Chapter 3: Perfect-Information Games

ECON 5001
Prof. Paul J. Healy
Based on: *Game Theory* by Giacomo Bonanno

Updated 2026-09-01 at 01:12:50

Everything So Far Happened at Once

In Chapter 2, players chose **simultaneously**, seeing nothing beforehand.

But most interesting situations aren't like that:

- chess
- an offer and a counteroffer
- a firm entering a market and the incumbent responding
- calling a play after seeing the defense line up

Now: players move **in turn**, and everyone sees the whole history.

Log In: Take It or Leave It

- MobLab, our session. You'll be paired with someone here.
- A pile of money sits on the table, and it **grows** every round
- On your turn: **Take** the pile, or **Pass** it back
- Taking ends the game. Whoever takes gets the larger share.

No talking. Play it the way you'd actually play it.

How Far Did We Get?

Pairs that ended at the first decision: ________________

Average round the pile was taken: ________________

Pairs that made it to the end: ________________

Hold onto these numbers. We'll come back to them.

Trees

Perfect Information

In a **dynamic game** (or **extensive form**) players move in sequence.

This chapter covers the case of **perfect information**:

whenever it is a player's turn to move,
she knows every move made so far

Chess has it. Poker does not.

Rooted Directed Trees

A **rooted directed tree** is a set of **nodes** joined by directed **edges**.

Alex and Blair are out to dinner. The check arrives.

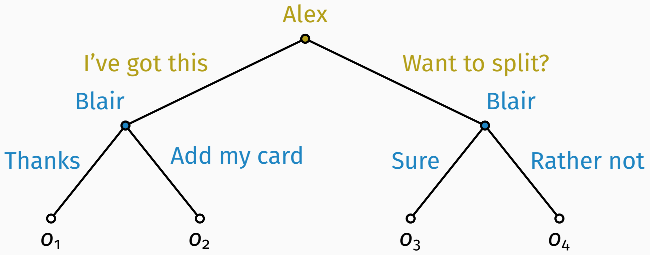

Rooted Directed Trees

A **rooted directed tree** is a set of **nodes** joined by directed **edges**:

- the **root** has no edge leading into it
- every other node has **exactly one** edge leading into it
- there is a **unique path** from the root to any other node

Nodes with no edges out are **terminal** (Z), all others **decision nodes** (D).

$$X = D \cup Z$$

We write x for a particular node, and $S(x)$ for x together with everything that follows it.

Unique paths are what make "the history so far" well defined.

Definition: Extensive Form With Perfect Information

A **finite extensive form with perfect information** consists of:

- a finite rooted directed tree
- a set of players $I = \{1, \ldots, n\}$, one at **each decision node**
- a set of actions A, one assigned to every edge, with no two edges out of the same node getting the same action
- a set of outcomes O, one assigned to every terminal node

Note what's missing: **preferences**. This is a frame, not a game.

Dissolving the Partnership

Amy and Beth are dissolving the partnership behind their Short North coffee shop. The assets are valued at $100,000.

- The charter says the senior partner, **Amy**, proposes a split
- **Beth** can Accept, or Reject and take it to court
- In court: each side pays $20,000 in legal fees, and the verdict (no uncertainty) gives 60% to Amy, 40% to Beth
- Amy has only two possible offers: **50-50** or **70-30**

So rejecting always yields Amy $\$60{,}000 - \$20{,}000 = \$40{,}000$ and Beth $\$40{,}000 - \$20{,}000 = \$20{,}000$.

The Frame

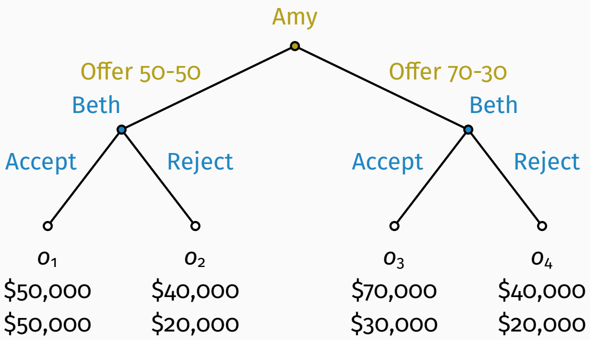

Top number is Amy's, bottom number is Beth's.

Backward-Induction Reasoning

Start at the *end* and work back toward the root.

- Offered 50-50, Beth compares $50,000 to $20,000 $\Rightarrow$ accept
- Offered 70-30, Beth compares $30,000 to $20,000 $\Rightarrow$ accept
- Anticipating both, Amy compares $50,000 to $70,000 $\Rightarrow$ offer 70-30

Clean. Convincing. And built on an assumption we never stated.

