Lecture Note 9: Dynamic Games

So far we've studied **static games**:

- Both players move at the same time (and only once).
- Each player has to pick her strategy before she knows which strategy the other player chose.

Sometimes it's reasonable to assume players move simultaneously:

- Prisoner's Dilemma.
- Rock-Paper-Scissors.
- Penalty kicks in soccer.

But sometimes players move sequentially:

- Tesla enters electric car market, Toyota follows (or not).
- N. Korea tests a nuclear bomb, US applies sanctions (or not).

We'll solve these problems by using **backward induction**.

Key questions:

- When is it better to move first? When is it better to move second?
- When is a threat a credible threat? When is it a bluff?
- When will self-interested players cooperate with each other?

Game trees

We represent static games with a payoff matrix. Here's one we've seen:

		Novartis: Enter	Novartis: Pass
Pfizer	Enter	$-1, -1$	$3, 0$
	Pass	$0, 3$	$0, 0$

For dynamic games, we use a **game tree** that shows who moves when. For example, suppose that Pfizer gets to move first:

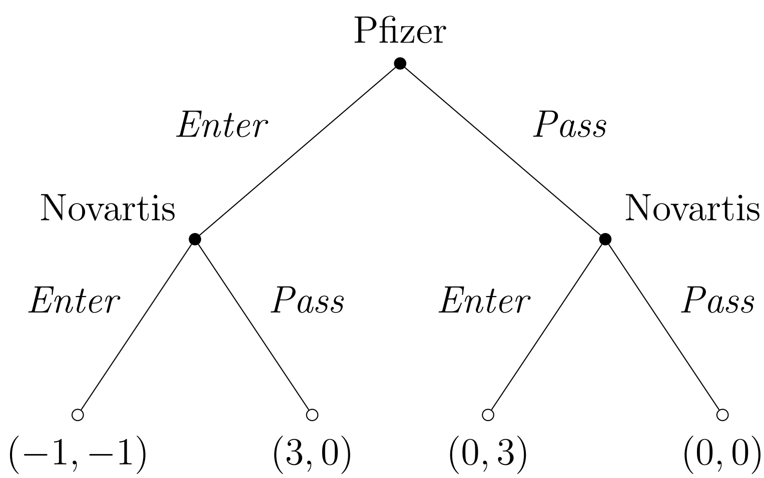

A game tree tells us:

- The order in which the players move.
- The actions available to each player at each **decision node**. (The decision nodes are the black circles in the game tree above.)
- The payoffs after each sequence of actions (player 1's, player 2's).

When studying dynamic games, we will assume that:

- Each player observes all prior moves. (Novartis sees Pfizer's choice.)
- Each player *knows* that the other player will see its move.

Actions, strategies, and backwards induction

In dynamic games, we need to distinguish "actions" from "strategies".

An **action** is a single "move" in the game.

- Prisoner's Dilemma: actions are "Flip" and "Deny".
- Pharmaceutical entry: actions are "Enter" and "Pass".

A **strategy** is a complete "if-then" plan for playing the game.

- Since Pfizer goes first, its only strategies are "Enter" and "Pass".
- Since Novartis goes second, its strategy can depend on Pfizer's move. Here is one of Novartis's possible strategies: "If Pfizer chooses Enter, we'll play Enter. If Pfizer chooses Pass, we'll play Pass."

To predict each player's strategy, we will use **backward induction**.

- Remember how we solved the monopolist's entry problem: decide how much to produce *if I enter*, then decide whether to enter.
- We're doing something similar here, but with *two* decision-makers.

Why work backwards? Consider a game with 100 "turns".

- To pick the best action on turn 1, I need to predict what will happen on turn 2, turn 3, turn 4, ..., turn 100. Yikes!
- To pick the best action on turn 100, I can just compare final payoffs. We can solve this mini-problem (or "subgame") in isolation.
- Once we solve turn 100, it's easier to solve turn 99, then turn 98, ...
- By the time we get to turn 1, we effectively know what will happen in every subsequent stage of the game. That makes it easier to figure out what the player should choose on turn 1.

Solving games by backwards induction is much easier than trying to start on turn 1 and work your way forward to the last turn of the game.

