



COOPERATION: HOW TO MODEL IT, HOW TO FOSTER IT, AND HOW IT MIGHT HAVE EMERGED

TOMORROW & TOMORROW & TOMORROW

ITERATED GAMES

Adrian Haret
a.haret@uva.nl

December 7-14, 2023 & January 11, 2024

How did the Prisoner's
Dilemma come about?

Melvin and I came up with the idea behind the Prisoner's Dilemma in the 50's, while working for the RAND corporation.

MERRIL FLOOD



MELVIN DRESHER



MERRIL FLOOD

Melvin and I came up with the idea behind the Prisoner's Dilemma in the 50's, while working for the RAND corporation.



MELVIN DRESHER

Amidst concern about nuclear war and political instability, we devised a little game to show that the Nash equilibrium sometimes makes strange predictions.



MERRIL FLOOD



MERRIL FLOOD

Melvin and I came up with the idea behind the Prisoner's Dilemma in the 50's, while working for the RAND corporation.



MELVIN DRESHER

Amidst concern about nuclear war and political instability, we devised a little game to show that the Nash equilibrium sometimes makes strange predictions.



MERRIL FLOOD

We made two of our friends, AA and JW, play the game over 100 times, and recorded their reactions.



MELVIN DRESHER



Poundstone, W. (1993). *Prisoner's Dilemma: John Von Neumann, Game Theory and the Puzzle of the Bomb*. Anchor Books.

MERRIL FLOOD

Melvin and I came up with the idea behind the Prisoner's Dilemma in the 50's, while working for the RAND corporation.



MELVIN DRESHER

Amidst concern about nuclear war and political instability, we devised a little game to show that the Nash equilibrium sometimes makes strange predictions.



MERRIL FLOOD

We made two of our friends, AA and JW, play the game over 100 times, and recorded their reactions.



MELVIN DRESHER

For all the confusion, mutual cooperation occurred 60 out of the 100 trials.



Poundstone, W. (1993). *Prisoner's Dilemma: John Von Neumann, Game Theory and the Puzzle of the Bomb*. Anchor Books.

| Game | AA | JW | AA's comments | JW's comments |
|------|----|----|---|---|
| 1 | D | C | JW will play [D]—sure win. Hence if I play [C]—I lose. | Hope he's bright. |
| 2 | D | C | What is he doing?!! | He isn't but maybe he'll wise up. |
| 3 | D | D | Trying mixed? | Okay, dope. |
| 4 | D | D | Has he settled on [D]? | Okay, dope. |
| 5 | C | D | Perverse! | It isn't the best of all possible worlds. |
| 6 | D | C | I'm sticking to [D] since he will mix for at least 4 more times. | Oh ho! Guess I'll have to give him another chance. |
| 7 | D | C | | Cagey, ain't he? Well . . . |
| 8 | D | D | | In time he could learn, but not in ten moves so: |
| 9 | D | D | If I mix occasionally, he will switch—but why will he ever switch from [D]? | |
| 10 | D | D | Prediction. He will stick with [D] until I change from [D]. I feel like DuPont. | I can guarantee myself a gain of 5, and guarantee that Player AA breaks |

PRISONER'S DILEMMA 109

even (at best). On the other hand, with nominal assistance from AA, I can transfer the guarantee of 5 to Player AA and make 10 for myself, too. This means I have control of the game to a large extent, so Player AA had better appreciate this and get on the bandwagon.

With small amounts of money at stake, I would (as above) try (by using [C]) to coax AA into mutually profitable actions. With large amounts at stake I would play [D] until AA displayed some initiative and a willingness to invest in his own future. One play of [C] by AA would change me from [D] to [C], where I would remain until bitten.

On the last play, it would be conservative for me to switch to [D], but I wouldn't do so if the evidence suggested that AA was a nice stable personality

| Game | AA | JW | AA's comments | JW's comments |
|------|----|----|--|---|
| | | | | and not in critical need of just a little extra cash. |
| 11 | D | C | | Probably learned by now. |
| 12 | C | C | | I'll be damned! But I'll try again. |
| 13 | C | C | | That's better. |
| 14 | C | C | | Ha! |
| 15 | C | C | | (bliss) |
| 16 | D | C | | |
| 17 | C | D | | The stinker. |
| 18 | C | D | | He's crazy. I'll teach him the hard way. |
| 19 | D | D | I'm completely confused. Is he trying to convey information to me? | Let him suffer. |
| 20 | D | D | | |
| 21 | D | C | | Maybe he'll be a good boy now. |
| 22 | C | C | | Always takes time to learn. |

Are AA and JW irrational?

MERRIL FLOOD
What do you say to that, John?!



JOHN NASH



MERRIL FLOOD

What do you say to that, John?!



JOHN NASH

You know, playing the Prisoner's Dilemma one time is not the same as playing it 100 times.

Playing it over and over again is like playing a different, multi-round game.

In the one-shot game there's no room for things like loyalty, trust, threats, or revenge.

But in the iterated version, these things can be relevant!

This gives us the first way out of the pessimistic outlook of the Prisoner's Dilemma.

This gives us the first way out of the pessimistic outlook of the Prisoner's Dilemma.

Does the equilibrium change if the game is played repeatedly?

So far we've been assuming that players make moves simultaneously, in ignorance of the other players' actions.

So far we've been assuming that players make moves simultaneously, in ignorance of the other players' actions.

But, of course, some games are played over rounds.

GAMES IN EXTENSIVE FORM

In *perfect-information extensive-form games*,
players take turns deploying their
actions.

And are aware of actions taken at
previous rounds: perfect memory!

Player 1 takes an action

...out of their action set: $\{a, a'\}$

Player 2 follows up

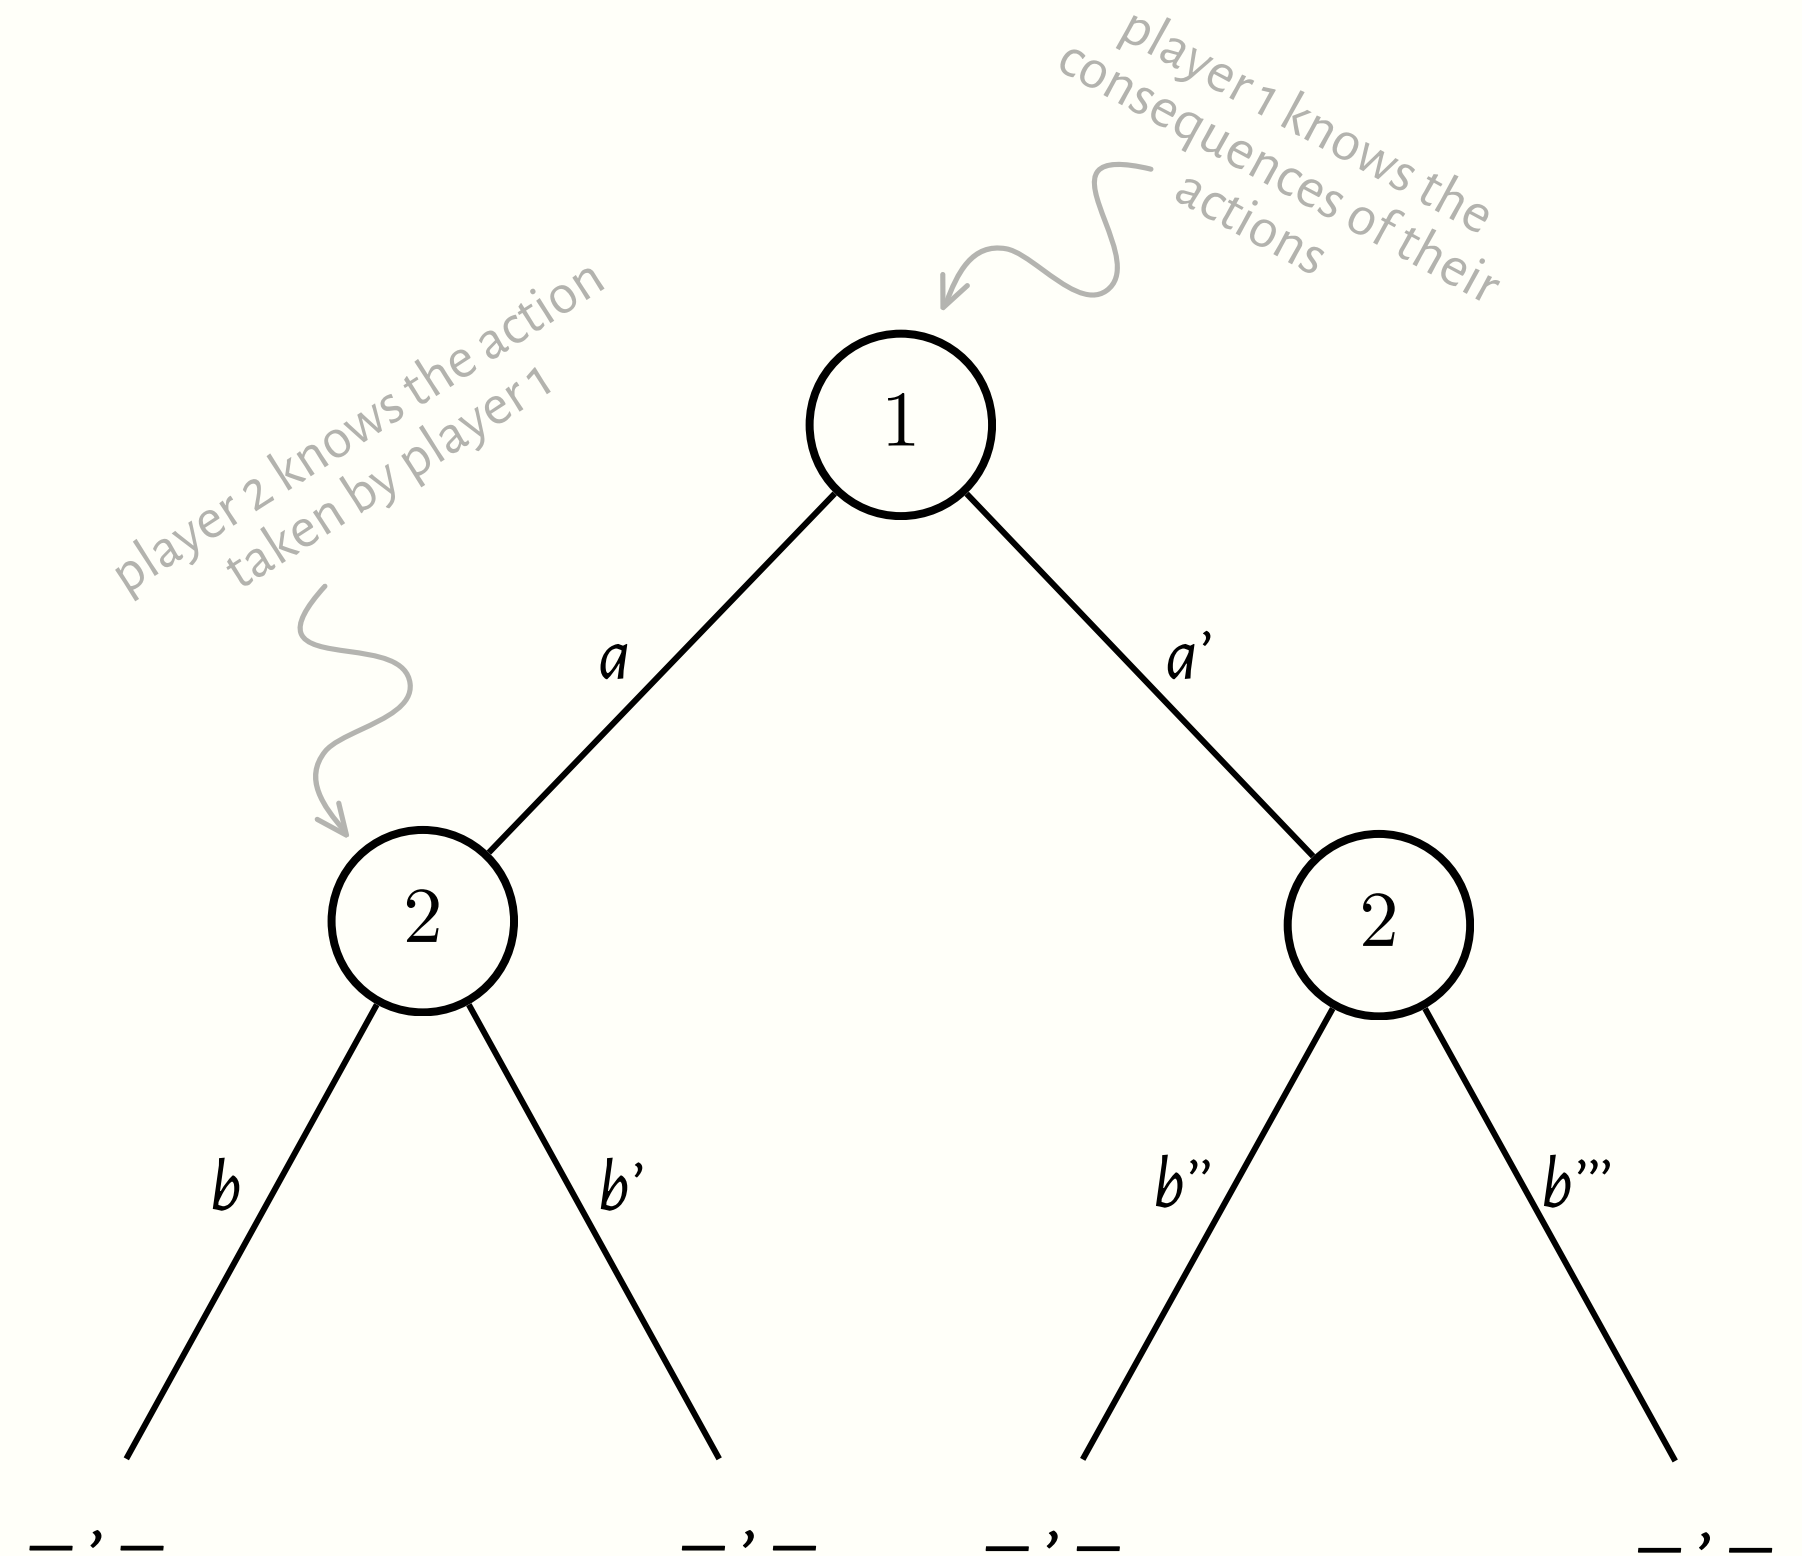
... knowing the action player 1 has taken