The Same Flaw as Golden Balls

We assumed Beth cares only about her own money.

Suppose instead Beth thinks she worked just as hard as Amy, and that the only fair split is 50-50. She may feel strongly enough that she would give up $10,000 to refuse an unfair offer.

That is: $o_4 \succ_{\text{Beth}} o_3$.

Then the reasoning above gives the wrong answer.

A tree with outcomes at the leaves is a **frame**.

Add each player's ranking and you have a **game**.

Adding Preferences

Amy is selfish and greedy: $o_3 \succ_1 o_1 \succ_1 o_2 \sim_1 o_4$

Beth cares about fairness: $o_1 \succ_2 o_2 \sim_2 o_4 \succ_2 o_3$

	o_1	o_2	o_3	o_4
U_1 (Amy)	2	1	3	1
U_2 (Beth)	3	2	1	2

Now replace each outcome with its payoff pair and rerun the algorithm.

The Game

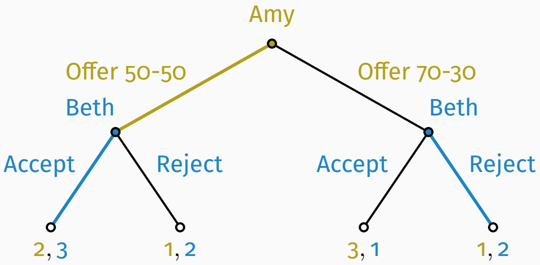

Highlighted edges are the backward-induction choices. Amy now offers **50-50** (she anticipates that 70-30 would be rejected).

Backward Induction

The Algorithm

Call a node **marked** once a payoff vector is attached to it. Initially, exactly the terminal nodes are marked.

1. Pick a decision node x whose immediate successors are **all** marked. Let i be the player who moves at x.
2. Select a choice leading to an immediate successor with the highest payoff *for player i*. Mark x with that successor's whole payoff vector.
3. Repeat until every node is marked.

Finiteness guarantees this terminates: start at nodes followed only by terminal nodes, and work outward.

When Choices Tie

Step 2 says "select a choice." If two choices tie for a player, the algorithm doesn't say which, so it can produce **several** solutions.

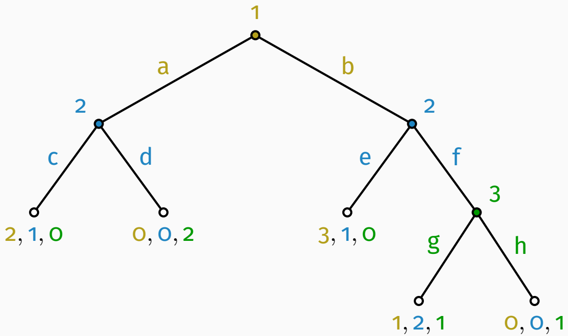

Payoffs in order: Player 1, Player 2, Player 3.

When Choices Tie

Step 2 says “select a choice.” If two choices tie for a player, the algorithm doesn’t say which, so it can produce **several** solutions.

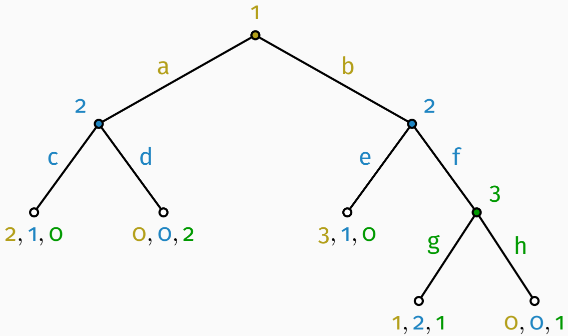

Payoffs in order: Player 1, Player 2, Player 3.

Strategies

A Strategy Is a Complete Plan

Definition
A **strategy** for a player in a perfect-information game is a list of choices, **one for each decision node of that player**.

Not "what I will do" but "what I would do *everywhere* I might move."

In the game two slides back, Player 2 has two decision nodes, so a strategy is a pair. Before the game starts she doesn't know whether Player 1 will pick a or b, so she must plan for both:

$$(c, e) \qquad (c, f) \qquad (d, e) \qquad (d, f)$$

Four strategies, not two.