Finding the Nash equilibrium through backward induction

Here's our game tree again:

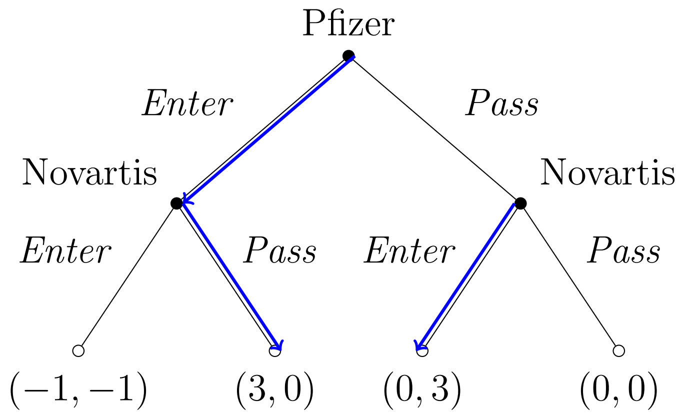

After Pfizer moves, what is Novartis's best response?

- If Pfizer picks Enter, Novartis's best response is Pass.
- If Pfizer picks Pass, Novartis's best response is Enter.

If Pfizer expects this, Pfizer will Enter. This gives us Nash equlibrium #1:

- Pfizer's strategy is simple: "Enter."
- Novartis's strategy has two parts, one for each possible situation: "If Pfizer plays Enter, we'll Pass. If Pfizer plays Pass, we'll Enter."

But there's also a Nash equilibrium #2, where Novartis "acts crazy":

- Pfizer: "Pass."
- Novartis: "No matter what Pfizer does, we'll Enter."

To see whether this is a Nash equilibrium, we ask: can either player benefit from **deviating**, *assuming that the other player "sticks to the script"*?

- Can Pfizer increase its payoff by deviating? No: if Pfizer expects Novartis to Enter no matter what, then Pfizer should Pass.
- Can Novartis increase its payoff by deviating? No: as long as Pfizer's strategy is "Pass", then Novartis's "crazy" strategy works just fine.

So this is a Nash equilibrium too! But it's an implausible one.

Subgame-perfect Nash equilibrium

A **subgame** is a "game within a game". Each subgame consists of one particular decision node and everything that follows it.

- By convention, the entire game is considered a subgame of itself—namely, it's the subgame that starts with the first decision node.
- The pharmaceutical game has three subgames: the whole game itself, the subgame starting with Novartis's decision after Pfizer enters, and the subgame starting with Novartis's decision after Pfizer passes.

A pair of strategies is a **subgame-perfect Nash equilibrium** (SPNE) if they form a Nash equilibrium in every subgame (including the full game).

- Every subgame-perfect Nash equilibrium is a Nash equilibrium.
- But not every Nash equilibrium is subgame-perfect.
- Intuition: a Nash equilibrium is subgame-perfect as long as nobody is making any "empty threats" (or "bluffs") about what they would do if their opponent tried to deviate from the equilibrium "script".

In the pharmaceutical game,

- Which Nash equilibrium is subgame-perfect?
  Answer—Pfizer: "Enter." Novartis: "Only Enter if Pfizer Passes."
- Which Nash equilibrium is <u>not</u> subgame-perfect?
  Answer—Pfizer: "Pass." Novartis: "Enter no matter what."

Good news: backward induction always finds subgame-perfect equilibria.

- Why? When we use backwards induction, we're assuming that each player always chooses to play her best response whenever it's time for her to move (even if she previously "threatened" that she would do something else). This procedure automatically rules out bluffs.
- In practice, we simply solve dynamic games by backward induction. If we find a Nash equilibrium that way, we know it's subgame-perfect.

First-mover advantage: the power of commitment

Since both players know the "rules of the game", it seems unlikely that a player will be able to bluff successfully: a bluff is not a "credible threat", and the other player will ignore it. In other words, we should expect the players to end up playing a SPNE, not some other Nash equilibrium.

In the pharmaceutical game, Pfizer can "call Novartis's bluff": Pfizer will Enter, Novartis will Pass, and Pfizer will get the whole market.

Moving first is very advantageous in the pharmaceutical entry game. When both players would rather move first, we say that a game exhibits **first-mover advantage**. Why is better to move first in this game?

- If Pfizer and Novartis move at the same time, they are playing a game of "chicken"—which player will "chicken out" and Pass?
- But if Pfizer gets to move first, it has the ability to commit to a particular action. Once Pfizer plays Enter, there is no turning back: once a player moves, they cannot change their action.
- That might seem risky—but if the other player is rational, it's not. With common knowledge of rationality, Pfizer knows that Novartis will stay out of the market, and Pfizer will make lots of profit.

This is a profound insight about strategic interactions.

- In most decision problems, it never hurts to keep your options open.
  - In fact, we often say that there is **option value** to postponing a decision to a later date, when you will have more information about the consequences of each possible action.
- But in strategic games, keeping your options open can be very costly.
  - If you leave yourself a way out, other players will call your bluff.
  - If you "burn your boats", other players will know you're serious.