Every player receives a payoff

... specific to the branch taken

The whole game tree is known

... to all players



Extensive-form games with perfect information are modeled as *trees*, where non-terminal nodes (called *choice nodes*) correspond to *players*.

Extensive-form games with perfect information are modeled as *trees*, where non-terminal nodes (called *choice nodes*) correspond to *players*.

At every one of its choice nodes, an agent has some *actions* available.

Extensive-form games with perfect information are modeled as *trees*, where non-terminal nodes (called *choice nodes*) correspond to *players*.

At every one of its choice nodes, an agent has some *actions* available.

Each edge is labeled with the action taken by the parent agent at that node.

Extensive-form games with perfect information are modeled as *trees*, where non-terminal nodes (called *choice nodes*) correspond to *players*.

At every one of its choice nodes, an agent has some *actions* available.

Each edge is labeled with the action taken by the parent agent at that node.

Terminal nodes are labeled with the *utilities* of the players for the combination of actions that led to that particular outcome.

Extensive-form games with perfect information are modeled as *trees*, where non-terminal nodes (called *choice nodes*) correspond to *players*.

At every one of its choice nodes, an agent has some *actions* available.

Each edge is labeled with the action taken by the parent agent at that node.

Terminal nodes are labeled with the *utilities* of the players for the combination of actions that led to that particular outcome.

A *strategy* for an agent is a combination of actions, one for each node corresponding to that agent.

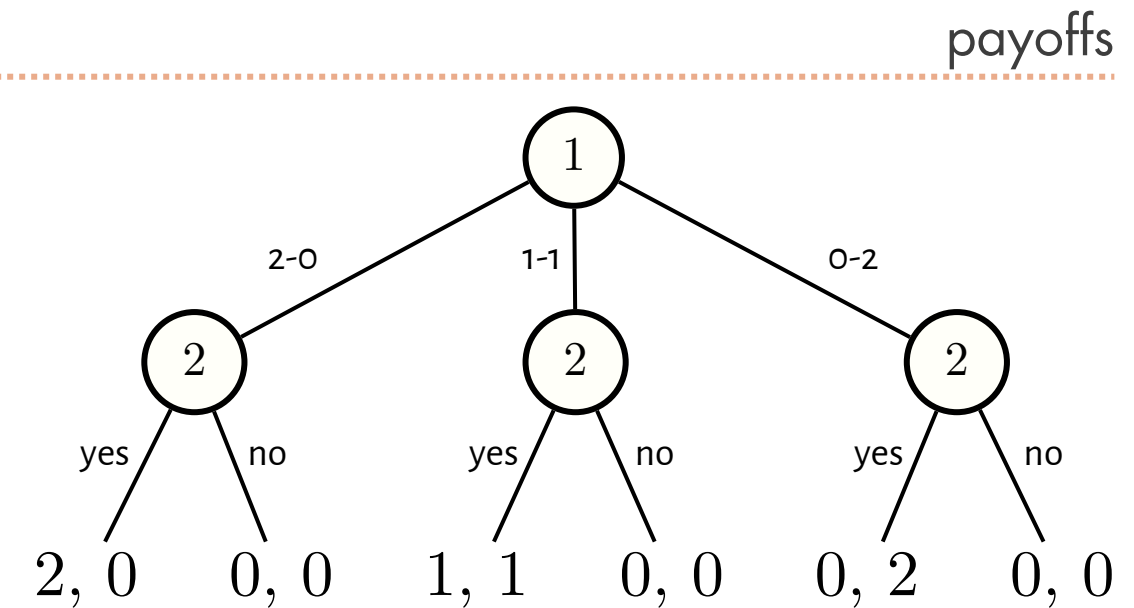
The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.



strictly dominant strategies


Pareto optimal strategy profiles

pure Nash equilibria

mixed Nash equilibria

There are interesting cultural differences in the offers people from different cultures accept and reject when playing The Ultimatum Game.

JOE HENRICH



Henrich, J. (2020). *The WEIRDest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous*. Farrar, Straus and Giroux.

Players

$$N = \{1, 2\}$$

Strategies of player 1

$$\{2-0, 1-1, 0-2\}$$

Strategies of player 2

(yes, yes, yes), (yes, yes, no), (yes, no, yes), (no, yes, yes),
(yes, no, no), (no, yes, no), (no, no, yes), (no, no, no)


Strategy profiles

(2-0, (yes, yes, yes)), (2-0, (yes, yes, no)), ...

Payoffs (aka utilities)

$$u_1(1-1, (\text{yes}, \text{no}, \text{yes})) = 0$$

...



Note that there is a subtlety in the definition of strategies.


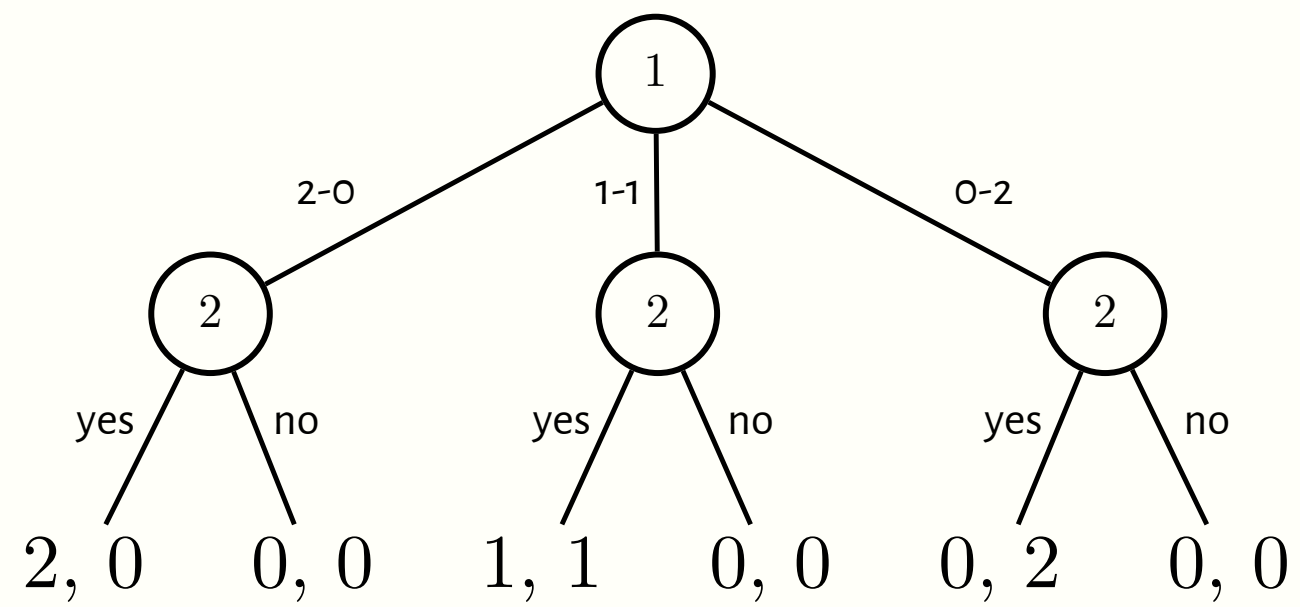
The strategies of each player need to be defined at every choice node of that player.

Even if there is no way to reach that node, given the other choice nodes.

To reason our way through a perfect-information game in extensive form, we just turn it into a *normal-form game*.

To reason our way through a perfect-information game in extensive form, we just turn it into a *normal-form game*.

Yes, we can always do it.



| | yyy | yyn | yny | ynn | nyy | nyn | nny | nnn |
|-------|------|------|------|------|------|------|------|------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

Nash equilibria and everything else is
computed with respect to the
induced normal-form game.


The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.



strictly dominant strategies

?

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.

payoffs

| | yyy | yy \bar{y} | y \bar{y} y | y \bar{y} \bar{y} | \bar{y} yy | \bar{y} y \bar{y} | \bar{y} \bar{y} y | \bar{y} \bar{y} \bar{y} |
|-------|------|--------------|---------------|-----------------------|--------------|-----------------------|-----------------------|-------------------------------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

strictly dominant strategies

?

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.

payoffs

| | yyy | yyN | yNy | yNn | nyy | nyn | nyy | nnn |
|-------|------|------|------|------|------|------|------|------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

strictly dominant strategies

none

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.

payoffs

| | yyy | yy \bar{y} | y \bar{y} y | y \bar{y} \bar{y} | \bar{y} yy | \bar{y} y \bar{y} | \bar{y} \bar{y} y | \bar{y} \bar{y} \bar{y} |
|-------|------|--------------|---------------|-----------------------|--------------|-----------------------|-----------------------|-------------------------------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

strictly dominant strategies

none

Pareto optimal strategy profiles

everything except (0, 0)

pure Nash equilibria

?

mixed Nash equilibria

?

The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.

payoffs

| | yyy | yy \bar{y} | y \bar{y} y | y \bar{y} \bar{y} | \bar{y} yy | \bar{y} y \bar{y} | \bar{y} \bar{y} y | \bar{y} \bar{y} \bar{y} |
|-------|------|--------------|---------------|-----------------------|--------------|-----------------------|-----------------------|-------------------------------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

strictly dominant strategies

none

Pareto optimal strategy profiles

everything except (0, 0)

pure Nash equilibria

see above

mixed Nash equilibria

?

The Ultimatum Game

Player 1 has two euros, which it has to divide between themselves and player 2.

Player 1 makes an offer, which player 2 can *accept* or *reject*.

If player 2 accepts, money is divided according to player 1's offer.

If player 2 rejects, no one gets anything.

payoffs

| | yyy | yyN | yNy | yNn | nyy | nyn | nyy | nnn |
|-------|------|------|------|------|------|------|------|------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

strictly dominant strategies

none

Pareto optimal strategy profiles

everything except (0, 0)


pure Nash equilibria

see above

mixed Nash equilibria


too lazy to figure out

What makes (2-0, nnn) a Nash equilibrium depends crucially on what Player 2 does at *all* nodes: including ‘irrelevant’ ones.