Counting Strategies

Suppose Player 1 has three decision nodes in some game:

- at the first, three choices a_1, a_2, a_3
- at the second, two choices b_1, b_2
- at the third, four choices c_1, c_2, c_3, c_4

A strategy fills in three blanks:

$$\Big(\underbrace{\rule{3cm}{0.4pt}}_{\text{one of } a_1, a_2, a_3}, \underbrace{\rule{3cm}{0.4pt}}_{\text{one of } b_1, b_2}, \underbrace{\rule{3cm}{0.4pt}}_{\text{one of } c_1, \ldots, c_4} \Big)$$

So her strategy set is a Cartesian product of the action sets she faces:

$$S_1 = \{a_1, a_2, a_3\} \times \{b_1, b_2\} \times \{c_1, c_2, c_3, c_4\}$$

$$3 \times 2 \times 4 = 24 \text{ strategies}$$

Strategy sets get large fast. That's the price of "complete plan."

A Redundancy Worth Noticing

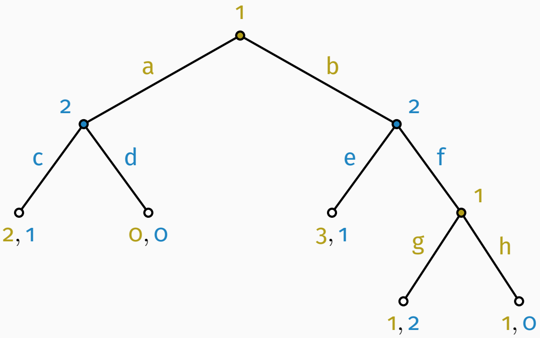

Strategy (a, g) says: play a at the root, and g at my other node.

But if she plays a, that other node is **never reached**. Why plan for it?

Why Keep the Redundancy?

Three ways to justify it:

1. **Mistakes.** She's cautious enough to plan for the possibility that she intends *a* but ends up playing *b*.

2. **Delegation.** A strategy is a set of instructions handed to someone playing on her behalf, who needs to know what to do everywhere.

3. **Beliefs.** A strategy isn't really Player 1's plan at all, but Player 2's *belief* about what Player 1 would do.

We'll set this aside for now and just use the definition.

On Path and Off Path

A strategy profile traces one path from the root to a terminal node.

- Nodes on that path are **on the path of play**
- Every other node is **off the path of play**

A strategy must still say what to do off path. Nobody ever sees those choices, but they are what make the on-path choices rational.

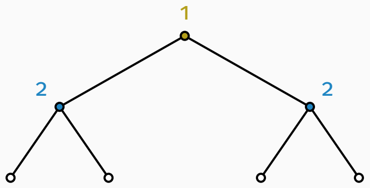

From Tree to Table

Every strategy profile determines a unique path, hence a unique terminal node, hence a unique payoff vector.

So every perfect-information game has an **associated strategic form**.

Player 1 \ Player 2	ce	cf	de	df
ag	2, 1	2, 1	0, 0	0, 0
ah	2, 1	2, 1	0, 0	0, 0
bg	3, 1	1, 2	3, 1	1, 2
bh	3, 1	1, 0	3, 1	1, 0

The top two rows are identical. That's the redundancy showing up.

From Tree to Table

Every strategy profile determines a unique path, hence a unique terminal node, hence a unique payoff vector.

So every perfect-information game has an **associated strategic form**.

Player 1 \ Player 2	ce	cf	de	df
ag	2, 1	2, 1	0, 0	0, 0
ah	2, 1	2, 1	0, 0	0, 0
bg	3, 1	1, 2	3, 1	1, 2
bh	3, 1	1, 0	3, 1	1, 0

The top two rows are identical. That's the redundancy showing up.

Five Equilibria, Two Solutions

That game has **five** Nash equilibria:

$$(ag, cf) \quad (ah, cf) \quad (bg, df) \quad (bh, ce) \quad (bh, de)$$

It has **two** backward-induction solutions:

- $\big((a, g), (c, f)\big)$ (because g and h tie for Player 1)
- $\big((b, h), (c, e)\big)$

Both are Nash equilibria.
The other three are not backward-induction solutions.

That gap is the whole point of the next section.

Solution vs. Outcome

A backward-induction **solution** is a *strategy profile*. It describes behavior everywhere, including at nodes that never get reached.

A backward-induction **outcome** is the *sequence of moves actually made*.

$$\text{Solution } \big((a,g),(c,f)\big) \quad \text{outcome } ac, \ \text{payoffs } (2,1)$$

$$\text{Solution } \big((b,h),(c,e)\big) \quad \text{outcome } be, \ \text{payoffs } (3,1)$$

Keep the distinction straight. Exam questions live here.

