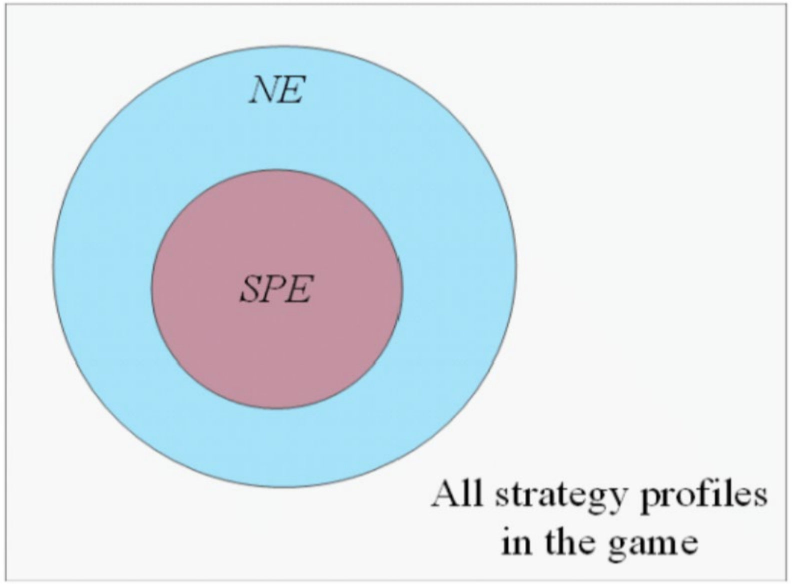

Figure 6.11. SPEs as a subset of NEs.

Example 6.3. Applying Backward Induction in the Modified Entry Game

- Consider the modified version of the Entry Game depicted in Figure 6.12.
- The top part of the game tree coincides with that in the original Entry game.
- However, if the entrant chooses to stay out of the industry, the incumbent can now respond investing in a new technology or not, with associated payoffs $(0,12)$ and $(0,10)$, respectively.

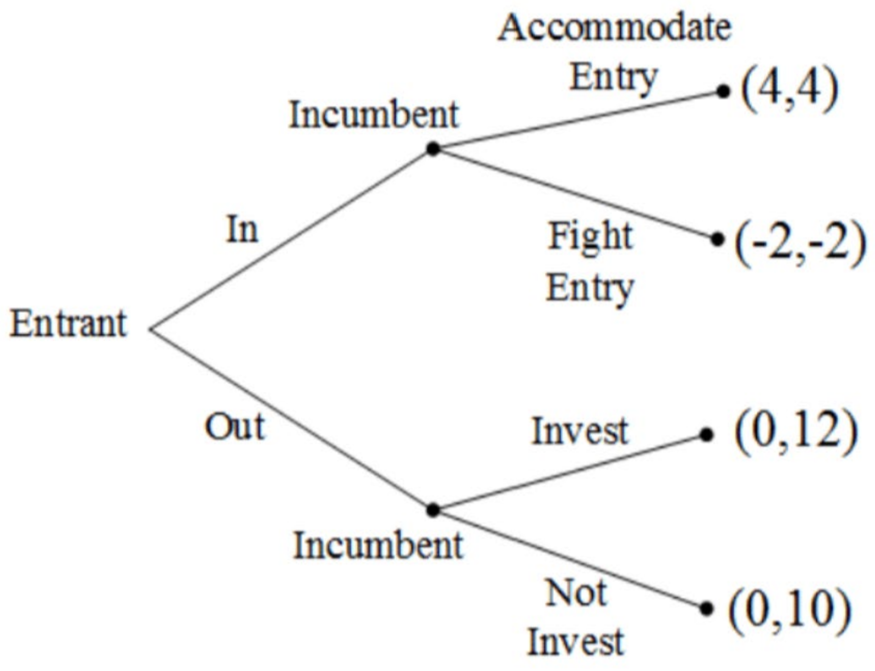

Figure 6.12. Modified entry game.

Example 6.3. Applying Backward Induction in the Modified Entry Game

Three subgames:

1. one initiated after the entrant joins the industry
2. another initiated after the entrant remains out and
3. the game as a whole

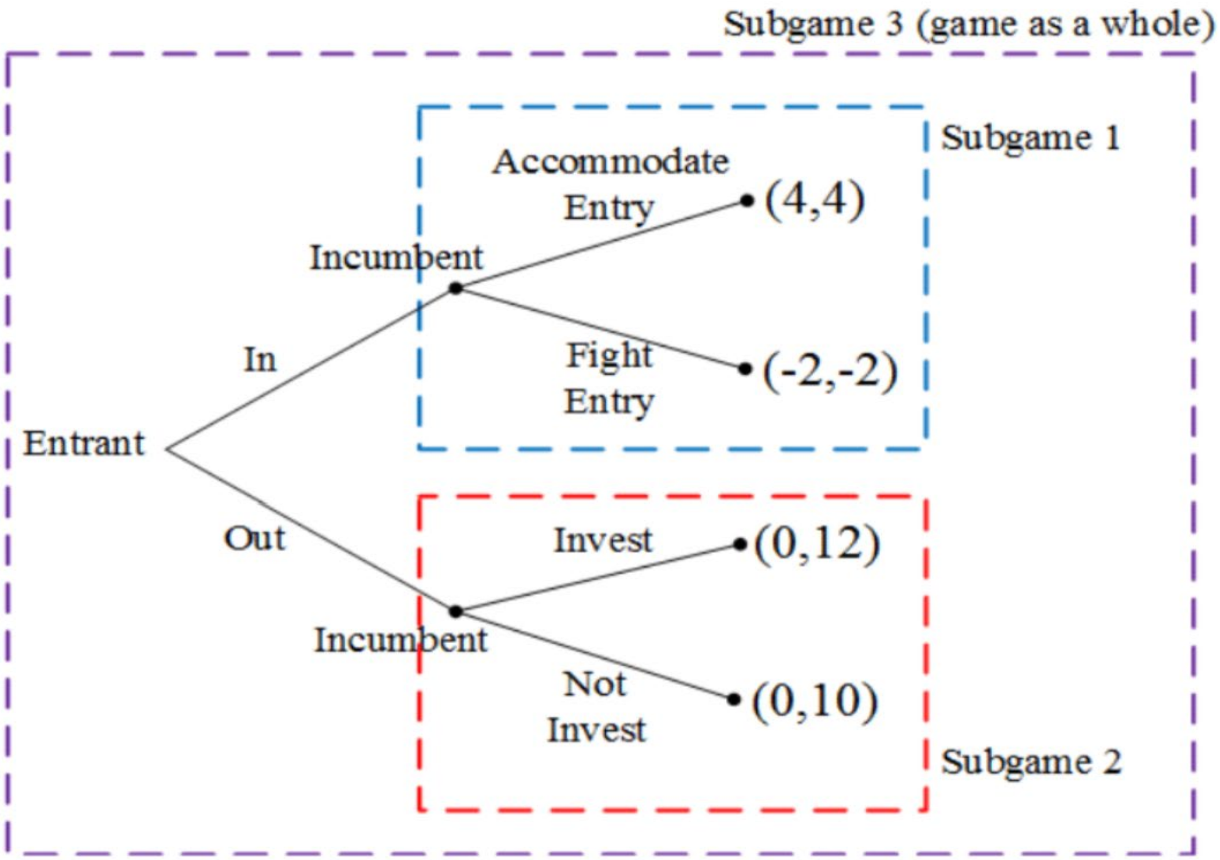

Figure 6.13. Modified entry game - Subgames.