Think: why does Player 1 not want to deviate?

Because Player 2 always says *no*, so there's no point!



| | yyy | yyn | yny | ynn | nyy | nyn | nny | nnn |
|-------|------|------|------|------|------|------|------|------|
| (2-0) | 2, 0 | 2, 0 | 2, 0 | 2, 0 | 0, 0 | 0, 0 | 0, 0 | 0, 0 |
| (1-1) | 1, 1 | 1, 1 | 0, 0 | 0, 0 | 1, 1 | 1, 1 | 0, 0 | 0, 0 |
| (0-2) | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 | 0, 2 | 0, 0 |

Games in extensive form afford a refinement of Nash equilibria: *subgame perfect equilibria*.


These involve playing a Nash equilibrium at every node of the game.

A subgame perfect equilibrium can be found by *backward induction*.

A subgame perfect equilibrium can be found by *backward induction*.

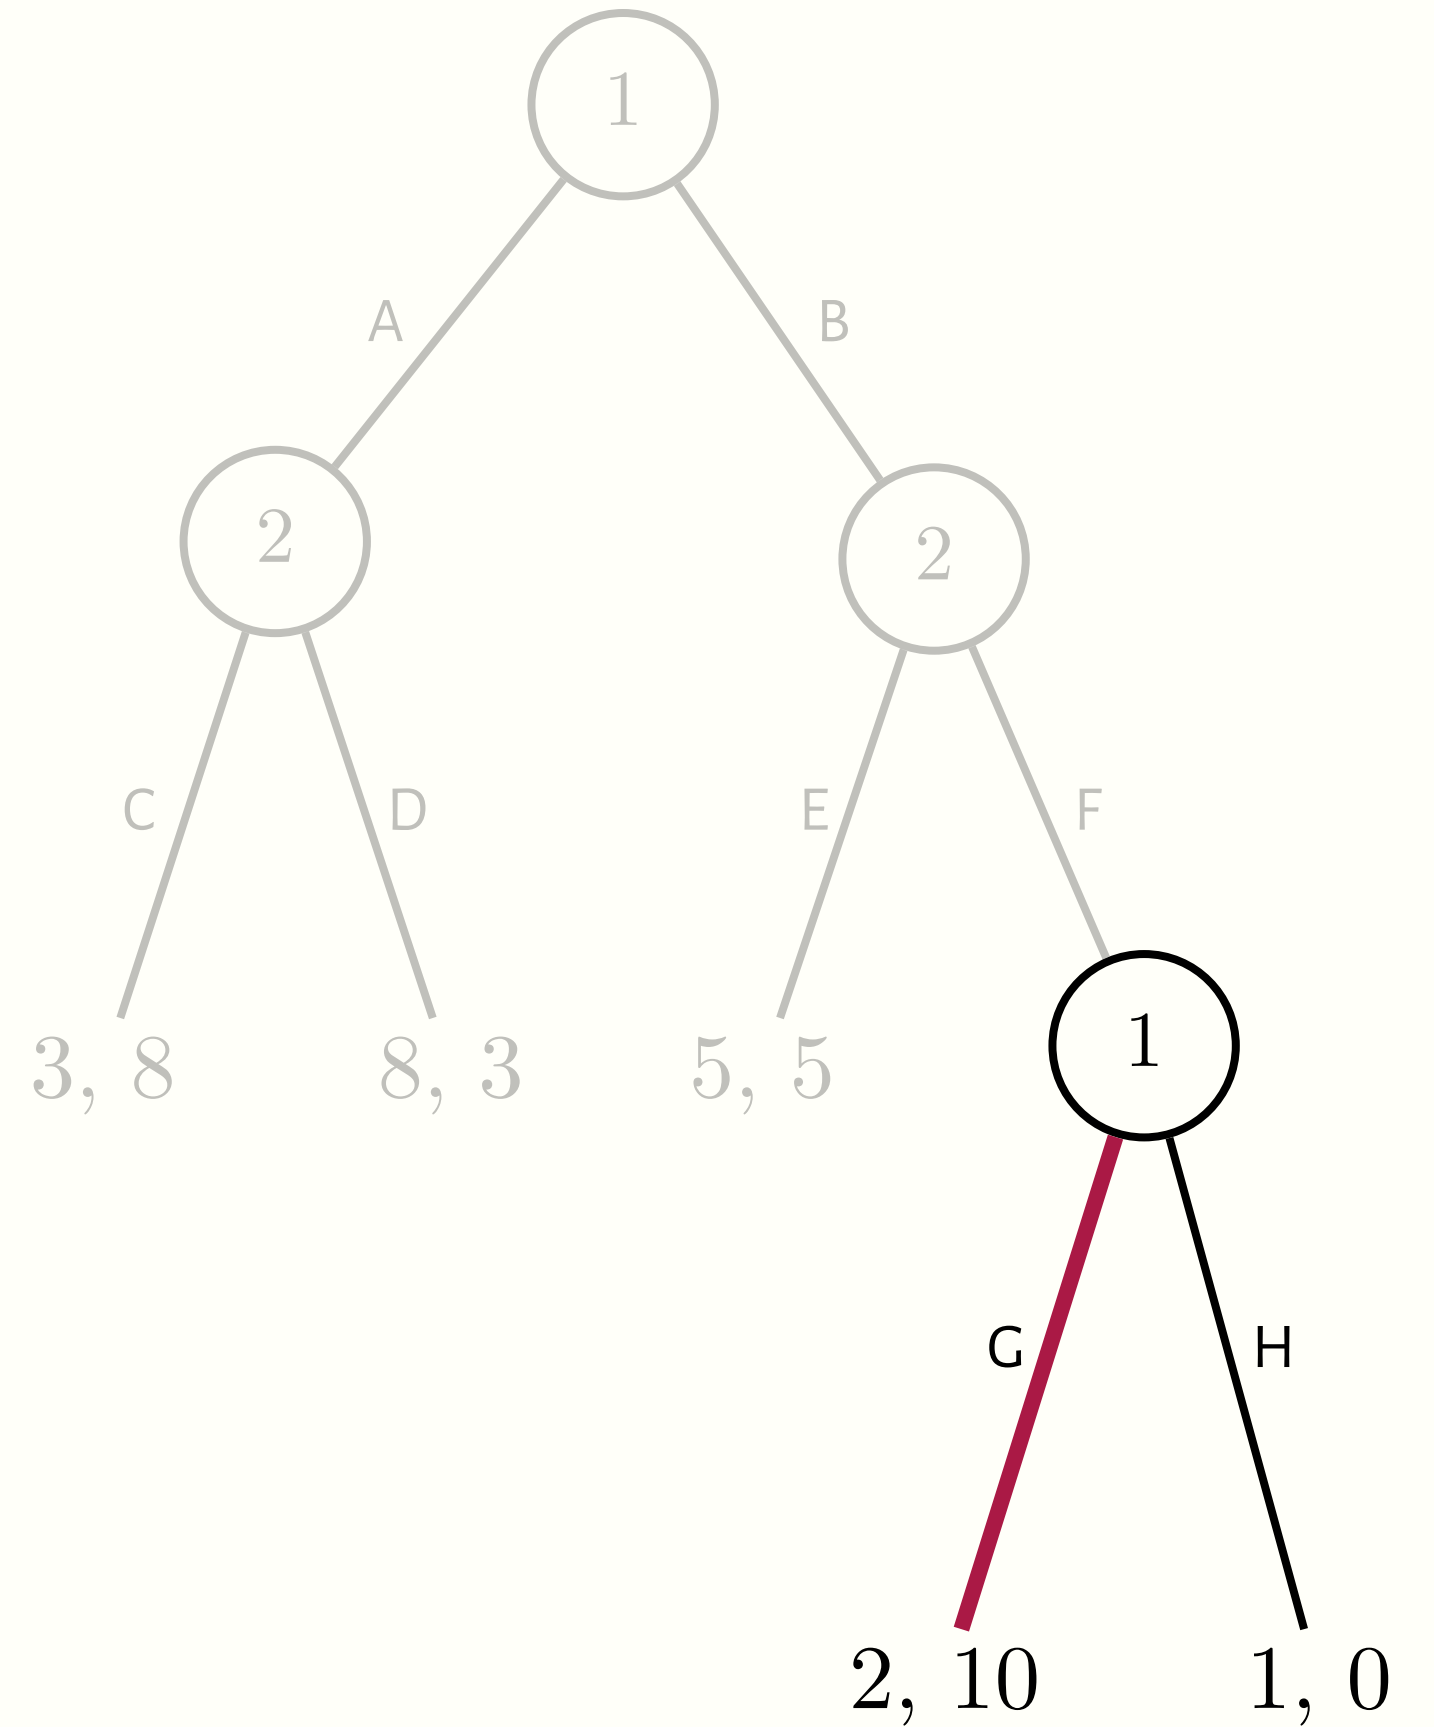
We reason backwards, from the end stages of a game, by finding the optimal action at every intermediate step.

Backward Induction: An Example



Backward Induction: An Example


Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).



Backward Induction: An Example

Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).




Backward Induction: An Example

Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

We infer that player 2 chooses F here.




Backward Induction: An Example

Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

We infer that player 2 chooses F here.

Which further means that player 1 sees a payoff of 2 if they go down this path.



Backward Induction: An Example

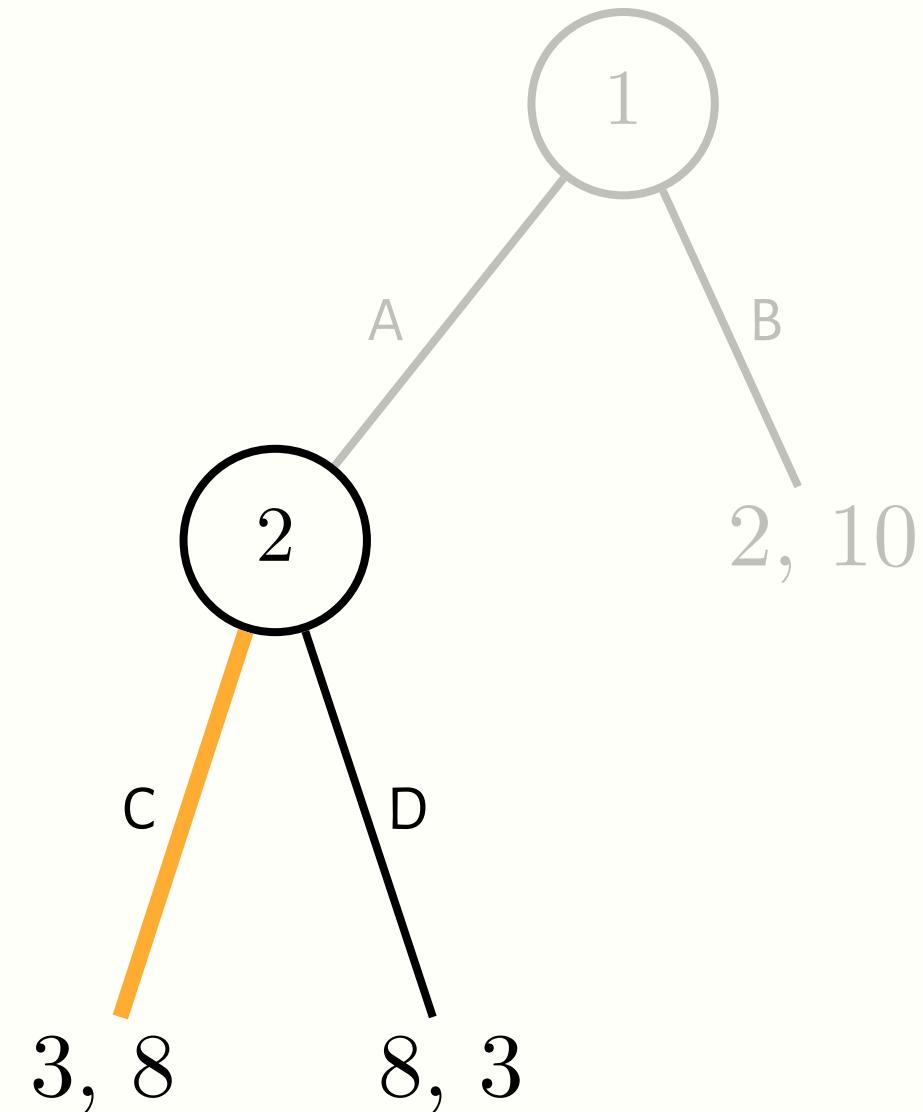
Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

We infer that player 2 chooses F here.