Backward Induction & Nash

The Theorem

Theorem. Every backward-induction solution of a perfect-information game is a Nash equilibrium of the associated strategic form.

Suppose player *i* deviates.

- If she only changes what she does off the path, the play doesn't change, so her payoff doesn't either
- If she changes what she does at a node *on* the path, the algorithm already checked that alternative and rejected it

So backward induction is a **refinement** of Nash equilibrium: fewer solutions, all of them Nash.

Which Nash Equilibria Get Cut?

Usually the ones resting on an **incredible threat**.

The Entry Game

IKEA has been the only big furniture store at Polaris for years, earning $5 million.
Wayfair is deciding whether to open across I-71.

- If Wayfair stays **out**, it earns $1 million elsewhere
- If Wayfair enters, IKEA can **fight**, a price war leaving both at zero
- Or IKEA can **accommodate** and share, and both make $2 million

Both firms are selfish and greedy.

IKEA announces: "open here and we will bury you on price."
Should Wayfair believe it?

The Entry Game

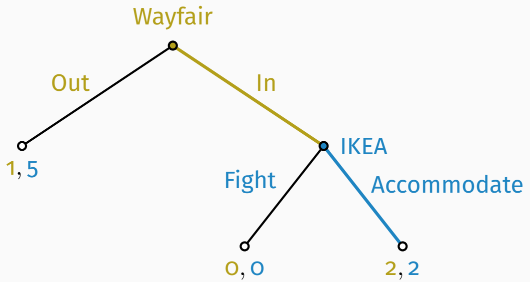

First number Wayfair, second number IKEA.

Backward induction: faced with entry, IKEA compares 0 to 2. It accommodates. Knowing that, Wayfair enters.

But There Are Two Nash Equilibria

		IKEA: Fight	IKEA: Accommodate
Wayfair	In	0, 0	2, 2
Wayfair	Out	1, 5	1, 5

- (In, Accommodate), the backward-induction solution
- (Out, Fight), also a Nash equilibrium!

It survives only because Wayfair never enters, so the threat is untested.

It is an **incredible threat**: if entry actually happened, IKEA would not want to carry it out.

Selten's Chain-Store Game

IKEA also has a store in Cincinnati, and Wayfair could expand there next. So the entry decision happens **one city at a time**, with everyone watching what happened in the earlier one.

IKEA makes the argument in CBus:

> *Fighting you costs me \$2 million today. But Cincy is watching. If they see me fight, they may stay out, and then I make* $\$(0+5) = \$5$ *million across the two cities. If I accommodate you here, they enter there too, and I make* $\$(2+2) = \$4$ *million. So it is in my interest to fight. Stay out.*

Sounds right. Does backward induction agree?

The Reply

No, and the reason is where backward induction starts.

- Whatever happens in CBus, Cincy is the **last** interaction
- Nobody is watching after that, so no reputation is left to build
- So Cincy is just the one-shot entry game, and IKEA accommodates

- Which means Cincy goes the same way *regardless* of CBus
- So fighting in CBus buys IKEA nothing: $\$(0+2) = \$2$ million against $\$(2+2) = \$4$ million

Entry happens in both cities. IKEA accommodates both.

What That Tells Us

The reputation argument is intuitive and backward induction rejects it.

- That's not a flaw in the arithmetic. The unraveling is airtight
- The reputation story needs **uncertainty**: the Cincy entrant must be unsure what kind of opponent IKEA is
- In a perfect-information game, uncertainty is ruled out by definition

We'll come back to this with better tools. For now, note the limitation.

Backward Induction and IDWDS

Backward induction resembles deleting dominated choices at each node.

Applying IDWDS to the strategic form of a perfect-information game gives profiles containing **at least one** backward-induction solution.

But that's all you can say in general. The IDWDS output:

- may also contain profiles that are **not** backward-induction solutions
- may **miss** some backward-induction solutions

So they're related, not the same. Don't use one to compute the other.

Win, Lose, or Draw

Win-Lose Games

Two players, two outcomes: Player 1 wins (W_1) or Player 2 wins (W_2).

Player 1 strictly prefers W_1; Player 2 strictly prefers W_2. Use payoffs $(1, 0)$ and $(0, 1)$.

Theorem. In every finite two-player win-lose game with perfect information, **one of the two players has a winning strategy.**

A winning strategy guarantees a win *no matter what the opponent does*.

Note what it doesn't say: *which* player, or what the strategy is.