Example 6.3. Applying Backward Induction in the Modified Entry Game

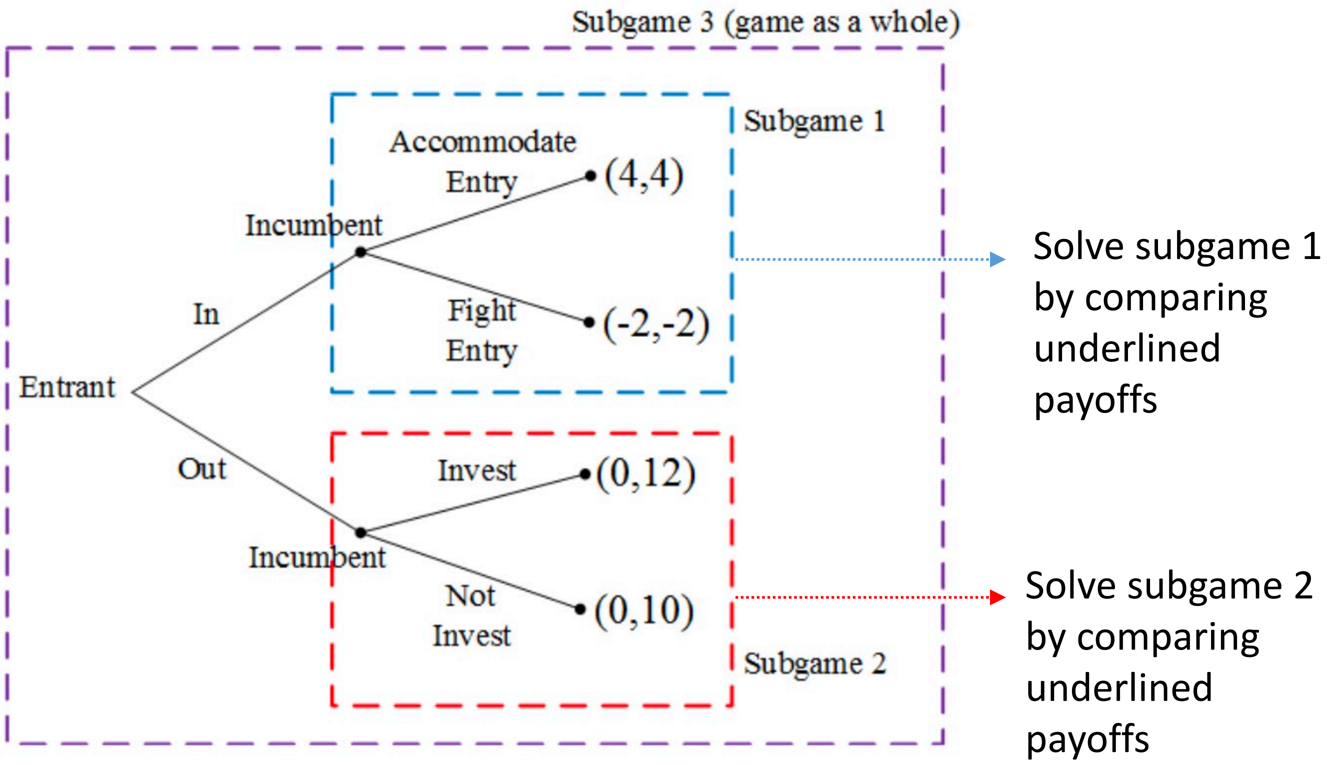

Figure 6.13. Modified entry game - Subgames.

- Anticipating Acc upon entry and $Invest$ upon no entry, the entrant can, in the initial node, compare its payoff from entry, 4, and from no entry 0, thus choosing to enter.
- Therefore:

$$SPE = \{(In, Accommodate/Invest)\}$$

which denotes that the entrant chooses In, and the incumbent responds with Acc after In, but with $Invest$ after Out.

Equilibrium Path vs. SPE

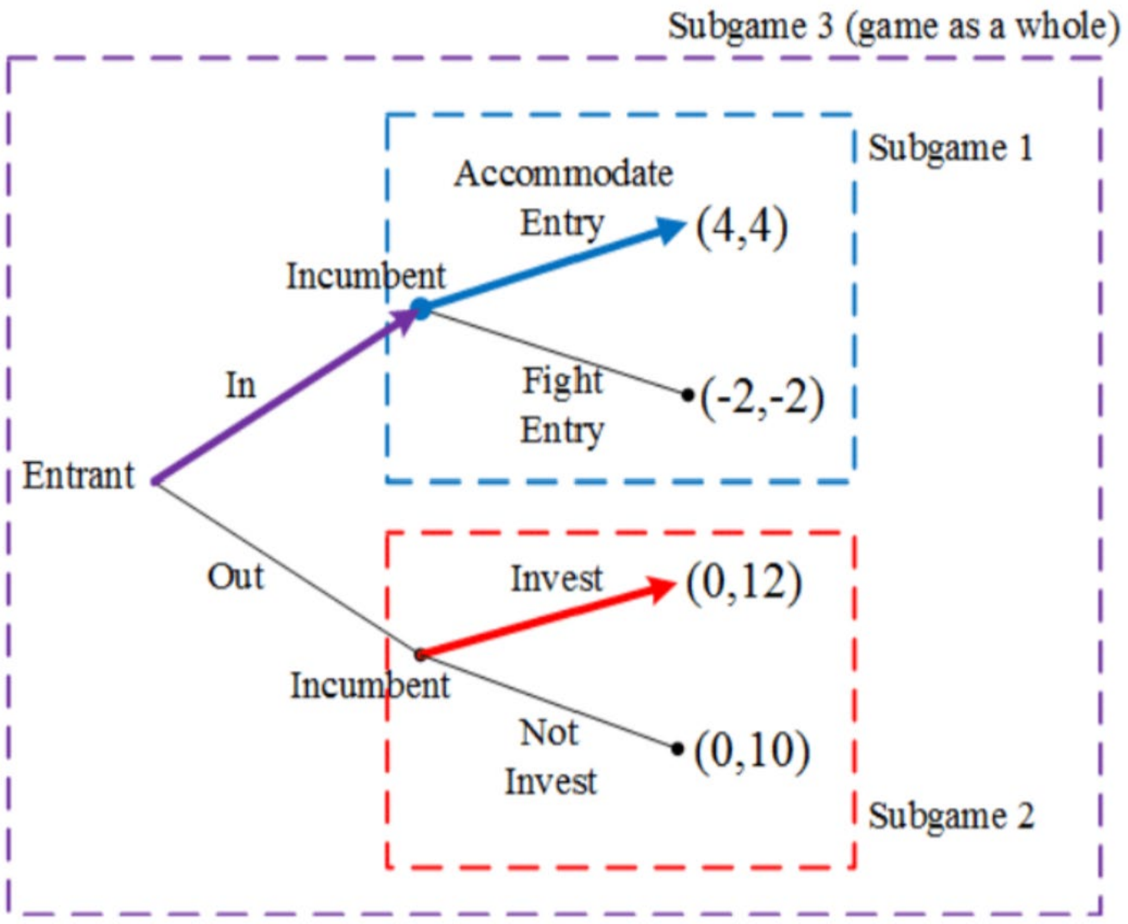

Figure: 6.14 Equilibrium Path vs. SPE

- Figure 6.14 shades the branches that players choose in equilibrium.
- The equilibrium path, coincides with that in the original Entry game.
- The SPE, however, is more intricate because it specifies the incumbent's equilibrium behavior both in the node that it reaches in equilibrium (when the entrant joins the market) and in the node that the incumbent does not even reach in equilibrium (when the entrant stays out).
- Thus, when describing the SPE of a sequential-move game, it must specify the equilibrium behavior for every player *at every node* where she is called to move, even in nodes that may not be reached in equilibrium.

Finding SPEs in Game Trees with Information Sets

- Consider Figure 6.15.
- If firms 1 chooses Up, this firm gets to play again, choosing between A and B.
- Firm 2 is then asked to respond, but without seeing whether firm 1 chose A or B.
- Firm 2's uncertainty is graphically represented by the dotted line connecting the end of the branches that it doesn't distinguish, A and B.
- This dotted line is formally known as an "information set" for Firm 2, because this firm doesn't know which of these two actions was chosen by Firm 1.

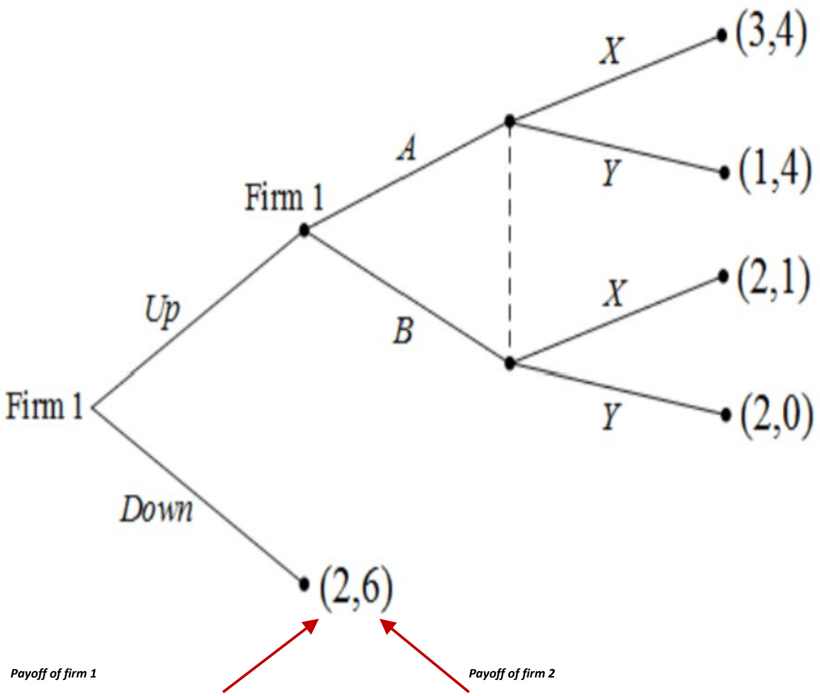

Figure 6.15. A more involved game tree.