Which further means that player 1 sees a payoff of 2 if they go down this path.

On the other branch player 2 chooses C.



Backward Induction: An Example


Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

We infer that player 2 chooses F here.

Which further means that player 1 sees a payoff of 2 if they go down this path.

On the other branch player 2 chooses C.



Backward Induction: An Example

Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

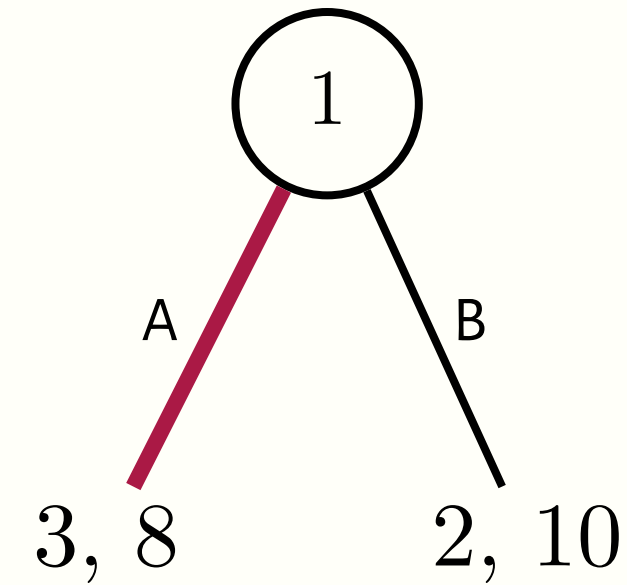
Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

We infer that player 2 chooses F here.

Which further means that player 1 sees a payoff of 2 if they go down this path.

On the other branch player 2 chooses C.

Which means player 1 chooses A.



Backward Induction: An Example

Faced with a choice only between G and H, player 1 surely chooses G, leading to a payoff of (2, 10).

Player 2 takes that into account when making their own decision one step earlier, i.e., they know that choosing F leads to a payoff of (2, 10).

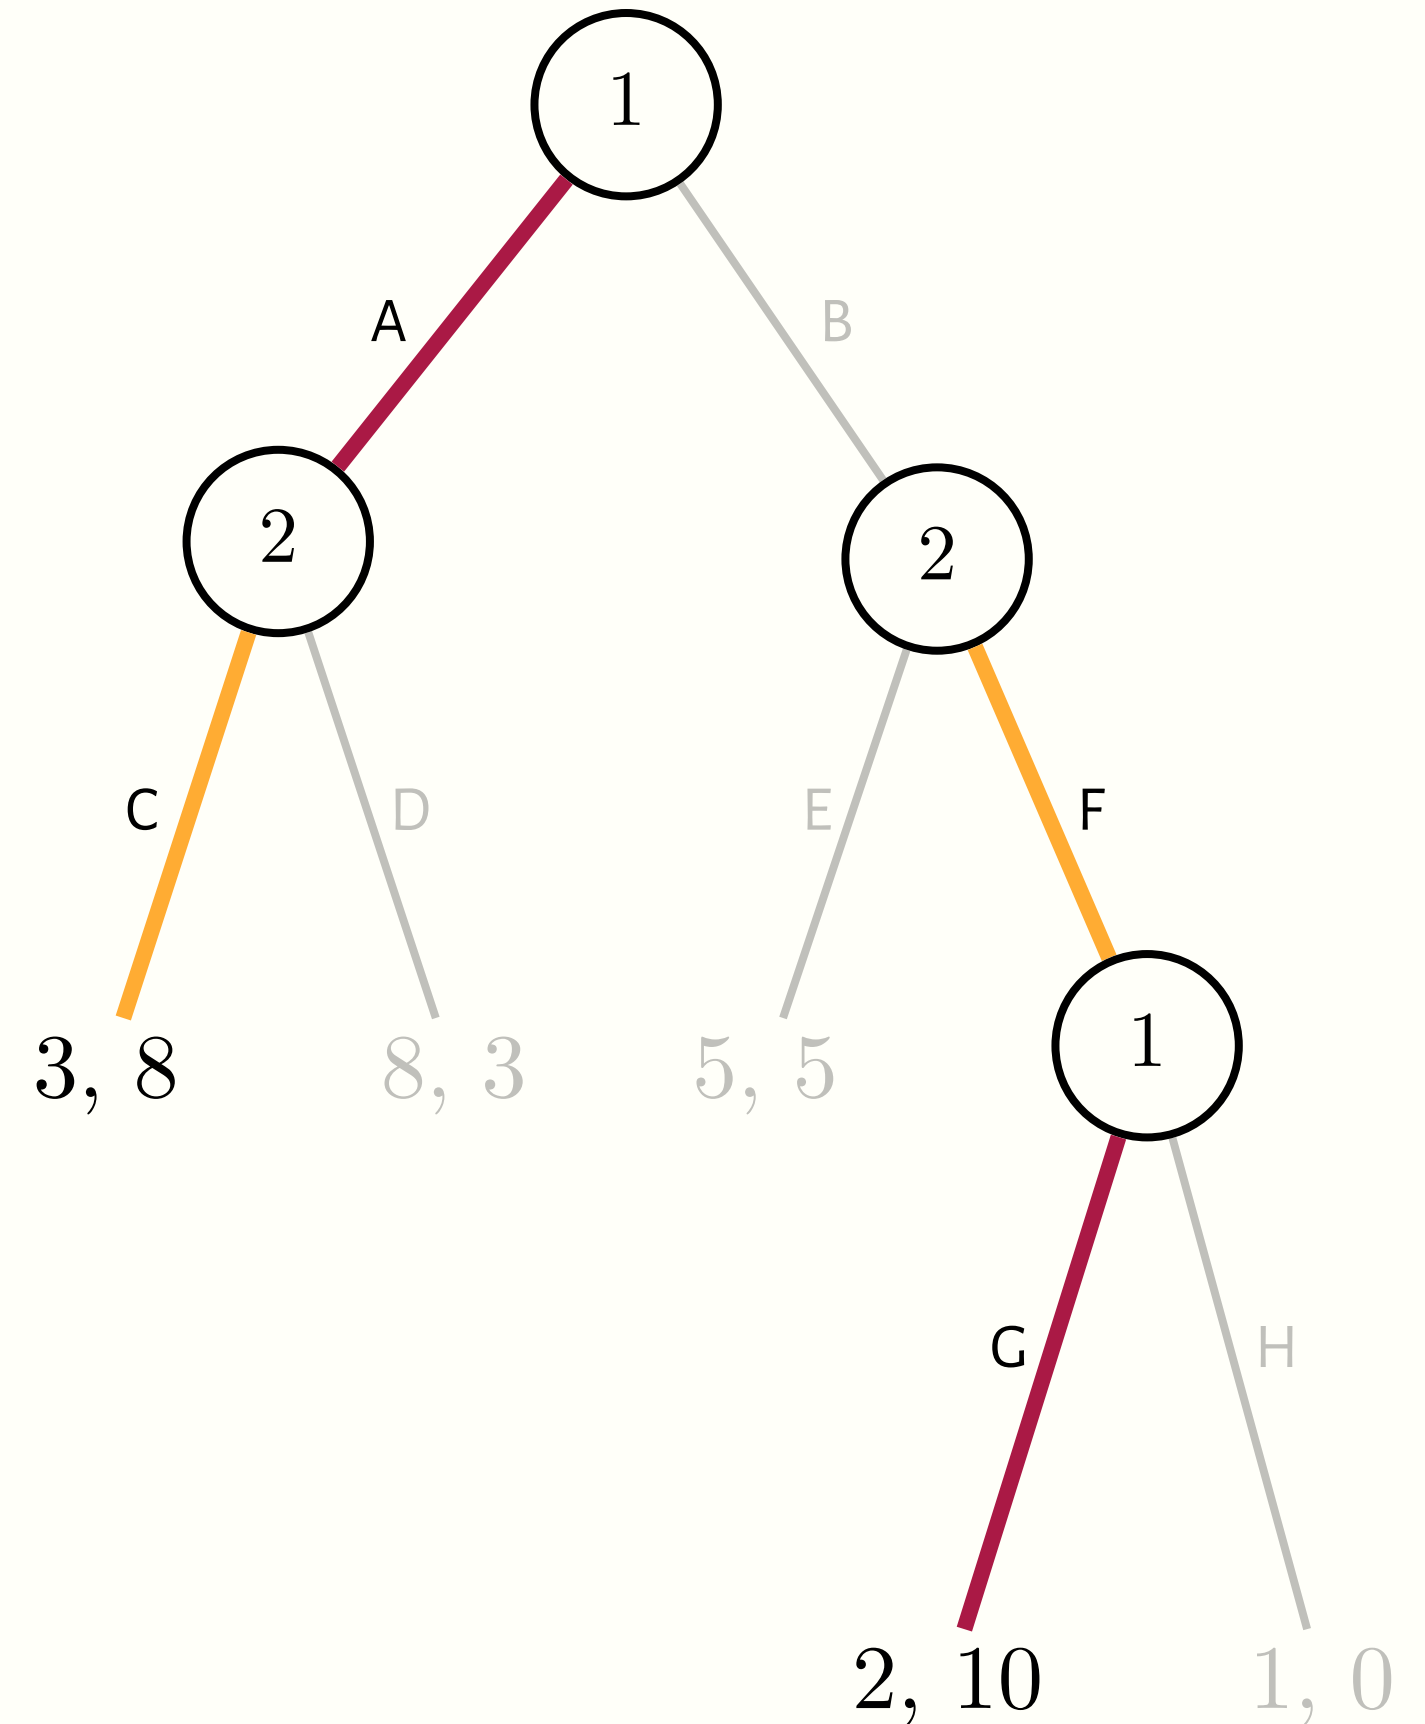
We infer that player 2 chooses F here.

Which further means that player 1 sees a payoff of 2 if they go down this path.

On the other branch player 2 chooses C.

Which means player 1 chooses A.

After which we can just read off the subgame-perfect equilibrium: ((A, G), (C, F)).



Backward induction is well-defined and terminates, if the game tree is finite.

So what have we shown?

THEOREM (SELTEN, 1965)

Every finite extensive-form game has at least one subgame-perfect equilibrium.

Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertraegheit. *Zeitschrift fuer die Gesamte Staatswissenschaft*, 121(2):301–324.

THEOREM (ZERMELO, 1913)

Every finite extensive-form game has at least one pure Nash Equilibrium.

Zermelo, E. (1913). Uber eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. *Proceedings of the 5th International Congress of Mathematicians*.