Predator-Prey

Is going first always advantageous? Let's look at a "predator-prey" game. A wolf is hunting a deer. The deer can hide in the woods or in a field, and the wolf can look for the deer in either location:

		Wolf: Woods	Wolf: Field
Deer	Woods	1, 1	3, 0
	Field	3, 0	0, 3

If the Deer moves first, what's the subgame-perfect Nash equilibrium?

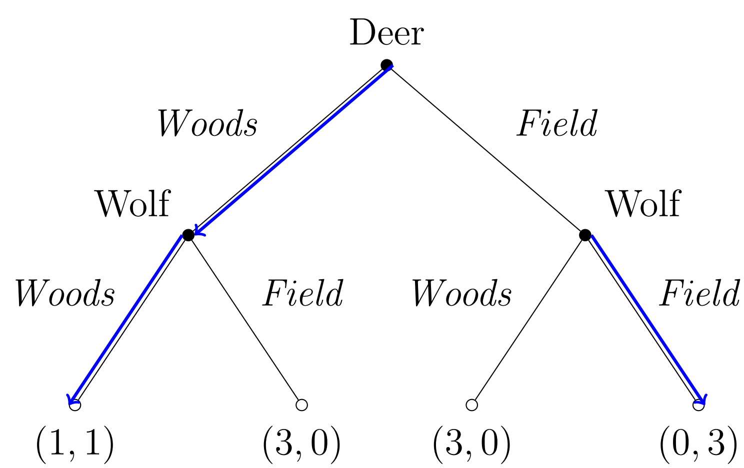

Remember: when we're looking for a SPNE, we need to describe each player's complete "if-then" strategy—their "battle plan" specifying what they'll do in every possible situation. We describe the SPNE as follows:

- Deer's strategy: "Woods"
- Wolf's strategy: "If Deer goes to the Woods, I'll go to the Woods. If Deer goes to the Field, I'll go to the Field."

Wolf's plan to go to the Field if Deer goes to the Field is never actually implemented. But the Wolf has to prepare for every possible situation.

Second-mover advantage: the power of information

In this predator-prey game, it's advantageous to move second.

If both players would rather move second, we say that a game exhibits **second-mover advantage**.

Why is it better to move second in this game?

- The second mover has more information.
- If the Wolf moves second, it can find the Deer.
- If the Deer moves second, it can escape the Wolf.

Whoever moves last always has an informational advantage.

Some of the games we've studied have first-mover advantage, some of them have second-mover advantage, and some games have neither one.

- Rock-Paper-Scissors: second-mover advantage.
- Pharmaceutical Entry: first-mover advantage.
- Prisoner's Dilemma: neither. The outcome is the same no matter who moves first: each player plays their dominant strategy.
- Cuban Missile Crisis: neither. In the dynamic version of this game, since nobody wants a war, the first-mover would choose Negotiate, and the second-mover would respond by choosing Negotiate as well.
- Name the Biggest Number: second-mover advantage.

Cooperation?

The **centipede game** is usually drawn as a horizontal game tree, with the first decision node shown on the left. (It gets its name from the fact that the game tree looks like a centipede.) We can use it to analyze two business partners deciding whether to work together or cheat each other.

- Players 1 and 2 take turns deciding whether to “stop” or “go”.
- If no one stops it early, the game ends after player 2’s third move.
- The longer the game goes on, the bigger the “pie” becomes.
- But if I stop the game early, I get to keep a bigger share of the pie.

What’s the subgame-perfect Nash equilibrium?

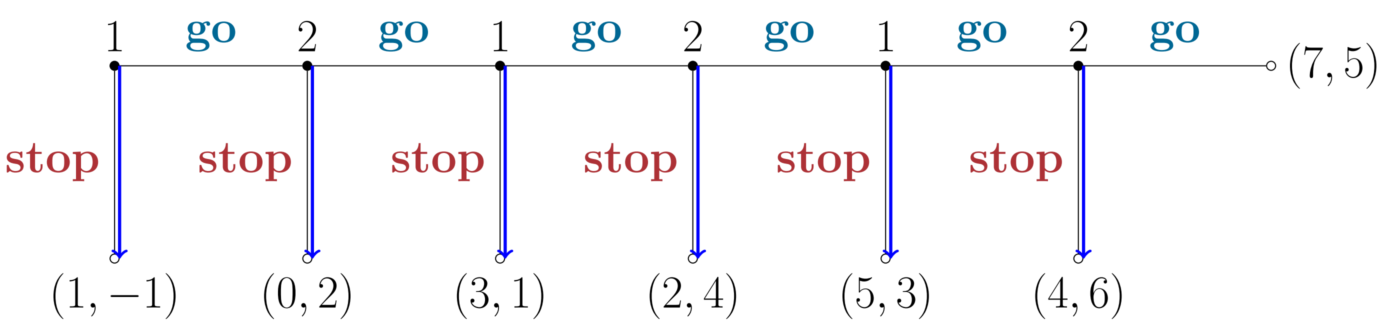

Solving by backward induction …

- At the last decision node, player 2 chooses “stop” since “stop” gives her a payoff of 6 and “go” gives her a payoff of 5.
- At the second-to-last decision node, player 1 chooses “stop” since otherwise player 2 will stop the game next turn.
- Repeating this logic, we find that both players choose “stop” on their turn. The (only) subgame-perfect Nash equilibrium is “player 1 always chooses stop, and player 2 always chooses stop”.