Finding SPEs in Game Trees with Information Sets

- Before applying backward induction to this game, a usual trick is to find all the subgames (i.e., circling the portions of the tree that do not break any information set)

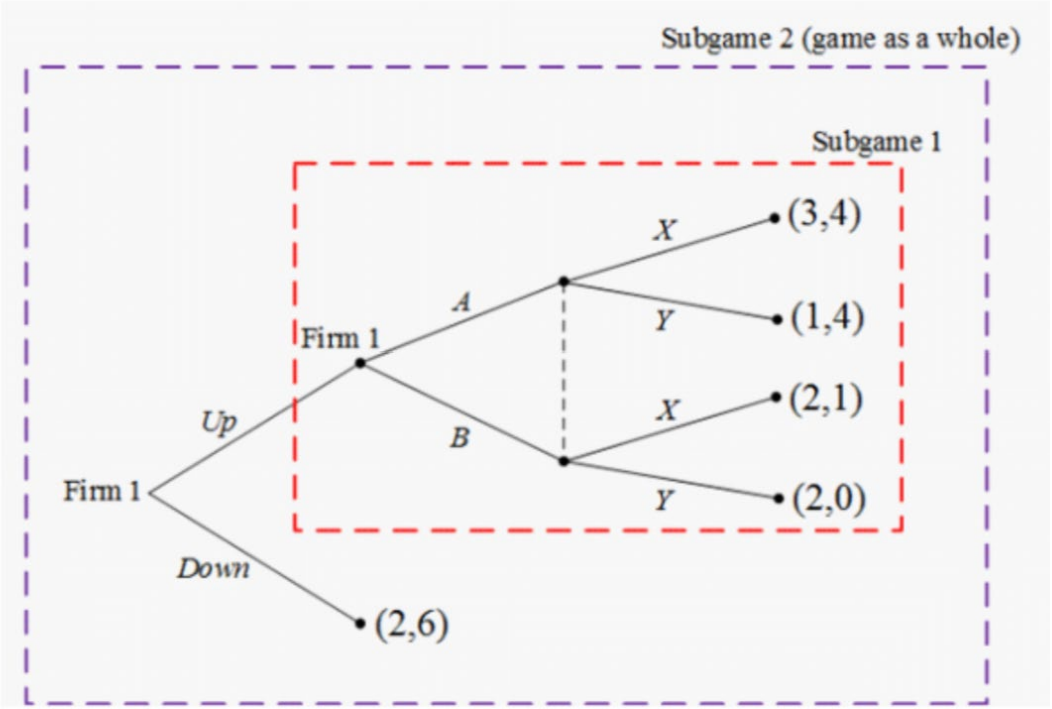

Figure: 6.15(a) Finding subgames

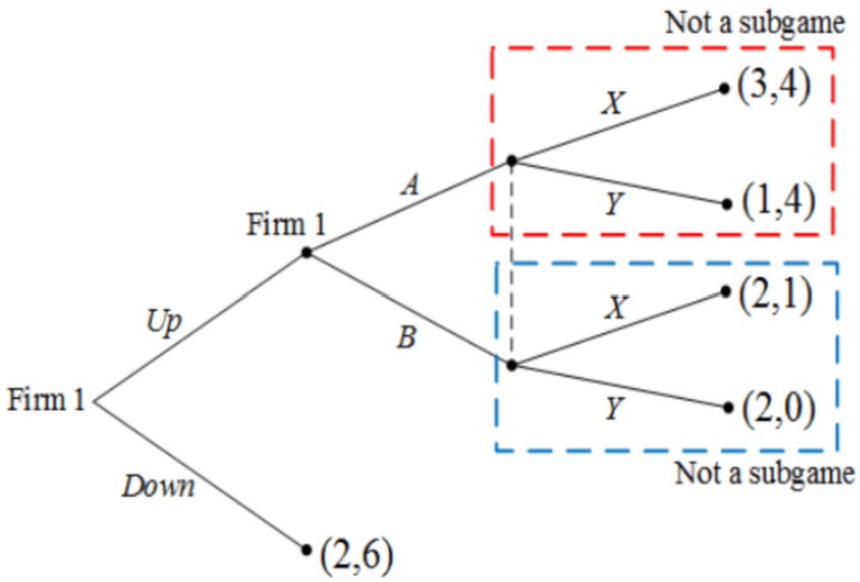

Figure: 6.15(b) Not subgames

Finding SPEs in Game Trees with Information Sets

- **Subgame 1.**
  - Firm 2 does not observe which action firm 1 chose (either A or B). Therefore, subgame 1 can be represented using Matrix 6.2.
  - Matrix 6.3 underlines the best response payoffs.
  - $NEof\ subgame\ 1 = \{(A, X)\}$

		Firm 2	
		X	Y
Firm 1	A	3,4	1,4
	B	2,1	2,0

Matrix 6.2. Representing Subgame 1 in matrix form

		Firm 2	
		X	Y
Firm 1	A	$\underline{3},\underline{4}$	$1,\underline{4}$
	B	$2,\underline{1}$	$\underline{2},0$

Matrix 6.3. Finding the NE of Subgame 1

Finding SPEs in Game Trees with Information Sets

- **The Game as a whole.**
  - Firm 1 must choose between Up and $Down$, anticipating that if it chooses Up, subgame 1 will start.
  - Firm 1 can anticipate equilibrium behavior in subsequent stages of the game; that is, the NE of subgame 1 is (A, X), with payoff $(3,4)$.
  - Firm 1 can then simplify its decision problem to the tree depicted in Figure 6.16. Therefore, firm 1 only needs to compare the following: if it chooses $Down$, the game is over and its payoff is 2, whereas if it chooses Up, subgame 1 is initiated, obtaining a payoff of 3.
  - Because $3 > 2$, firm 1 prefers to choose Up rather than $Down$, illustrated by the thick arrows on the branch corresponding to Up.
  - The SPE of this game is $\left(Up, (A, X)\right)$.

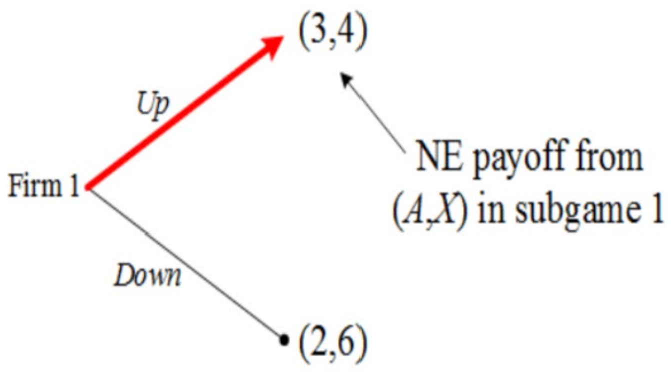

Figure 6.16. The reduced game tree from figure 6.15a.

Evaluating SPE as a Solution Concept

1. Existence? Yes.
    - When we apply backward induction in any game tree, we find that at least one equilibrium exist.
        - This is often known as “Zermelo’s theorem”, after Zermelo (1913) article on chess, and later on extended by Kuhn (1953).
    - Intuitively, this means that every player selects on strategy at every node where she is called to move.
    - If a player is indifferent between two available actions *a* and *b*, for instance, then her best response is to choose either.
        - But every player selects on strategy at every node, yielding at least one SPE strategy profile.

Evaluating SPE as a Solution Concept

2. Uniqueness? No.
   - While the examples in this chapter display game trees that produce a unique SPE, we cannot guarantee that applying backward induction will always induce a unique SPE, as required by this criterion.
   - For instance, if in the Entry game of Example 6.2 the incumbent's payoff from both accommodating and fighting entry was 4, then this firm would be indifferent between Acc and $Fight$, leading to two SPEs:

$$SPE = \{(In, Acc), (In, Fight)\}$$

   - In this case, the SPEs exactly coincide with those that are NEs.
   - Interesting point:
     - If players are not indifferent about any of their actions at any of the nodes (or information sets) where they are called to move…
     - equilibrium behavior in each subgame must specify a *unique strategy for every player in that subgame*,
     - so the SPE must be unique.

Evaluating SPE as a Solution Concept

3. Robust to small payoff perturbations? Yes.
   - If we change the payoff of one of the players by a small amount (e.g. 0.001, but generally, for any ε that approaches zero), backward induction still provides us with the same equilibrium outcome, implying that SPE is robust to small payoff perturbations.
   - This is due the fact that:
     - If a strategy s_i yields a strictly higher payoff than another strategy s_i',
     - it must still yield a strictly higher payoff than s_i' after we apply a small payoff perturbation (remember that ε can be infinitely small),
     - meaning that play i still chooses s_i in the subgame where this action was optimal.

Evaluating SPE as a Solution Concept

4. Socially optimal? No.
   - The application of backward induction does not necessarily produce a socially optimal outcome, which goes in line with our evaluation of IDSDS and NE, which did not yield socially optimal outcomes either.

Application: Stackelberg Game of Sequential Quantity Competition

- Consider two firms produce a homogeneous good facing a linear demand function $p(Q) = 1 - Q$, where $Q = q_1 + q_2$ denotes aggregate output.
- Assume that all firms have a constant marginal cost of production c, where $1 > c \geq 0$.
- Firms interact in a sequential-move game:
  1. In the first stage, firm 1 (the industry leader) chooses its output q_1.
  2. In the second stage, firm 2 (the industry follower) observes that the leader's output q_1 and responds with its own output level q_2.
- Solving the game by backward induction, we start by solving the second stage.

Application: Stackelberg Game of Sequential Quantity Competition

- **Second stage, follower – firm 2.**

Firm 2 maximizes its profit as follows

$$\max_{q_2 \geq 0} \pi_2 = (1 - q_1 - q_2)q_2 - cq_2$$

Differentiating with respect to q_2:

$$1 - 2q_2 - q_1 - c = 0$$

Solving for q_2 yields:

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

which is positive for all $\frac{1-c}{2} - \frac{1}{2}q_1 \geq 0 \Rightarrow q_1 \leq 1 - c$.

Application: Stackelberg Game of Sequential Quantity Competition

Firm 2's BRF:

$$q_2(q_1) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_1 \text{ if } q_1 \leq 1-c \\ 0 \text{ otherwise.} \end{cases}$$

- Intuitively, the follower produces an output of $q_2 = \frac{1-c}{2}$ when the leader is inactive $(q_1 = 0)$, but decreases its output in q_1 at a rate of $\frac{1}{2}$.
- When the leader's output is sufficiently large, $q_1 > 1-c$, the follower responds staying inactive, $q_2 = 0$.
- This BRF coincides with that under the Cournot model of simultaneous quantity competition in chapter 4.