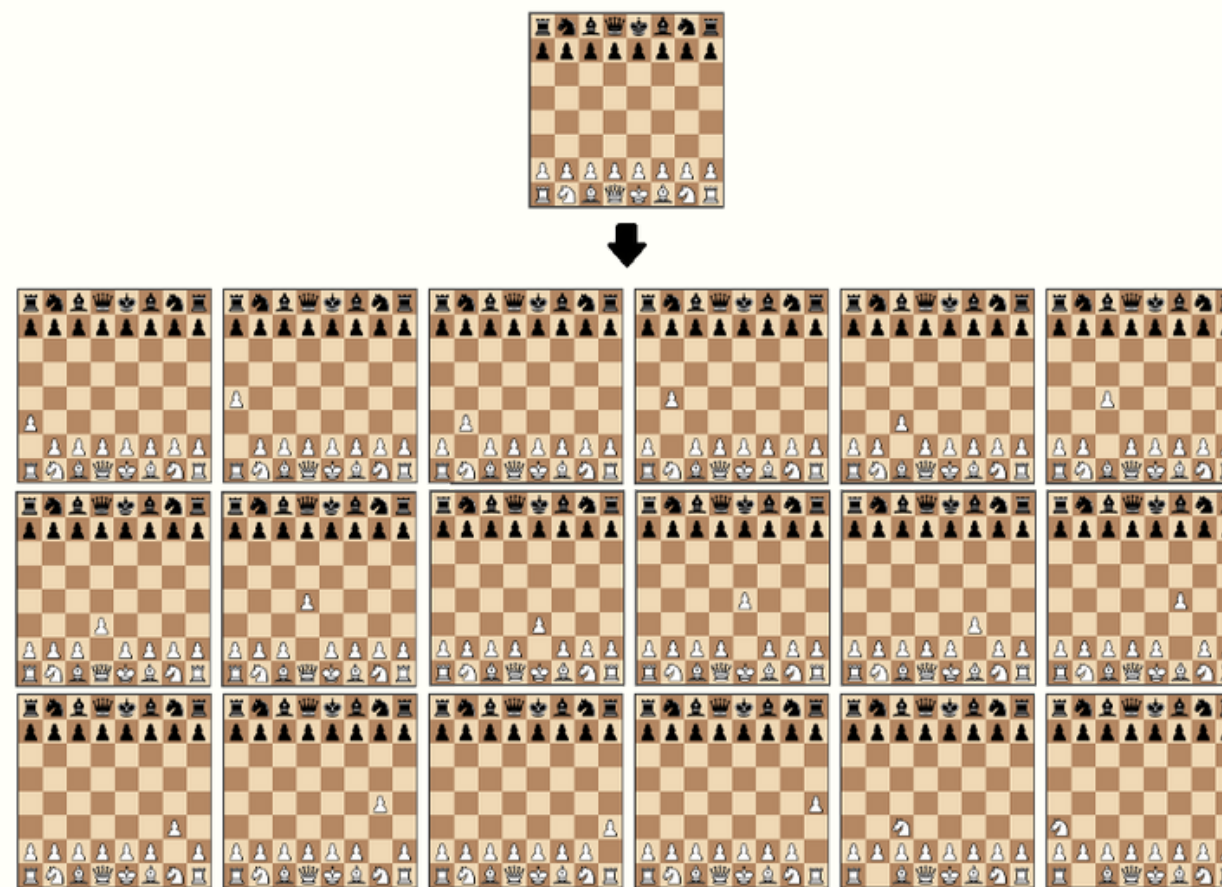
ERNST ZERMELO




I arrived at these ideas while thinking about whether chess is determined, i.e., whether either white or black has a winning strategy, or can force a draw.

Which is true if we can bound the length of a game.

At the same time, the game tree of chess is too large to actually survey the strategies, let alone represent it explicitly.



alpha-beta pruning to disregard parts of the tree



With extensive-form games, we can go even a bit more general.

With extensive-form games, we can go even a bit more general.

In perfect-information games players know what actions were played at previous rounds.

And thus, what nodes they are in.

With extensive-form games, we can go even a bit more general.

In perfect-information games players know what actions were played at previous rounds.

And thus, what nodes they are in.


But in many other situations, players have only partial knowledge.

Enter extensive-form games
with *imperfect* information.

We represent an agent's uncertainty over what choice node they're at by an *information set*.

Adding Uncertainty: A Dashed Line

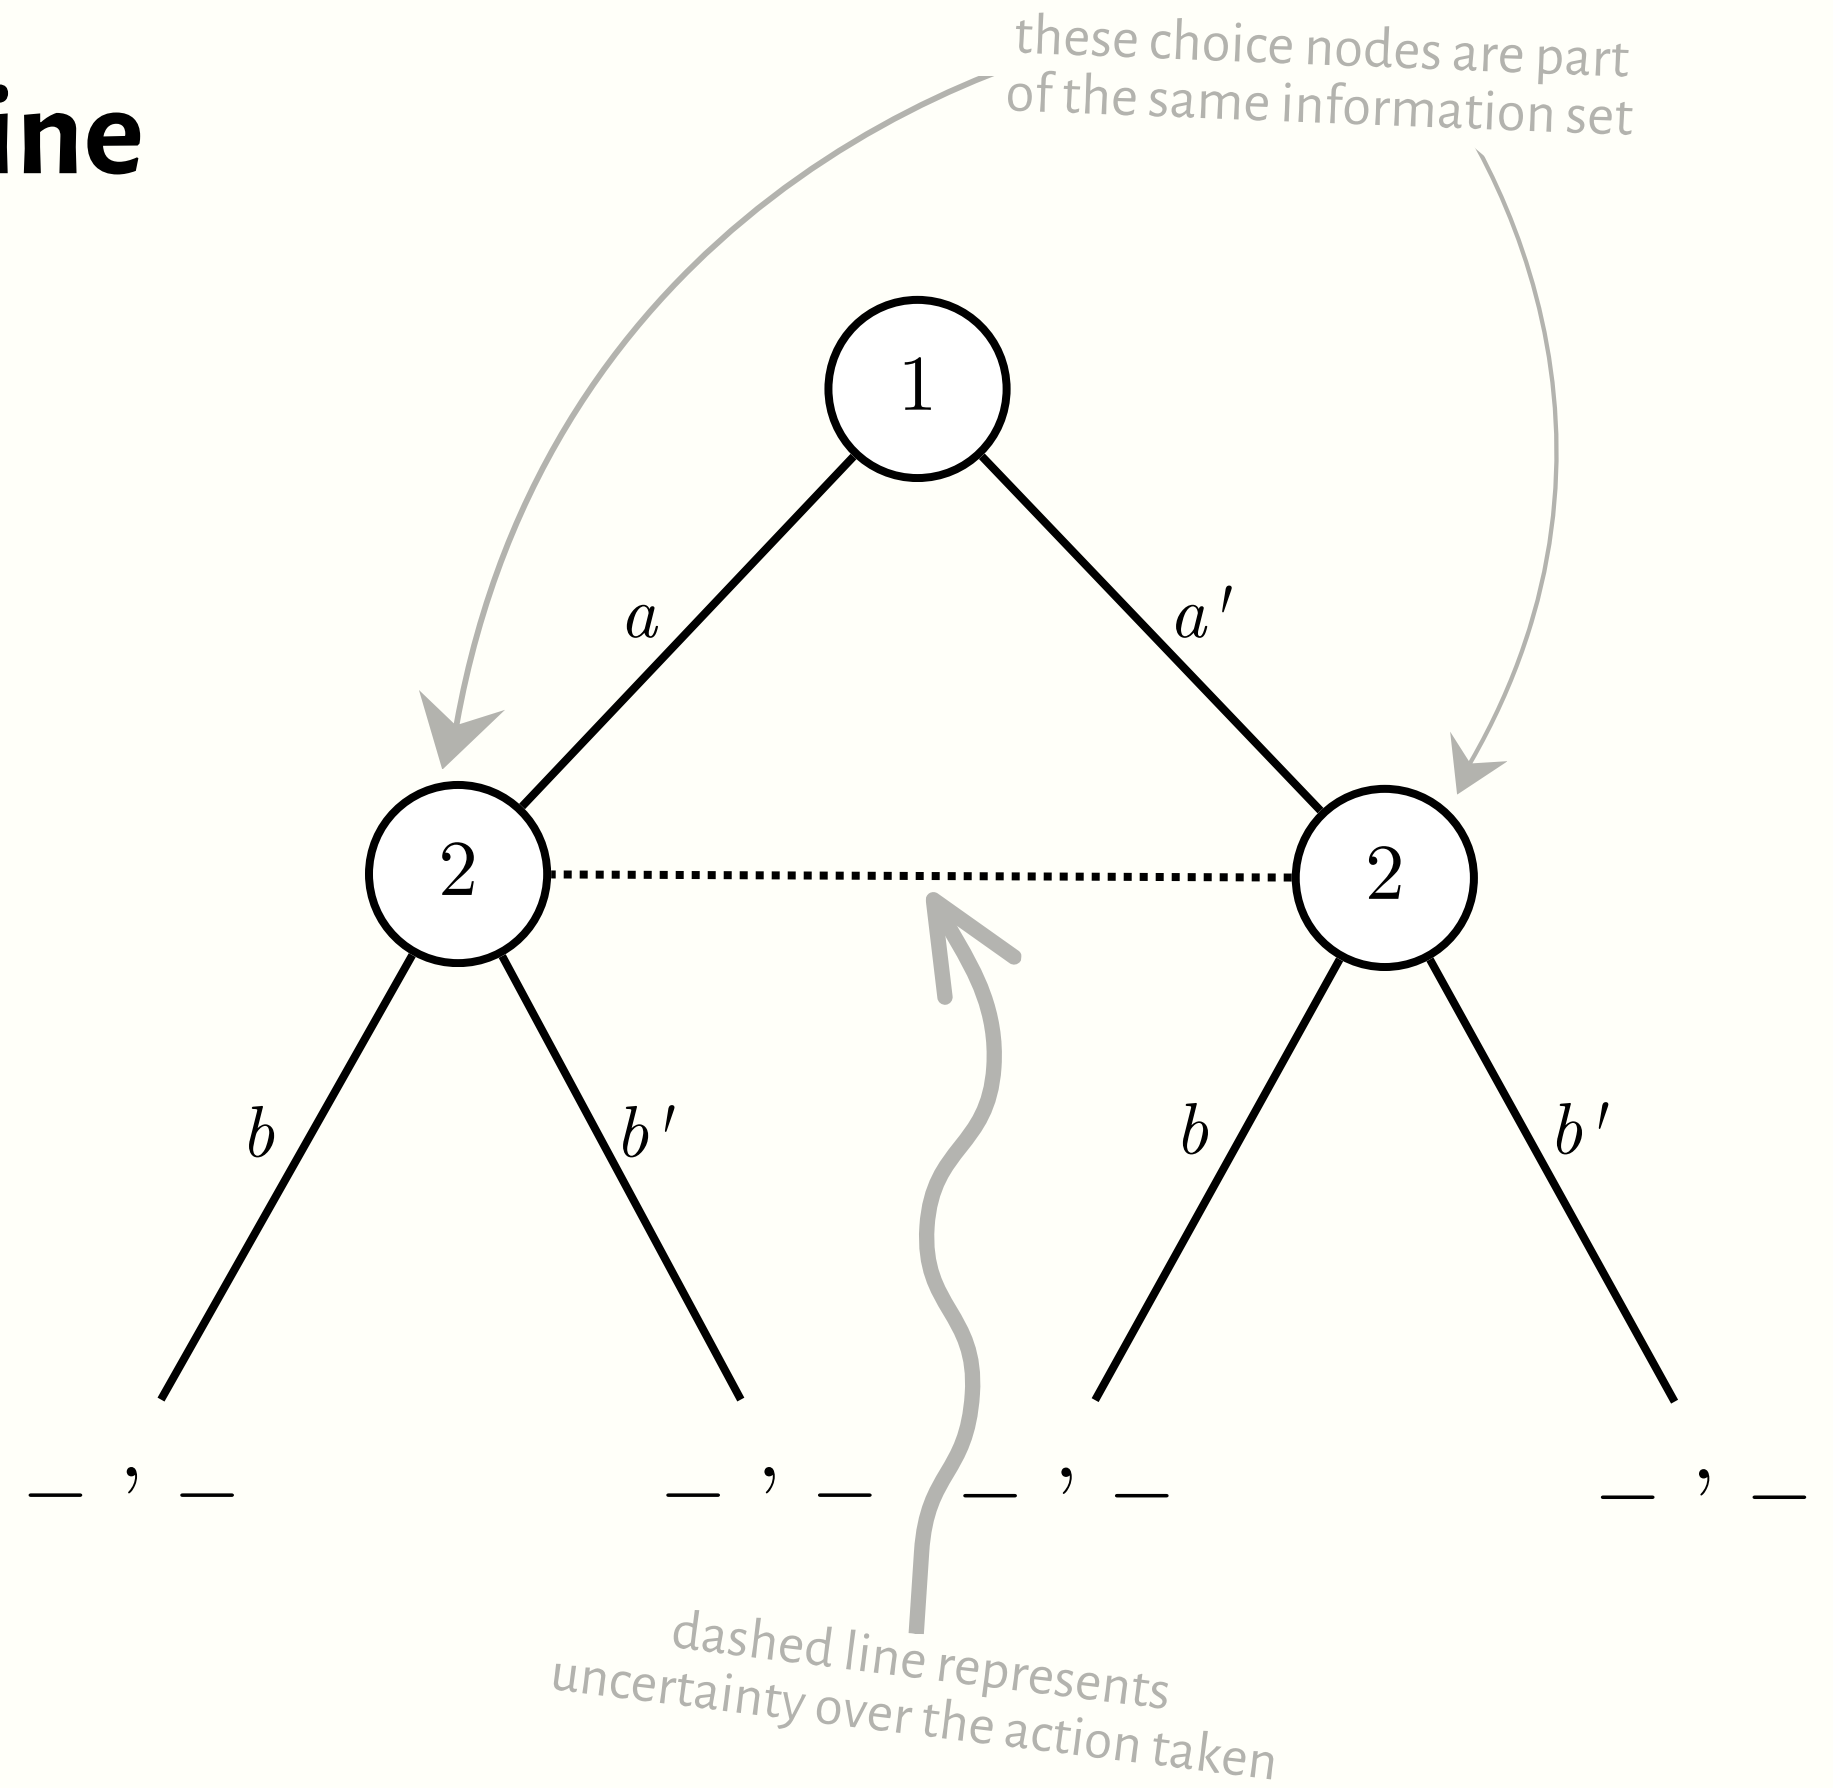
Player 1 takes an action: a or a' .