What if the centipede game went on for 100 periods, with a huge payoff for each player if they make it to the end? Answer: the same thing!

Appendix: repeated games (not on exams)

Economists often study games in which a single static game (called the **stage game**) is played repeatedly. We call these **repeated games**. We won't be covering repeated games in this class (and they won't appear on the exams), but I'm including this appendix in case you're curious.

The most famous example is the Repeated Prisoner's Dilemma, which we use to think about when **cooperation** is sustainable and when it isn't.

The US and China are drafting an agreement to limit their CO_2 emissions.

- Each country decides whether to "Honor" the agreement or "Cheat".
- It's in their collective interest to work together, but each one is tempted to prioritize its own immediate GDP over the global climate.

The story is different, but this is still a Prisoner's Dilemma-style game:

		China	
		Honor	Cheat
United States	Honor	$1, 1$	$-1, 2$
	Cheat	$2, -1$	$0, 0$

If we play this game just once, both countries Cheat, since "Cheat" is a dominant strategy in the one-period Prisoner's Dilemma.

But what if we play this stage game twice?

- First, we play the stage game as usual.
- Then we see what happened and play it a second time.
- Each country's total payoff is the sum of its two stage payoffs.

Will the threat of retaliation enable the US and China to work together?

Finitely repeated Prisoner's Dilemma (not on exams)

No! In the 2-period Prisoner's Dilemma, both players still Cheat.

The subgame-perfect Nash equilibrium is (Cheat, Cheat) in both periods.

Why isn't there a way for the players to cooperate, at least in period 1?

- No matter what happens in period 1, Cheating is still a dominant strategy in period 2.
  - The game is about to end, so there's no incentive to cooperate. The other player can't reward or punish me next period.
  - Therefore, both players will Cheat in period 2.
- No matter what I do today, the other player will Cheat next period.
  - Even if I "play nice" today, I'll still get "punished" tomorrow.
  - I might as well "play dirty" today.
- As a result, both countries Cheat in period 1 as well as period 2.

Maybe 2 periods isn't enough ... what if they play for 100 periods?

- Period 100: each country Cheats since doing so has no consequences.
- Period 99: each country Cheats since it gets punished either way.
- ⋮
- Period 1: each country Cheats since it gets punished either way.

In Finitely Repeated Prisoner's Dilemma, the only subgame-perfect Nash equilibrium is for both players to Cheat in every period.

This is a depressing conclusion! It suggests that self-interested decision-makers—whether two business partners or two governments—will never cooperate with each other when doing so requires short-term sacrifices.

Infinitely repeated Prisoner's Dilemma (not on exams)

In games where cooperation involves short-term "sacrifice", will self-interested players ever cooperate with each other? Is cooperation possible?

Maybe, if one of the following is true:

- The game lasts forever, or;
- The game ends eventually, but the players don't know exactly when.

Now the players never think: "The world is ending, so I should Cheat."

Instead, they weigh short-term temptations against the long-term benefits of cooperation. If they're patient enough, they may be able to cooperate.

Suppose the US and China pursue the following strategies.

- In period 1: both countries Honor the agreement.
- In all other periods: if anyone has ever Cheated before, both players Cheat. Otherwise, both players keep playing Honor.

Under this so-called **grim trigger strategy**, the players cooperate until someone cheats. But if somebody cheats even once, they suffer eternal punishment (hence the name "grim trigger").

Suppose both countries have discount factors $\delta < 1$, so that the **present discounted value** of each country's payoff is

$$PDV = x_0 + \delta x_1 + \delta^2 x_2 + \delta^3 x_3 + \ldots$$

If the US thinks China will play grim trigger, what will the US do?

- If the US Cheats: $PDV = 2 + 0 + 0 + \ldots = 2$.
- If the US Honors: $PDV = 1 + \delta + \delta^2 + \delta^3 + \ldots = \frac{1}{1-\delta}$ (where the last step uses the fact that this is a "geometric sum").

So the US will Honor if $\frac{1}{1-\delta} > 2 \implies \delta > \frac{1}{2}$. There's still hope!