Application: Stackelberg Game of Sequential Quantity Competition

- **First stage, leader – firm 1.**

Firm 1 maximizes its profit as follows

$$\max_{q_1 \geq 0} \pi_1 = [1 - q_1 - q_2(q_1)]q_1 - cq_1$$

which is evaluated at $q_2(q_1)$. Inserting firm 2's BRF, we obtain

$$\max_{q_1 \geq 0} \pi_1 = \left[1 - q_1 - \overbrace{\left(\frac{1-c}{2} - \frac{1}{2}q_1\right)}^{q_2(q_1)}\right]q_1 - cq_1 = \left[\frac{1+c}{2} - \frac{1}{2}q_1\right]q_1 - cq_1$$

Differentiating with respect to q_1:

$$\frac{1+c}{2} - q_1 - c = 0$$

Solving for q_1 yields:

$$q_1^{Seq} = \frac{1-c}{2}$$

Application: Stackelberg Game of Sequential Quantity Competition

- The SPE in this Stackelberg game of quantity competition:

$$SPE = \left(q_1^{Seq}, q_2(q_1)\right) = \left(\frac{1-c}{2}, \frac{1-c}{2} - \frac{1}{2}q_1\right)$$

where firm 2 (the follower) response with its BRF, $q_2(q_1)$, for any output that the leader chooses (both its equilibrium output, $q_1 = q_1^{Seq}$, and any off-the-equilibrium output, $q_1 \neq q_1^{Seq}$).

- In equilibrium, we can claim that the follower observes the leader's output, q_1^*, and inserts it into its BRF, to obtain the follower's equilibrium:

$$q_2^{Seq} = \frac{1-c}{2} - \frac{1}{2}\overbrace{\frac{1-c}{2}}^{q_1^{Seq}} = \frac{1-c}{4}$$

- Importantly, *we do not say* that the SPE of this Stackelberg game with two symmetric firms is $\left(q_1^{Seq}, q_2^{Seq}\right) = \left(\frac{1-c}{2}, \frac{1-c}{4}\right)$. This output vector only describes firms' output along the equilibrium path. In contrast, the SPE must specify each firm's output decision, both in- and off-the-equilibrium path.

Equilibrium Output level in the Cournot and Stackelberg Games

- Recall that the equilibrium output in the Cournot model of simultaneous quantity competition was $q_i^{Sim} = \frac{1-c}{3}$ for every firm i (see section 4.2.1), where we assumed the same indirect demand function $p(Q) = 1 - Q$ and marginal cost c for both firms.
- Therefore,
  - the leader produces more units than when firms chooses their output simultaneously, $q_1^{Seq} = \frac{1-c}{2} > \frac{1-c}{3} = q_1^{Sim}$,
  - while the follower produces fewer units, $q_2^{Seq} = \frac{1-c}{4} < \frac{1-c}{3} = q_2^{Sim}$.

- Intuitively, the leader exercises its “First-mover-advantage” by increasing its output, relative to that under Cournot competition, anticipating that the follower will respond to this increase in q_1 by decreasing its own output q_2.

Equilibrium Output level in the Cournot and Stackelberg Games

- Figure 6.17 depicts the equilibrium output level in the Cournot and Stackelberg games, showing that the output pair moves form the $45° - line$ (where both firms produce the same output) to above this line, which indicates that firm 1 produces more units than firm 2.
- Graphically, firm 1 anticipates firm 2's BRF, $q_2(q_1)$, and chooses the point along this line that yields the highest profit.

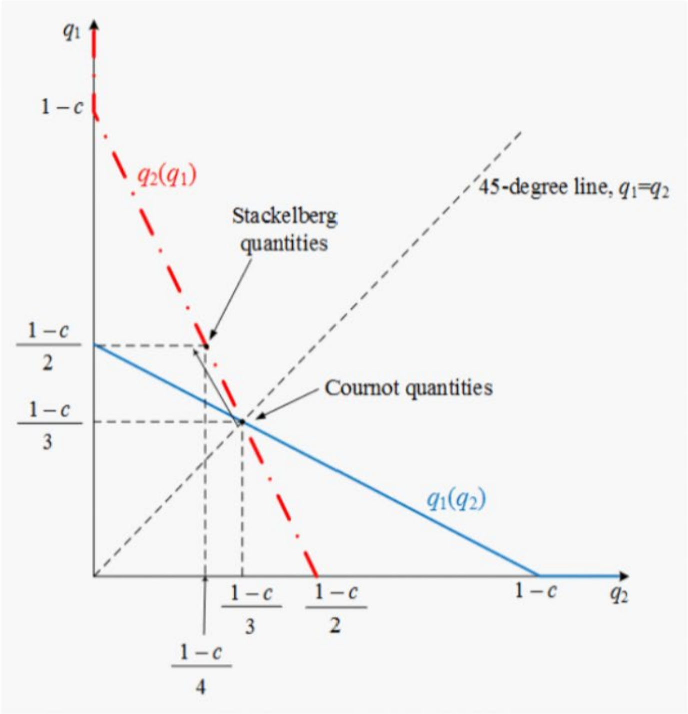

Figure 6.17. Equilibrium quantities in Cournot vs. Stackelberg games.

Equilibrium Profit level in the Cournot and Stackelberg Games

As expected, the leader's profit when firms compete sequentially,

$$\pi_1^{Seq} = \left[1 - q_1^{Seq} - q_2^{Seq}\right] q_1^{Seq} - cq_1^{Seq}$$
$$= \left[1 - \frac{1-c}{2} - \frac{1-c}{4}\right] \frac{1-c}{2} - c\,\frac{1-c}{2}$$
$$= \frac{(1-c)^2}{8},$$

is larger than its profit when they compete simultaneously i.e.,

$$\pi_1^{Seq} = \frac{(1-c)^2}{8} > \pi_1^{Sim} = \frac{(1-c)^2}{9}$$

Equilibrium Profit level in the Cournot and Stackelberg Games

In contrast, the follower's profits when firms compete sequentially,

$$\pi_2^{Seq} = \left[1 - q_1^{Seq} - q_2^{Seq}\right] q_2^{Seq} - c q_2^{Seq}$$
$$= \left[1 - \tfrac{1-c}{2} - \tfrac{1-c}{4}\right] \tfrac{1-c}{4} - c\, \tfrac{1-c}{4}$$
$$= \tfrac{(1-c)^2}{16},$$

is lower than its profits under simultaneous quantity competition, i.e.,

$$\pi_2^{Seq} = \frac{(1-c)^2}{16} < \pi_2^{Sim} = \frac{(1-c)^2}{9}$$

Sequential Public Good Game

- Let us now return to the public good game we examined in section 4.4, but rather than assuming players simultaneously choose their contributions to the public good, we now consider that
  - Player 1 makes her contribution x_1 in the first period, and after observing this contribution, player 2 responds with her contribution x_2 in the second period.
  - Because this is a sequential-move game of complete information, we solve it by backward induction, focusing on player 2 first, and then moving on to player 1.

Sequential Public Good Game

Player 2, follower. Player 2 takes player 1's contribution, x_1, as given, and chooses her contribution $x_2$to solve the utility maximization problem:

$$\max_{x_2} \; u_2(x_1, x_2) = (w_2 - x_2)\sqrt{m(x_1 + x_2)}$$

Differentiating with respect to x_2, we obtain

$$-\sqrt{m(x_1 + x_2)} + \frac{m(w_2 - x_2)}{2\sqrt{m(x_1 + x_2)}} = 0$$

which simplifies to $\frac{m(w_2 - 2x_1 - 3x_2)}{2\sqrt{m(x_1 + x_2)}} = 0$ which holds when $3x_2 = w_2 - 2x_1$

Sequential Public Good Game

Solving for x_2, we find that player 2's BRF is:

$$x_2(x_1) = \begin{cases} \frac{w_2}{3} - \frac{2}{3}x_1 & if\, x_1 < \frac{w_2}{2} \\ 0 & otherwise. \end{cases}$$

This BRF coincides with that found in section 4.4:

- Originating at $x_2 = \frac{w_2}{3}$ when player 1 does not donate to the public good.
- But decreasing at a rate of $^2/_3$ for every dollar that player 1 contributes.