Adding Uncertainty: A Dashed Line

Player 1 takes an action: a or a' .

Player 2 follows up, *not* knowing what action Player 1 has actually taken: the two nodes connected by a dashed line are in the same information set.



Adding Uncertainty: A Dashed Line

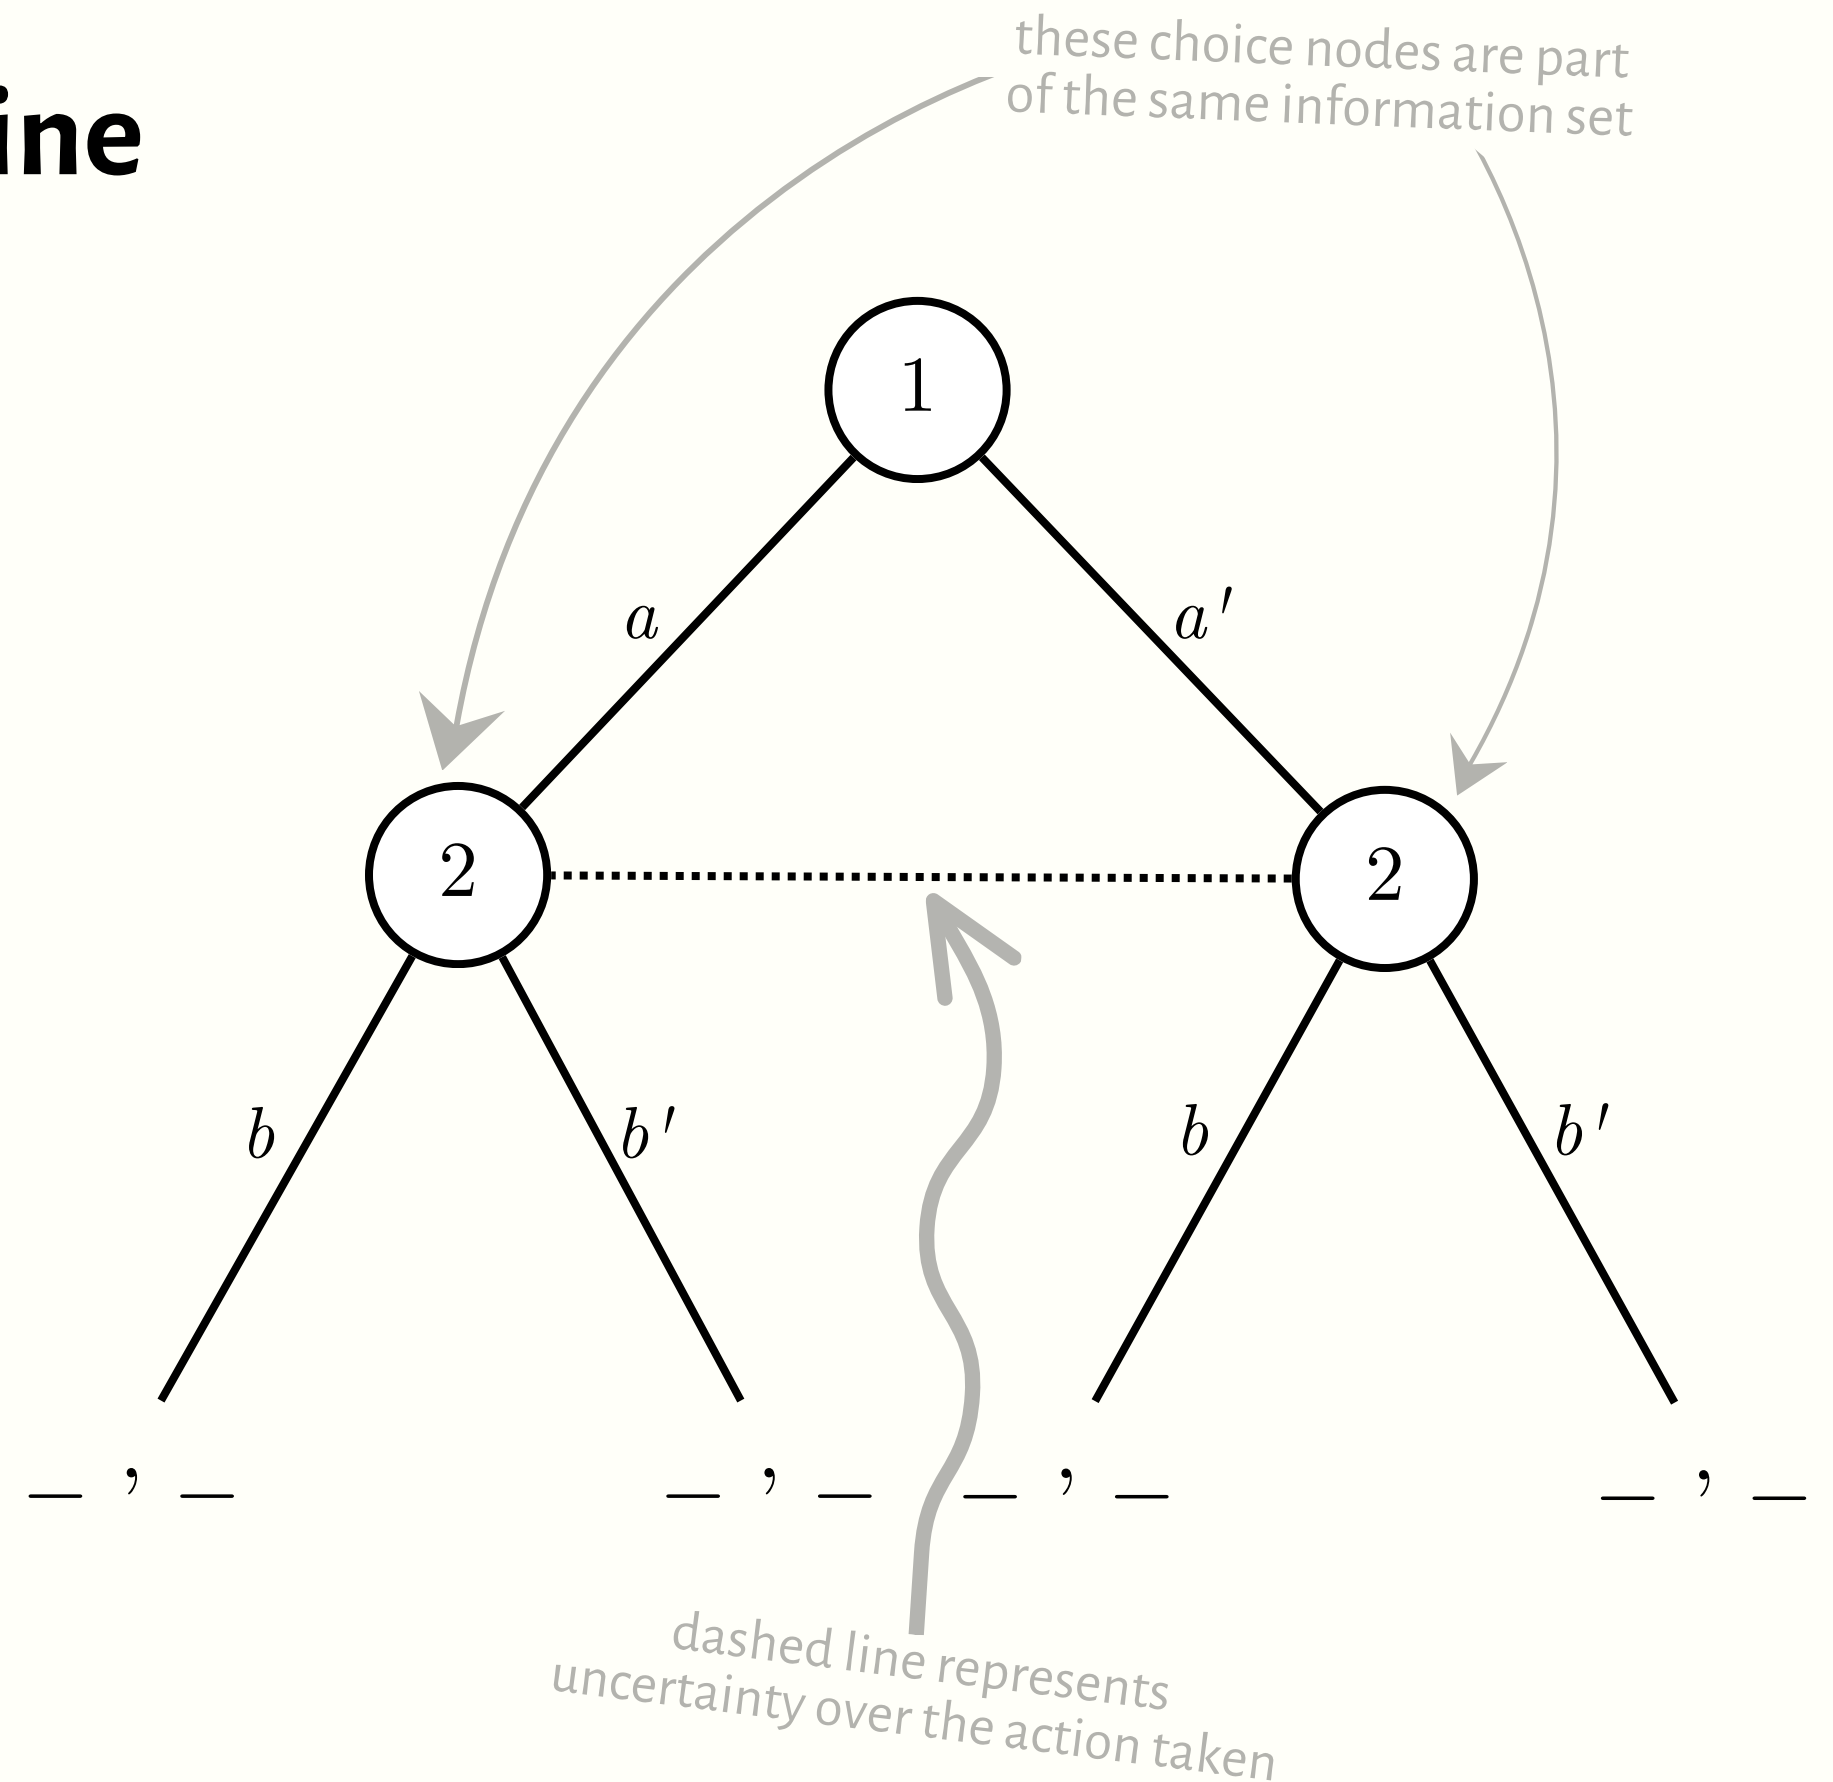
Player 1 takes an action: a or a' .