Let's Play

Two players take turns choosing a number from $\{1, 2, \ldots, 10\}$.
Running total starts at 0. **First to reach 100 or more wins.**

Player 1 moves first.

Volunteer? I'll go second.

Why I Was Always Going to Lose That

The tree is far too big to draw: 10,000 nodes just for the first four moves. But we can reason backward anyway.

Call a total a **losing position** if whoever has to move from it cannot win.

- **89** is a losing position: from 89 you can reach only $90, \ldots, 99$, and from any of those I can hit exactly 100
- **78** is the next one down: from 78 you must reach $79, \ldots, 88$, and from any of those I can get to 89
- Keep going: $89,\ 78,\ 67,\ 56,\ 45,\ 34,\ 23,\ 12,\ \mathbf{1}$

Player 1's Winning Strategy

Start with **1**.

Then every turn, choose $\mathbf{11-n}$, where n is what your opponent just chose.

Each pair of moves adds exactly 11, so the total marches $1, 12, 23, \ldots, 89, 100$.

Notice this is a genuine *strategy*: it says what to do at every position Player 1 could face, not just the ones on the winning path.

Adding a Draw

Now three outcomes: W_1, W_2, and a draw D, with

$$W_1 \succ_1 D \succ_1 W_2 \qquad \text{and} \qquad W_2 \succ_2 D \succ_2 W_1$$

Theorem. Every such game falls into exactly one of three categories:

1. Player 1 has a strategy guaranteeing W_1
2. Player 2 has a strategy guaranteeing W_2
3. Player 1 can guarantee *at least* a draw, and so can Player 2, so if both play well it's a draw

The proof is the same backward-induction argument, with three possible markings instead of two.

Which Category?

- **Tic-Tac-Toe:** category 3. Solved: perfect play draws.
- **Checkers:** category 3. Solved in 2007, after 18 years of work.
- **Chess:** ______________________________

The theorem says chess is in *one* of the three categories. It does not say which, and nobody knows.

Existence proofs are like that.

Back to the Centipede

Apply backward induction to the game you played:

- At the **last** decision, taking beats passing. So take.
- At the second-to-last, you know the other player will take next, so passing gets you the smaller share. So take.
- Keep unraveling … all the way to the first move.

Prediction: the very first player takes immediately.

Our numbers said: ________________________________

So Who's Wrong?

Three honest possibilities:

- The **payoffs** are wrong. Players may care about the other person or about fairness, so the game isn't the one we wrote down
- The **reasoning** is wrong. Players may not carry the unraveling through that many steps
- The **beliefs** are wrong. I might reason perfectly and still pass if I doubt that *you* will

All three are live. We'll get the machinery to separate them later.

Let’s Work One Out

Planning Spring Break

Ann, Ben, and Cal share a house and settle the trip in turn, each seeing every decision made before theirs.

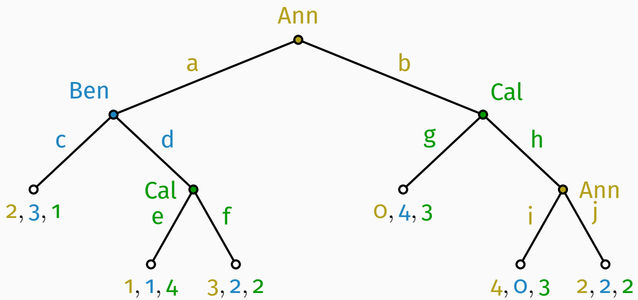

Payoffs in order: Ann, Ben, Cal.

Planning Spring Break

1. How many strategies does each player have?

2. Run backward induction. Where does a player face a tie?

3. Write out **both** backward-induction solutions, as strategy profiles.

4. What are the two backward-induction outcomes?

What Have We Learned, Charlie Brown?

- Perfect-information games are drawn as **rooted trees**; a frame becomes a game once preferences are added
- **Backward induction** works from the leaves inward; ties give several solutions
- A **strategy** covers all your decision nodes, even unreachable ones
- Every backward-induction solution is a **Nash equilibrium**; the extra Nash equilibria typically rest on **incredible threats**
- In finite two-player win-lose games, **somebody** has a winning strategy

Next: what happens when players *don't* see the whole history.

Reading & Practice

- Bonanno, *Game Theory*, Chapter 3
- Exercises: §3.7.1 (trees and frames), §3.7.2 (backward induction), §3.7.3 (strategies), §3.7.4 (two-player games)
- Exercise 3.8 shows IDWDS missing a backward-induction solution
- Solutions in §3.8