Sequential Public Good Game

Player 1, leader. In the first period, player 1 anticipates player 2's BRF in the subsequent period, $x_2(x_1)$, and inserts it into her utility maximization problem as follows:

$$\max_{x_1 \geq 0} \quad u_1(x_1, x_2) = (w_1 - x_1)\sqrt{m[x_1 + x_2(x_1)]}$$

$$= (w_1 - x_1)\sqrt{m\left[x_1 + \underbrace{\left(\frac{w_2}{3} - \frac{2}{3}x_1\right)}_{x_2(x_1)}\right]}$$

which simplifies to

$$\max_{x_1 \geq 0} \ (w_1 - x_1)\sqrt{m\frac{w_2 + x_1}{3}}$$

Sequential Public Good Game

Player 1, leader. In the first period, player 1 solves

$$\max_{x_1 \geq 0} \ (w_1 - x_1)\sqrt{m\frac{w_2 + x_1}{3}}$$

Differentiating with respect to x_1, we obtain

$$\frac{m(w_1 - x_1)}{2\sqrt{m(w_2 + \ x_1)}} - \frac{\sqrt{m(w_2 + x_1)}}{\sqrt{3}} = 0$$

$$\Rightarrow \quad \frac{m(w_1 - 2w_2 - 3x_1)}{2\sqrt{3}\sqrt{m(w_2 + \ x_1)}} = 0$$

Then solving for x_1, we find player 1's equilibrium contribution in this sequential-move version of the public good game

$$x_1^{Seq} = \frac{w_1 - 2w_2}{3}.$$

Sequential Public Good Game

Therefore, the SPE of the game is:

$$SPNE = \left(x_1^{Seq}, x_2(x_1)\right) = \left(\frac{w_1 - 2w_2}{3}, \frac{w_2}{3} - \frac{2}{3}x_1\right)$$

As a consequence, players contributions evaluated in equilibrium are:

- $x_1^{Seq} = \frac{w_1 - 2w_2}{3}$ for player 1, and
- $x_2^* = x_2\left(x_1^{Seq}\right)$ for player 2, that is

$$x_2\left(x_1^{Seq}\right) = \frac{w_2}{3} - \frac{2}{3}\left(\frac{w_1 - 2w_2}{3}\right) = \frac{7w_2 - 2w_1}{9}$$

Sequential and Simultaneous Public Good Game Comparison

- We find that the leader contributes less to the charity than when players submit donations simultaneously, i.e.,

$$x_1^{Seq} = \frac{w_1 - 2w_2}{3} < \frac{3w_1 - 2w_2}{5} = x_1^{Sim}$$

because this inequality simplifies to $-w_2 < w_1$, which always hold given that wealth levels are positive by assumption.

- The follower, however, contributed more in the sequential- than simultaneous-move game since:

$$x_2^{Seq} = \frac{7w_2 - 2w_1}{9} > \frac{3w_2 - 2w_1}{5} = x_2^{Sim}$$

simplifies to $-w_1 < w_2$, which also holds for all wealth levels.

- Intuitively, the leader anticipates that her decreased contributions will be responded by the follower with an increase in her own donation.
- In other words, the leader exploits her first-move advantage, which in this case means decreasing her contribution to free-ride the follower’s donation.

Application: Ultimatum Bargaining Game

- Bargaining is prevalent in many economic situations where two or more parties negotiate how to divide a certain surplus, such as the federal budget or a salary increase.
- We first examine the shortest bargaining setting, the so-called "ultimatum bargaining game" depicted in Figure 6.18a:
  - Player 1 makes a division of the surplus to player 2, d, where $d \in [0,1]$ can be interpreted as a share of the total surplus.
  - Observing the offer d, player 2 only has the ability to accept it or reject it.
- The arc at the top of the game represents the continuum of offers that player 1 can make to player 2, since $d \in [0,1]$.

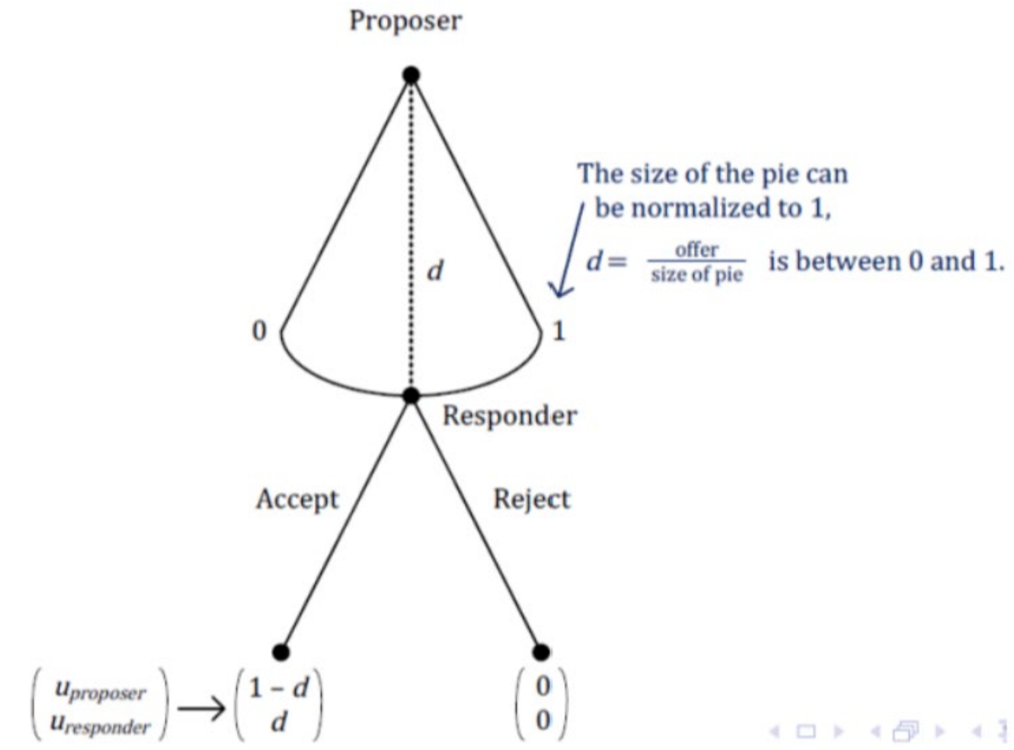

Figure 6.18a. Ultimatum bargaining game.

Application: Ultimatum Bargaining Game

- If player 2 rejects the offer, both players earn a payoff of zero.
- If player 2 accepts the offer:
  - She receives d
  - Player 1 earns the remaining surplus, $1 - d$.
- *Example.*
  - If player 1 offers 20 percent of the surplus to player 2, $d = 0.2$, and player 2 accepts it, player 1 keeps the remaining 80 percent of the surplus, $1 - d = 0.8$.
- This bargaining game is, then, equivalent to a "take-it-or-leave-it" offer, or an ultimatum, from player 1, explaining the game's name.

Application: Ultimatum Bargaining Game

- The smallest subgame that we can identify in Figure 6.18b is that initiated after the responder (player 2) observes the proposer's offer, d.
- The second smallest subgame is the game as a whole.

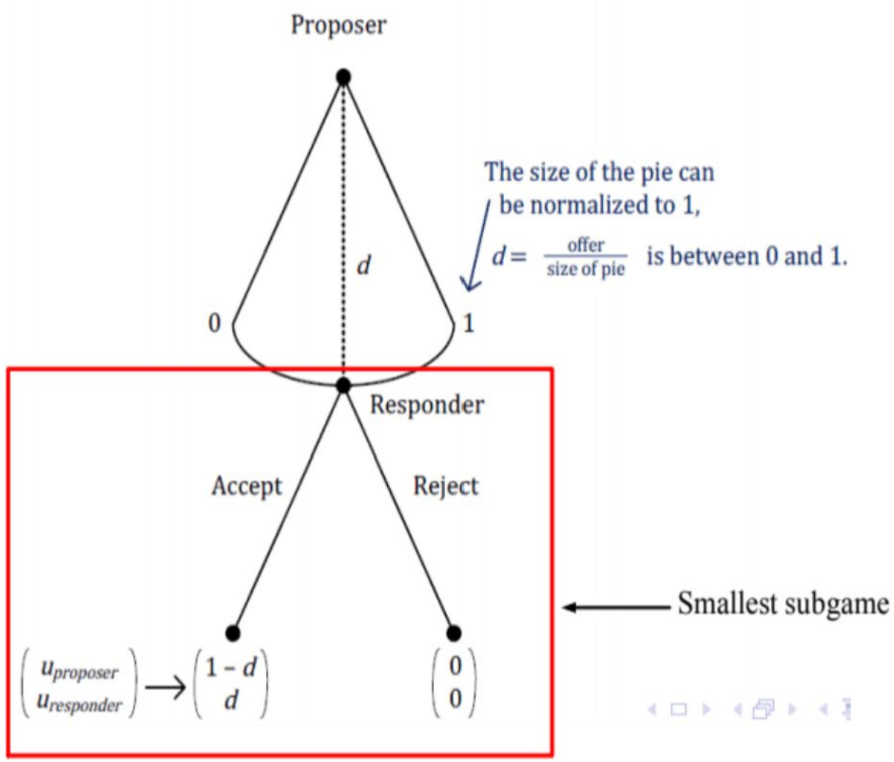

Figure 6.18b. Ultimatum bargaining game - Smallest subgame.

Application: Ultimatum Bargaining Game

- Applying backward induction
  - **Smallest subgame.**
    - Upon receiving an offer d, the responder accepts it if d satisfies $d \geq 0$.
  - **First Stage.**
    - Proposer anticipates the responder's decision rule, $d \geq 0$**,** and makes an offer that maximizes her payoff conditional on that being accepted, i.e., proposer solves

$$\max_{d \geq 0} \; 1 - d$$

    which considers the constraint $d \geq 0$ to induce the responder to accept the offer.
    - Differentiating with respect to d, yields -1 (corner solution).
    - Intuitively, the proposer seeks to reduce d as much as possible (Figure with d on the horizontal axis).

  - Therefore, the proposer reduces d all the way to $d^* = 0$, which still satisfies acceptance, making the responder indifferent between accepting and rejecting the offer.

Application: Ultimatum Bargaining Game

- For simplicity, we assume that the responder accepts offers that make her indifferent.
  - If, instead, she rejects this type of offers, the proposer could offer her an extremely small division of the surplus, $d \rightarrow 0$, still yielding similar equilibrium results as above.
- Therefore, the SPE of the ultimatum game is:

$$SPE = \{d^* = 0, d \geq 0\}$$

which indicates that the proposer makes an offer $d^* = 0$, and the responder accepts any offer d that satisfies $d \geq 0$.