Player 2 follows up, *not* knowing what action Player 1 has actually taken: the two nodes connected by a dashed line are in the same information set.

Payoffs are specific to the branch taken.

Players know the actions available to all players, and the payoffs corresponding to each sequence of actions, i.e., the structure of the game.

But *do not know* which node from a particular information set they're in.



Intuitively, an agent cannot distinguish between the actions in one of their information sets.

Like their perfect-information counterparts, extensive-form games with perfect information are modeled as *trees*.

Like their perfect-information counterparts, extensive-form games with perfect information are modeled as *trees*.

The main difference is that every agent's choice nodes are partitioned into *information sets*.

Like their perfect-information counterparts, extensive-form games with perfect information are modeled as *trees*.

The main difference is that every agent's choice nodes are partitioned into *information sets*.

With the added proviso that the actions available at every information set are the *same* for all actions in that set.

Like their perfect-information counterparts, extensive-form games with perfect information are modeled as *trees*.

The main difference is that every agent's choice nodes are partitioned into *information sets*.

With the added proviso that the actions available at every information set are the *same* for all actions in that set.

A *strategy* for an agent is a combination of actions, one for each information set corresponding to that agent.

rather than for each choice node

Players

$N = \{1, 2\}$

Information sets of Player 1

$\{1a\}, \{1b, 1c\}$

Information sets of Player 2

$\{2\}$

Strategies of player 1

$(L, X), (L, Y), (R, X), (R, Y)$

Strategies of player 2

A, B


Strategy profiles

$((L, X), A), ((L, X), B), \dots$

Payoffs (aka utilities)



you can figure this out

not $(L, X, X), (L, X, Y), \dots$, which would be the case with perfect information




Now we can finally get back to
the Prisoner's Dilemma!


| | Cooperate | Defect |
|-----------|-----------|--------|
| Cooperate | 2, 2 | 0, 3 |
| Defect | 3, 0 | 1, 1 |



no



Cooperate



Cooperate

2, 2

Defect


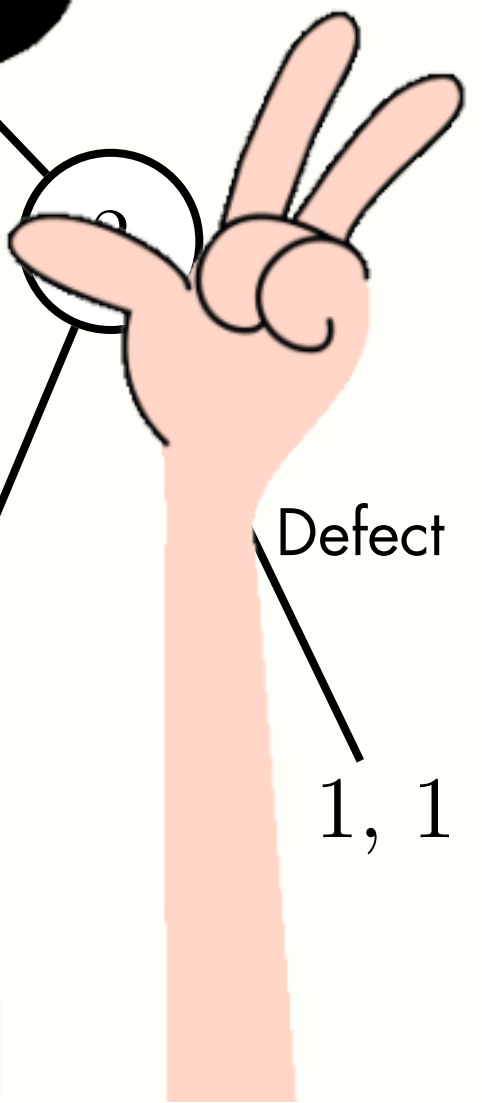
Cooperate

Defect

Cooperate

Defect

1, 1



| | Cooperate | Defect |
|-----------|-----------|--------|
| Cooperate | 2, 2 | 0, 3 |
| Defect | 3, 0 | 1, 1 |