- Equilibrium path:
  - The proposer offers $d^* = 0$
  - The responder accepts that offer, yielding equilibrium payoff $(1 - d, d) = (1,0)$ which implies that the proposer keeps all the surplus.
- Recall that this “equilibrium path” doesn’t describe how the responder reacts to an off-the-equilibrium offer from the proposer, that is, $d \neq 0$.

Application: Two-period alternating offers bargaining game

- Figure 6.19 extends the above bargaining:
  - allowing the responder (player 2) to make a counteroffer if she rejects the division that player 1 offers, d^1.

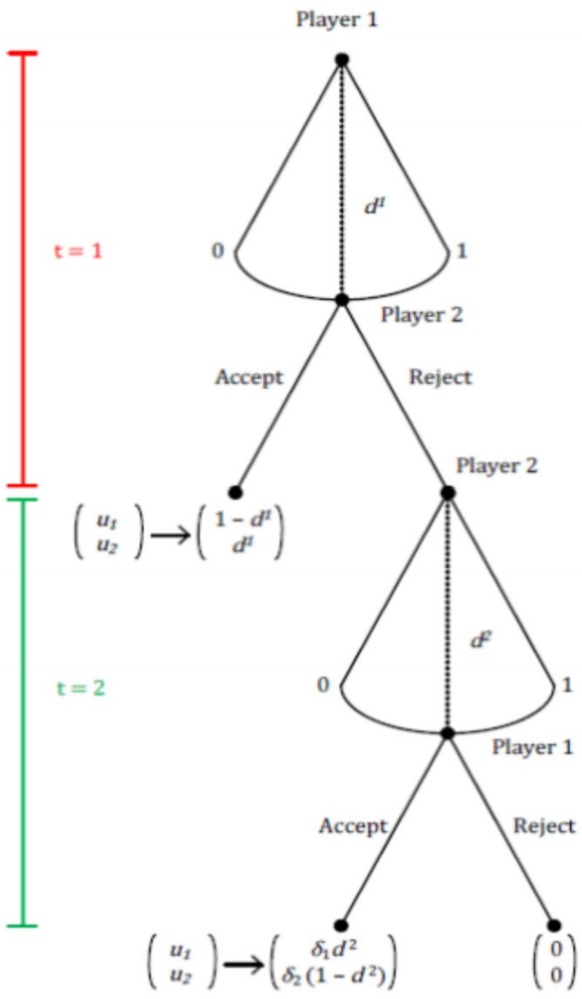

Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- Player 1 makes an offer d^1 to player 2, who accepts or rejects it.
- If player 2 accepts d^1, the game is over, and player 2 earns d^1.
  - Player 1 earns the remaining surplus, $1 - d^1$.
- If player 2 rejects, however, she can make an offer d^2 to player 1 (which we can interpret as a "counteroffer").
- Observing d^2, player 1 responds accepting or rejecting it.

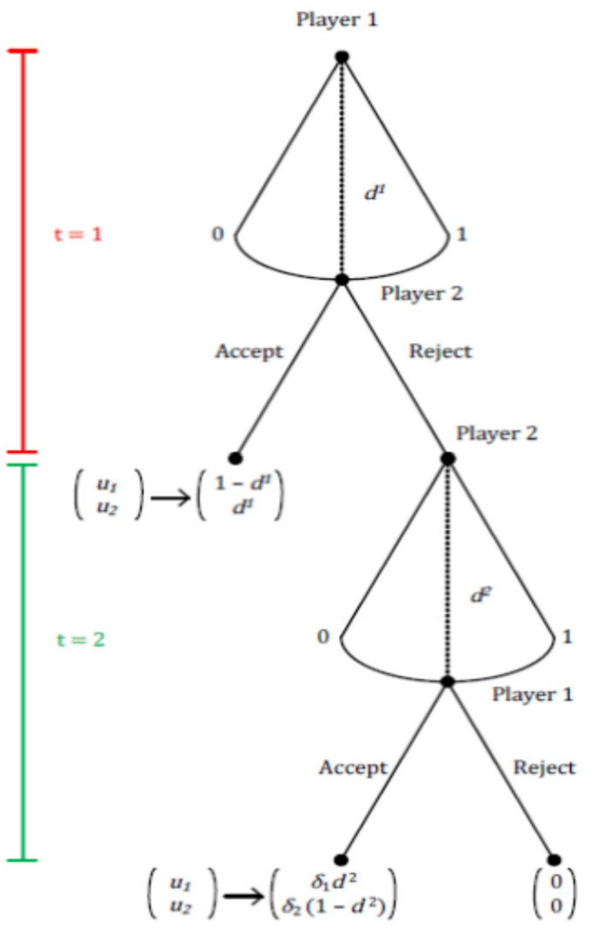

Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- As in previous stages, both players earn zero if the game ends without a division being accepted.
  - Otherwise, player 1 (the responder in the second stage) earns d^2, which has a discounted value of $\delta_1 d^2$ in today's terms.
- Discount factor $\delta_1 \in [0,1]$ indicates the relative importance that player 1 assigns to future payoffs.

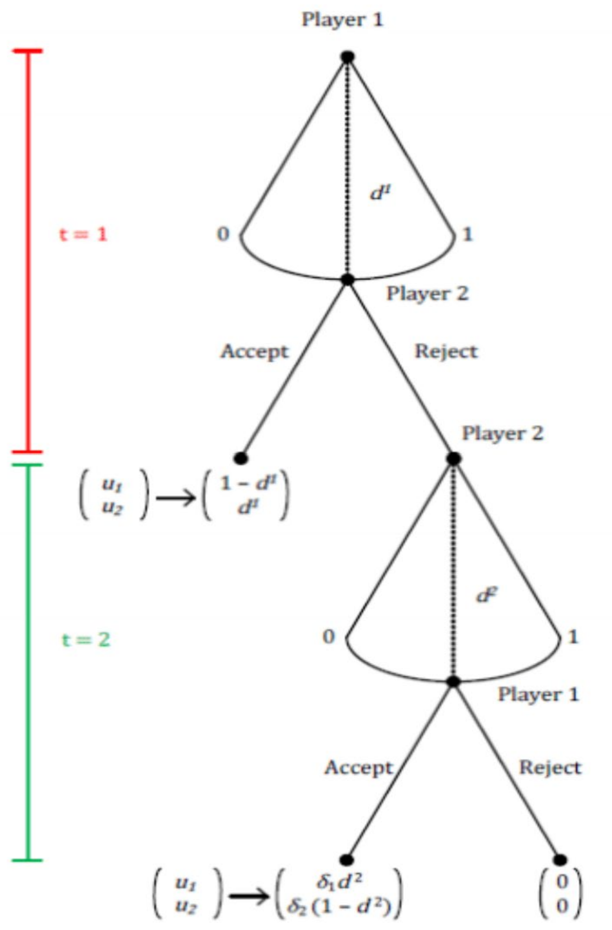

Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- Player 2 earns, in this case, the remaining surplus $1 - d^2$, whose discounted value in today's terms is $\delta_2(1 - d^2)$.
- In this context, we can find four subgames, those initiated after:
  1. Player 1 receives offer d^2
  2. Player 2 rejects offer d^1
  3. Player 2 receives offer d^1, and
  4. The game as a whole

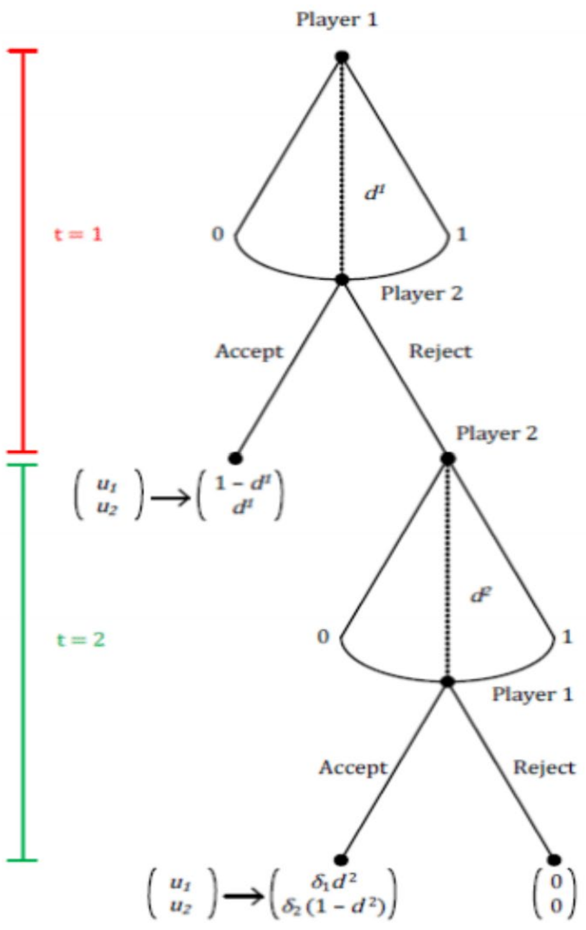

Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

By backward induction, we start with the smallest subgame (subgame (1)).

- **Subgame (1).** At that point of the game, player 1 accepts any offer d^2 from player 2 if and only if $d^2 \geq 0$.
- **Subgame (2).** In subgame (2), player 2 anticipates this best response by player 1, making an offer d^2 that solves

$$\max_{d^2 \geq \mathbf{0}} \quad 1 - d^2$$

which is analogous to the proposer's problem in the ultimatum bargaining game of section 6.4.3.

  - Intuitively, when players reject offers and reach their final period of interaction, they face a strategic setting like the ultimatum bargaining game:
    - Not reaching an agreement at that point will lead to a zero payoff for everyone, as in the ultimatum bargaining game.
  - The proposer (player 2) offers the lowest division that guarantees acceptance from player 1, $d^2 = 0$, earning a payoff of $1 - d^2 = 1 - 0 = 1$.

Application: Two-period alternating offers bargaining game

- **Subgame (3).** Player 2 chooses whether to:
  - Accept the offer from player 1, earning d^1, or
  - Reject it and become the proposer tomorrow.
    - In this setting, she anticipates that her equilibrium payoff in subgame (2) will be 1.
    - Discounted value of 1 today is $\delta_2 \times 1 = \delta_2$.
- Therefore, player 2 accepts offer d^1 if and only if $d^1 \geq \delta_2$.

Application: Two-period alternating offers bargaining game

- **Game as a whole.** Player 1 maximizes

$$\max_{d^1 \geq \delta_2} \; 1 - d^1$$

which yields an equilibrium offer $d^1 = \delta_2$. Therefore, the SPE of the two-period alternating offers bargaining game is

$$SPE = \left\{ \underbrace{(d^1 = \delta_2, d^2 \geq 0)}_{Player\ 1}, \underbrace{(d^2 = 0, d^1 \geq \ \delta_2)}_{Player\ 2} \right\}$$

  - Player 1 offers a division $d^1 = \delta_2$ in the first period, and accepts any offer $d^1 \ \geq 0$ in the second period.
  - Player 2 offers a division $d^2 = 0$ in the first period, and accepts any offer $d^1 \ \geq \delta_2$ in the second period.

Application: Two-period alternating offers bargaining game

- **Game as a whole.**
  - This SPE implies that, while players have two periods to negotiate how to divide the surplus, they reach an agreement in the first period (instantaneously!):
    - Player 1 offers a division $d^1 = \delta_2$, which player 2 accepts since it satisfies $d^2 \geq 0$, and the game is over (see left side of Figure 6.19).
  - In equilibrium,
    - Player 1 earns a payoff $1 - d^1 = 1 - \delta_2$.
    - Player 2 earns a payoff equal to the division she accepted from player 1, $d^1 = \delta_2$.
    - In summary, equilibrium payoffs are $(u_1^*, u_2^*) = (1 - \delta_2, \delta_2)$

Application: Two-period alternating offers bargaining game

Comparative Statics of equilibrium payoffs.

- As player 2 becomes more patient (higher δ_2):
  - Her equilibrium payoff increases while that of player 1 decreases.
  - Player 1 understands that player 2 does not discount future payoffs significantly, meaning that she can reject player 1's offer today, becoming the proposer tomorrow.
- If instead $\delta_2 \rightarrow 0$:
  - Player 1 anticipates that player 2 severely discounts future payoffs, having a stronger preference for today's payoffs.
  - In this case, player 1 can exploit player 2's impatience by offering her a lower division.
- Interestingly, player 1's patience, as captured by his discount factor δ_1, does not affect:
  - The offer she makes to player 2 in equilibrium, $d^2 = 0$, or
  - Players' equilibrium payoffs, $(u_1^*, u_2^*) = (1 - \delta_2,\ \delta_2)$, which are only affected by player 2's patience, δ_2.

Some tricks about solving alternating offer bargaining games

- Equilibrium payoffs in the last stage of the game are $(0,1)$ where the proposer in that stage makes a zero offer, which is accepted, helping the proposer keep the whole surplus.
  - This holds true when the game has only one period of possible negotiations (ultimatum bargaining game), two periods, or more.
- Players then anticipate that, if the game reaches the last period, their equilibrium payoffs will be $(0,1)$ if player 2 is the proposer in that period or $(1,0)$ if player 1 is the proposer.
  - (WLOG, assume that player 2 is the proposer in the last period).
- The proposer in the previous-to-last period (player 1), must then offer a division d^1 that makes player 2 indifferent between:
  - Accepting d^1 today or…
  - Rejecting it to become the proposer tomorrow and earn a payoff of 1, with discounted value $\delta_2 \times 1 = \delta_2$ today.
  - This means that player 1 offers exactly $d^1 = \delta_2$, keeping the remaining surplus $1 - \delta_2$ in equilibrium.

Some tricks about solving alternating offer bargaining games

Applying in longer alternating offer bargaining games

- For instance, if players can negotiate for four periods, starting with player 1 making an offer d^1, we can operate by backward induction as follows:
  1. *Fourth period. P*layer 2 is the proposer, so the equilibrium payoff in that subgame is $(0,1)$.
  2. *Third period.* Player 1 makes an offer, so she must make player 2 indifferent between her offer today or the whole surplus tomorrow, worth δ_2 today. Equilibrium payoffs are, then,
     - $(1-\delta_2, \delta_2)$.
  3. *Second period.* Player 2 is the proposer, so she must make player 1 indifferent between her offer today and payoff $1-\delta_2$ tomorrow (when she becomes the proposer), with discounted value $\delta_1(1-\delta_2)$ today. Therefore, equilibrium payoffs are

$$\left(\delta_1(1-\delta_2), 1-\delta_1(1-\delta_2)\right)$$

where $1-\delta_1\ (1-\delta_2)$ represents the remaining surplus that player 2 earns, after offering $\delta_1(1-\delta_2)$ to player 1.

Some tricks about solving alternating offer bargaining games

4. *First period.* Player 1 is the proposer, and she must make player 2 indifferent between her current offer and payoff $1 - \delta_1\,(1 - \delta_2)$ tomorrow, with discounted value $\delta_2[1 - \delta_1\,(1 - \delta_2)]$ today. As a consequence, equilibrium payoffs are

$$(1 - \delta_2[1 - \delta_1\,(1 - \delta_2)],\ \ \delta_2[1 - \delta_1\,(1 - \delta_2)])$$

where player 1's equilibrium payoff denotes the remaining surplus after she offers $\delta_2[1 - \delta_1\,(1 - \delta_2)]$ to player 2.

Some tricks about solving alternating offer bargaining games

- Table 6.1 summarizes our approach to solve this alternating-offer bargaining game, starting in the last period.
- Vertical arrows indicate that a player must be indifferent between her payoff in the period when she is a proposer and her payoff in the previous period when she is the responder.
- Horizontal arrows, however, denote that the proposing player's equilibrium payoff is the remaining surplus, i.e., the share not offered to her rival.
- Graphically, we move up and sideways, and then up and sideways again, until reaching the first period, arrows resemble a ladder, explaining why our students often refer to this approach as the "ladder".

Player	*Period*	*Player 1's payoff*	*Player 2's payoff*
Player 1	t=1	$1-\delta_2[1-\delta_1(1-\delta_2)]$	$\delta_2[1-\delta_1(1-\delta_2)]$
Player 2	t=2	$\delta_1(1-\delta_2)$	$1-\delta_1(1-\delta_2)$
Player 1	t=3	$1-\delta_2$	δ_2
Player 2	t=4	0	1

Table 6.1. An approach to solve alternating-offers bargaining games.

Alternating-offer bargaining game with infinite periods

- Consider an infinite-period bargaining game where player 1 makes offers to player 2, d^1, in odd numbered periods while player 2 makes to player 1, d^2, in even numbered periods

- We seek to find "Stationary strategies," meaning:
  - Player 1:
    - Makes the same offer to player 2 in every period when player 1 is the proposer (every odd-numbered period), and
    - Follows the same decision rule when he becomes the responder (in every even-numbered period).
  - Player 2:
    - Makes the same offer to player 1 in every period when player 2 is the proposer (every even-numbered period), and
    - Follows the same decision rule when he becomes the responder (in every odd-numbered period).

Alternating-offer bargaining game with infinite periods

Odd-numbered periods:

- Player 2 compares the payoff she gets from accepting the offer from player 1, d^1, against the payoff that she can earn tomorrow when she becomes the proposer.
- Tomorrow, player 2 offers d^2, keeping the remaining $1 - d^2$ of the surplus, with discounted value $\delta_2(1 - d^2)$ in today's terms.
- Player 1 accepts player 1's offer today, d^1, if and only if

$$d^1 \geq \delta_2(1 - d^2).$$

- Player 1 anticipates this decision rule and minimizes her offer, d^1, making player 2 indifferent: $d^1 = \delta_2(1 - d^2)$.

Alternating-offer bargaining game with infinite periods

Even-numbered periods:

- Player 1 compares her payoff from accepting player 2's offer, d^2, and the payoff she can earn tomorrow when she becomes the proposer.
- Player 1 offers d^1 to player 2 tomorrow, keeping the remaining $1 - d^1$ of the surplus, with discounted value $\delta_1(1 - d^1)$.
- Player 1 accepts d^2 today if and only if $d^2 \geq \delta_1 (1 - d^1)$, where player 2 reduces d^2, ultimately making player 1 indifferent between accepting and rejecting, that is,

$$d^2 = \delta_1 (1 - d^1).$$

Alternating-offer bargaining game with infinite periods

Solving the game.

- We have one equation each from odd and even-numbered periods with two unknowns d^1 and d^2.
- Solve for these two offers in equilibrium, we can insert one indifferent condition into the other yielding

$$d^1 = \delta_2 \left(1 - \underbrace{\delta_1 (1 - d^1)}_{d^2} \right)$$

which rearranges to $d^1 = \delta_2 - \delta_1 \delta_2 + \delta_1 \delta_2 \, d^1$, and solving for d^1, yields

$$d^1 = \frac{\delta_2 (1 - \delta_1)}{1 - \delta_1 \delta_2}$$

Alternating-offer bargaining game with infinite periods

- Inserting this result into indifference condition $d^2 = \delta_1(1 - d^1)$, we obtain

$$d^2 = \delta_1\left(1 - \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}\right) = \frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}$$

- We can then conclude that, in the first period, player 1 makes an offer $d^1$to player 2, who immediately accepts it, and the game is over, yielding equilibrium payoffs

$$(1 - d^1, d^1) = \left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}\right)$$

where $1 - d^1 = 1 - \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2} = \frac{1-\delta_2}{1-\delta_1\delta_2}$.

Alternating-offer bargaining game with infinite periods

- Therefore, player 2's equilibrium payoff, d^1, increases in her own discount factor δ_2, since

$$\frac{\partial d^1}{\partial \delta_2} = \frac{1-\delta_1}{(1-\delta_1\delta_2)^2} \geq 0$$

Implying that, as player 2 assigns more weight to her future payoff, she can reject player 1's offer and wait to become the proposer, which forces player 1 to make a more generous offer today.

- In contrast, player 2's equilibrium payoff decreases in her rival's discount factor, δ_1, because

$$\frac{\partial d^1}{\partial \delta_1} = \frac{\delta_2(\delta_2 - 1)}{(1 - \delta_1\delta_2)^2} < 0$$

which intuitively means that, as player 1, assigns more weight to her future payoffs, player 2 must offer her a larger surplus share in future periods (when player 2 becomes the proposer). This, in turn, reduces player 2's future payoffs, making her less attracted to reject the offer from player 1 today.

Appendix:

Mixed and Behavioral Strategies

Appendix – Mixed and Behavioral Strategies

- Definition. **Mixed Strategy.** In a sequential-move game, a mixed strategy is a probability distribution over all player $i's$ pure strategies $s_i \in S_i$.
- In Figure 6.20, the set of pure strategies is $S_1 = \{Ac, Ad, Bc, Bd\}$, and a mixed strategy is a randomization over all these pure strategies.

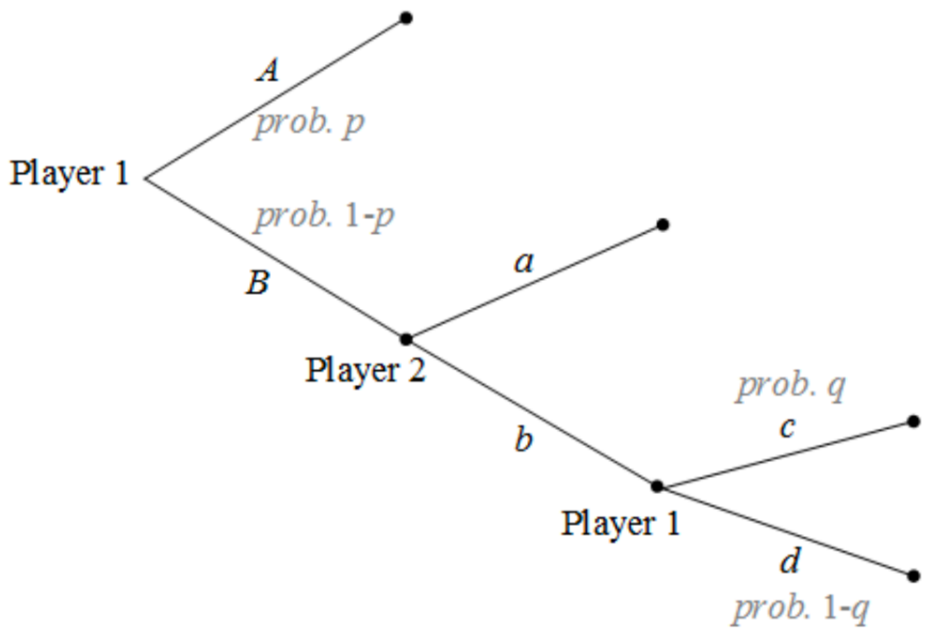

Figure 6.20. Mixed and behavioral strategies in a game tree.

Appendix – Mixed and Behavioral Strategies

- Intuitively, player i rolls a dice, before the game starts, and the outcome of this roll determines the path of play that she will follow (i.e., the pure strategy that she chooses).
- For instance, a mixed strategy for player 1 could be $\left(0Ac, 0Ad, \frac{1}{2}Bc, \frac{1}{2}Bd\right)$,
- which puts:
  - no probability weight on pure strategies Ac and Ad, but
  - assigns 50 percent probability on Bc and Bd.

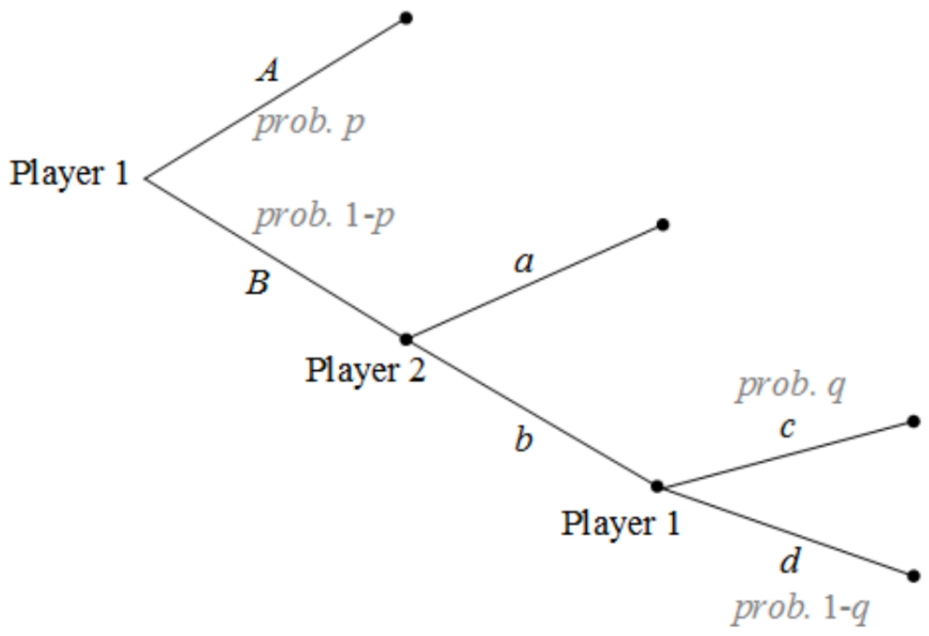

Figure 6.20. Mixed and behavioral strategies in a game tree.

Appendix – Mixed and Behavioral Strategies

- Definition. **Behavioral strategy.**
- A behavioral strategy is a mapping $b_i\colon\ H_i\ \rightarrow\ \Delta A_i(h_i)$, where $b_i\big(s_i(h_i)\big)$ is the probability that player i selects action $s_i(h_i)$ at information set h_i.

- In the context of Figure 6.20, an example of behavioral strategy for player 1 could specify that:
  - At the first node where she is called to move, she randomizes between the two available actions in this node, A and B, such as $pA + (1 - p)B$ where $p\ \in\ [0,1]$.
  - Similarly, in the second node where she is called to move, she randomizes between actions c and d, e.g., $qc + (1 - q)d$, where $q\ \in\ [0,1]$.

- In that context, an example of a behavioral strategy for player 1 is $(p, q) = \left(\frac{1}{3}, \frac{1}{4}\right)$, meaning that:
  - She assigns a probability of $\frac{1}{3}$ to action A in the first node where she is called to move, and
  - Probability $\frac{1}{4}$ to action c in the second node where she plays.

Appendix – Mixed and Behavioral Strategies

- **Equivalence.**
  - Are behavioral strategies equivalent to a corresponding mixed strategy?
  - Yes, we can construct a one-to-one correspondence between behavioral and mixed strategy.

- To see this with the above example, if probabilities p and q in the behavioral strategy satisfy $p = 0$ and $q = \frac{1}{2}$, then the behavioral strategy becomes

$$\left(B, \frac{1}{2}c + \frac{1}{2}d\right)$$

is equivalent to the mixed strategy

$$\left(0Ac, 0Ad, \frac{1}{2}Bc, \frac{1}{2}Bd\right)$$

Since, in one of them, player 1 starts choosing B in pure strategies, and then randomizes between c and d with equal probability.

- Otherwise, the behavioral and mixed strategy do not produce the same outcome.

Appendix – Mixed and Behavioral Strategies

- **“Manuals and Libraries.”**
  - Following Luce and Raiffa’s (1957) analogy, a pure strategy $s_i \in S_i$ can be understood as an “instructional manual” in which each page tells player i which action to choose when she is called on to move at information set h_i.
    - Every manual has as many pages as information sets player i can face.
  - In this analogy, strategy space S_i is a “library” with all possible instruction manuals.
  - A mixed strategy, then, randomly chooses a specific manual s_i from the library S_i (i.e., a specific complete contingent plan), which the player considers for the rest of the game.
  - In contrast, a behavioral strategy chooses pages from different manuals (actions in each information set) with positive probability.