Cooperate

Cooperate

Defect

2, 2

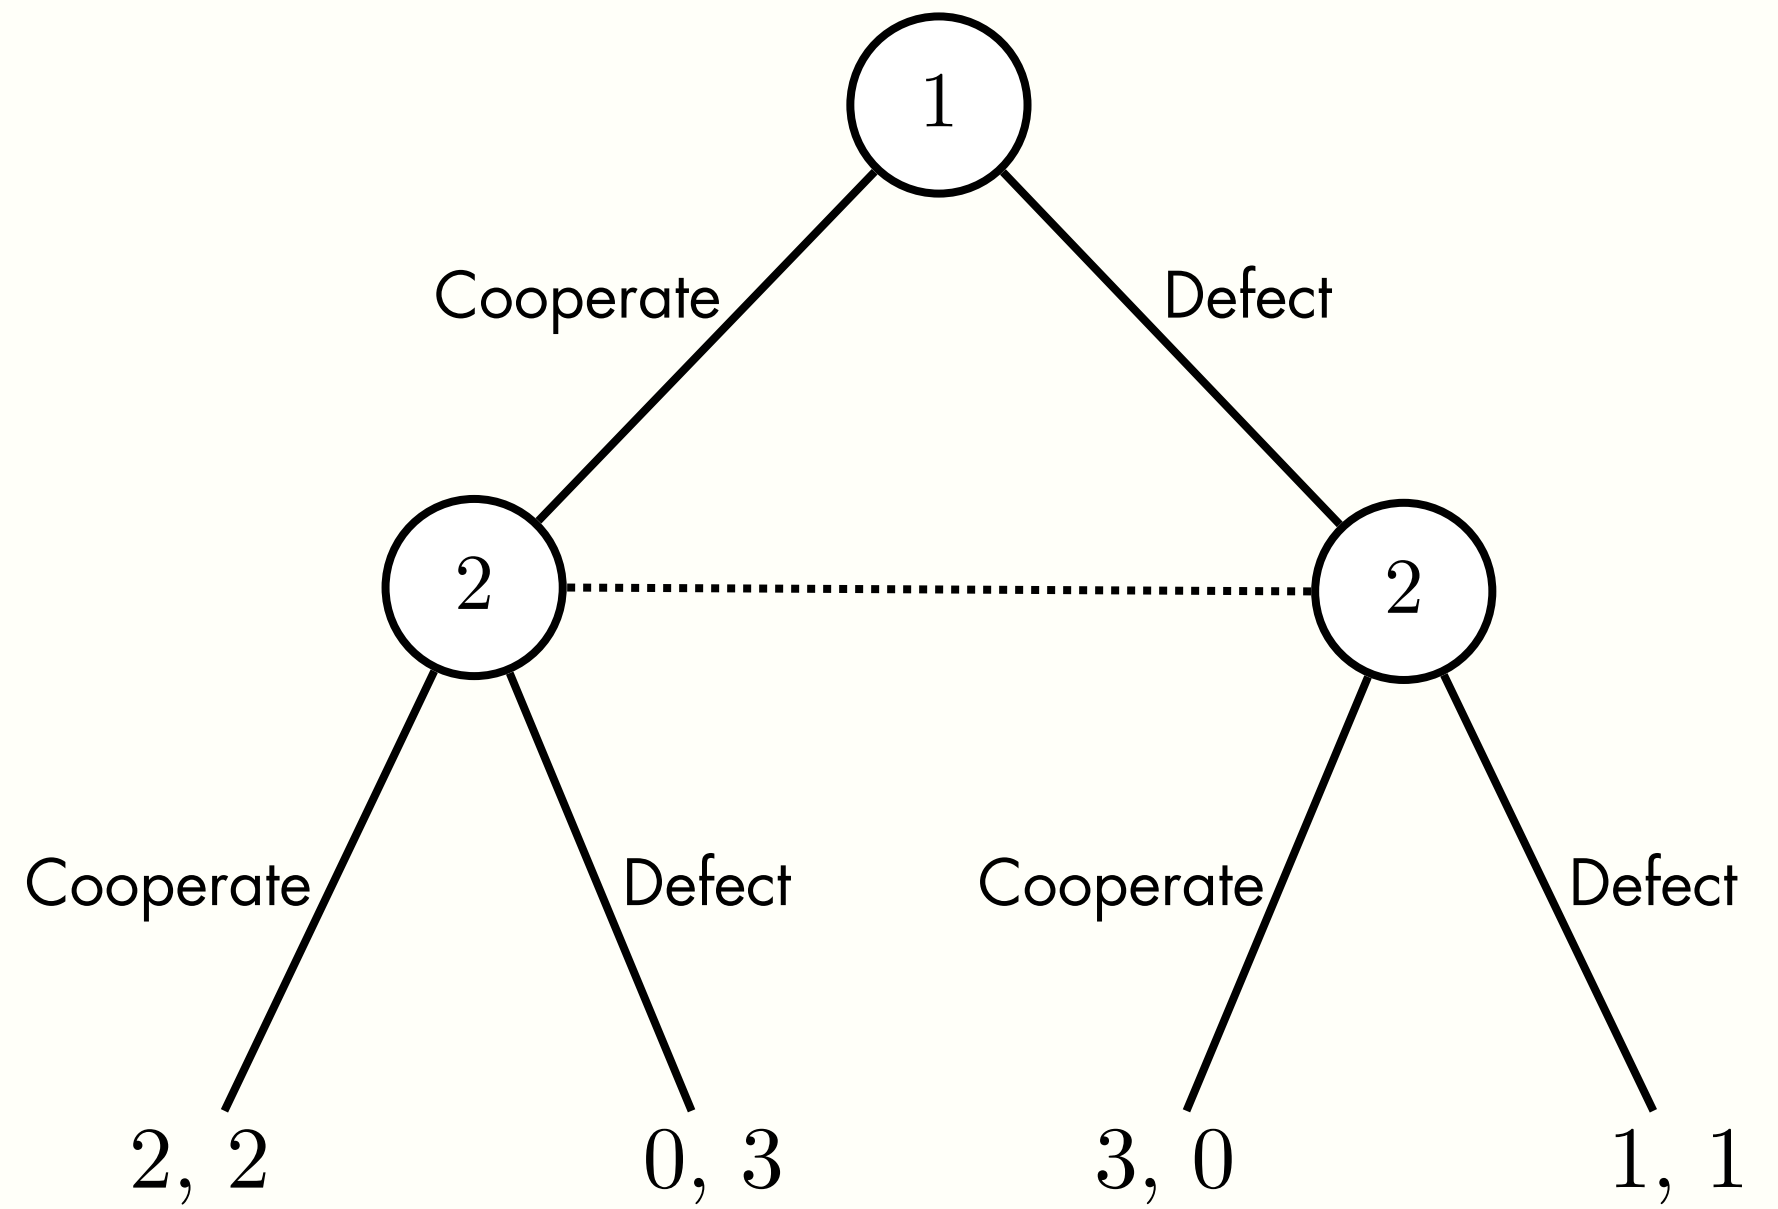

0, 3

Defect

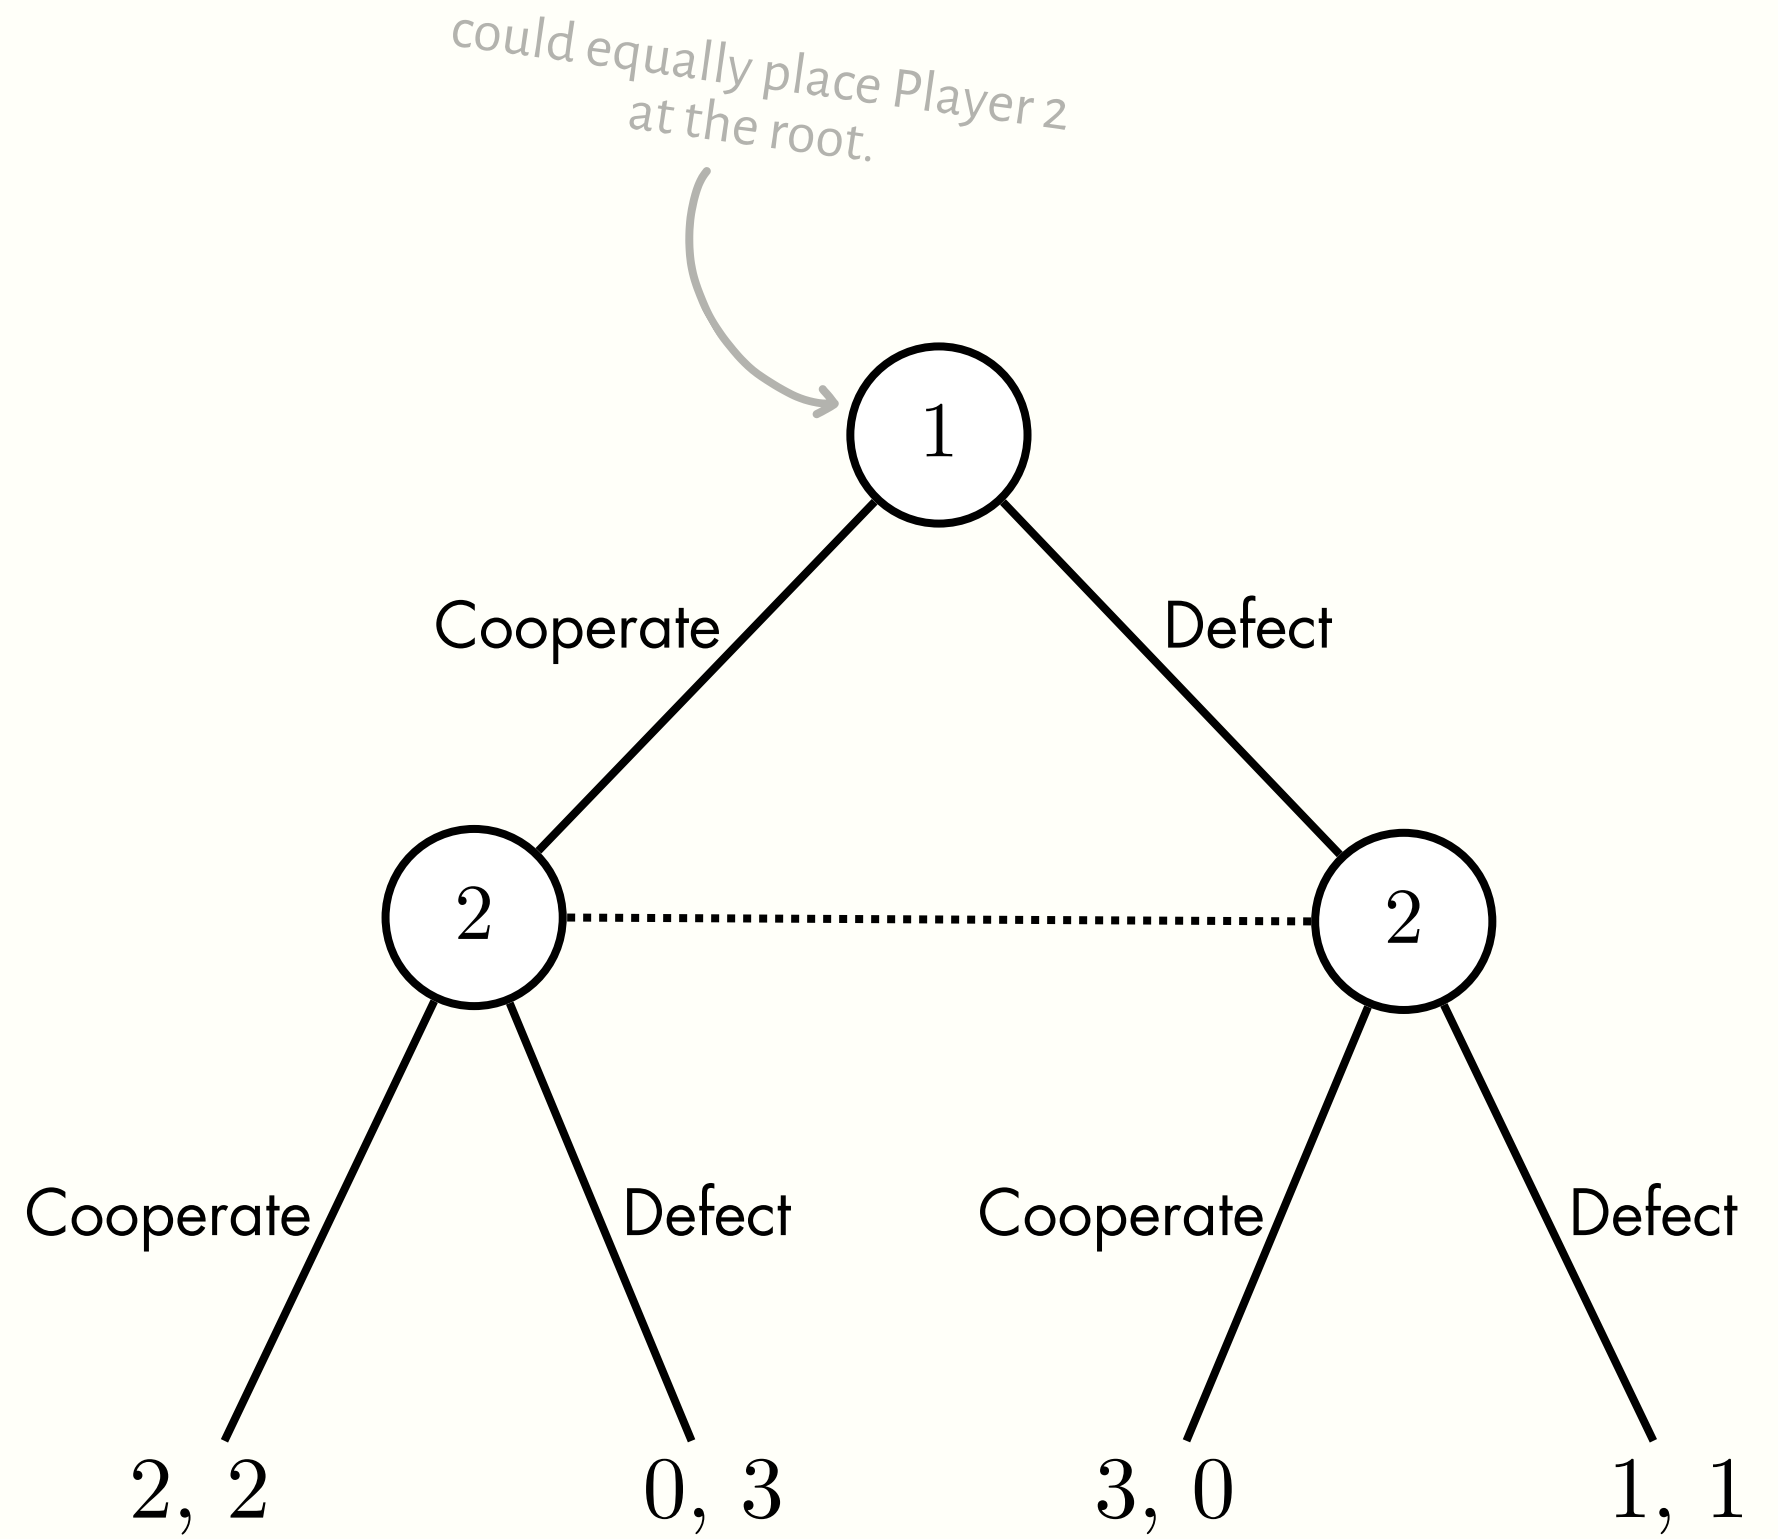

3, 0

1, 1

| | Cooperate | Defect |
|-----------|-----------|--------|
| Cooperate | 2, 2 | 0, 3 |
| Defect | 3, 0 | 1, 1 |



| | Cooperate | Defect |
|-----------|-----------|--------|
| Cooperate | 2, 2 | 0, 3 |
| Defect | 3, 0 | 1, 1 |



Note that we can't model the Prisoner's Dilemma as an extensive-form game with *perfect* information.

Because, well, players don't have perfect information.

But we *can* model it as a game of imperfect information.

Not only that, but now we
can even model the iterated
Prisoner's Dilemma!

A finite number of rounds.

A finite number of rounds.

Like, say, two.

Iterated Prisoner's Dilemma

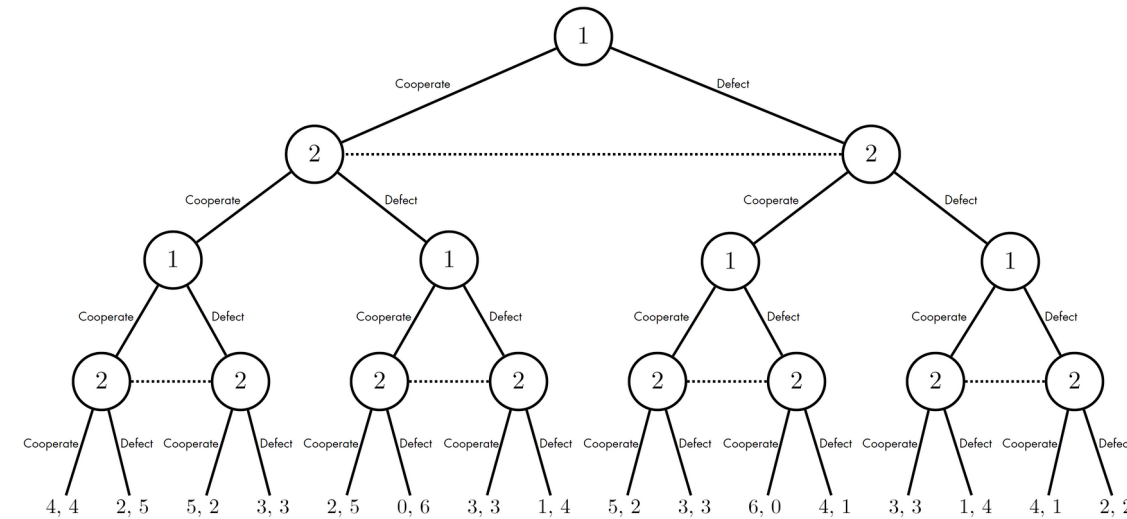


2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs



strictly dominant strategies

?

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

Iterated Prisoner's Dilemma

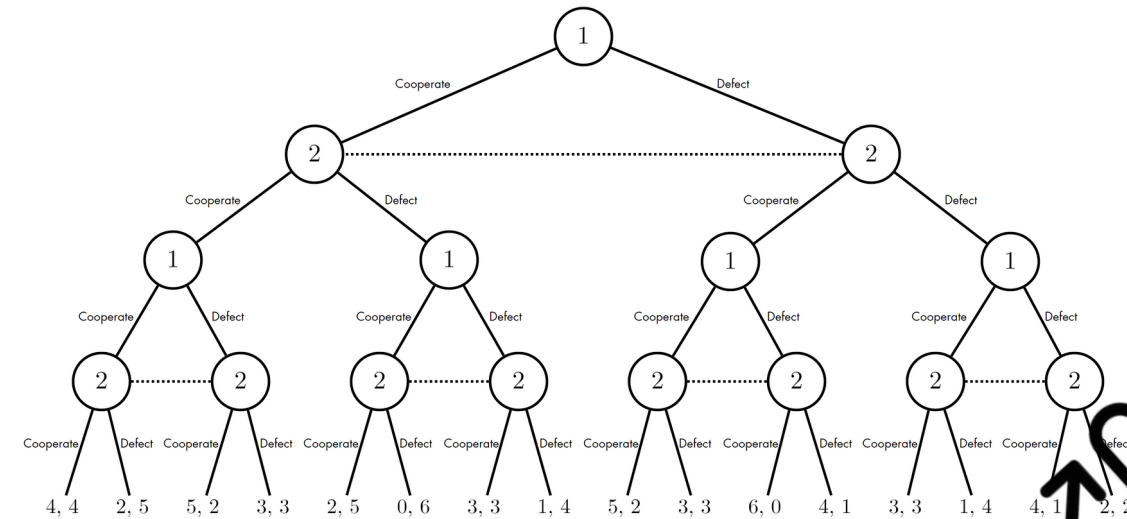


2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs



strictly dominant strategies

?

Pareto optimal strategy profiles

?

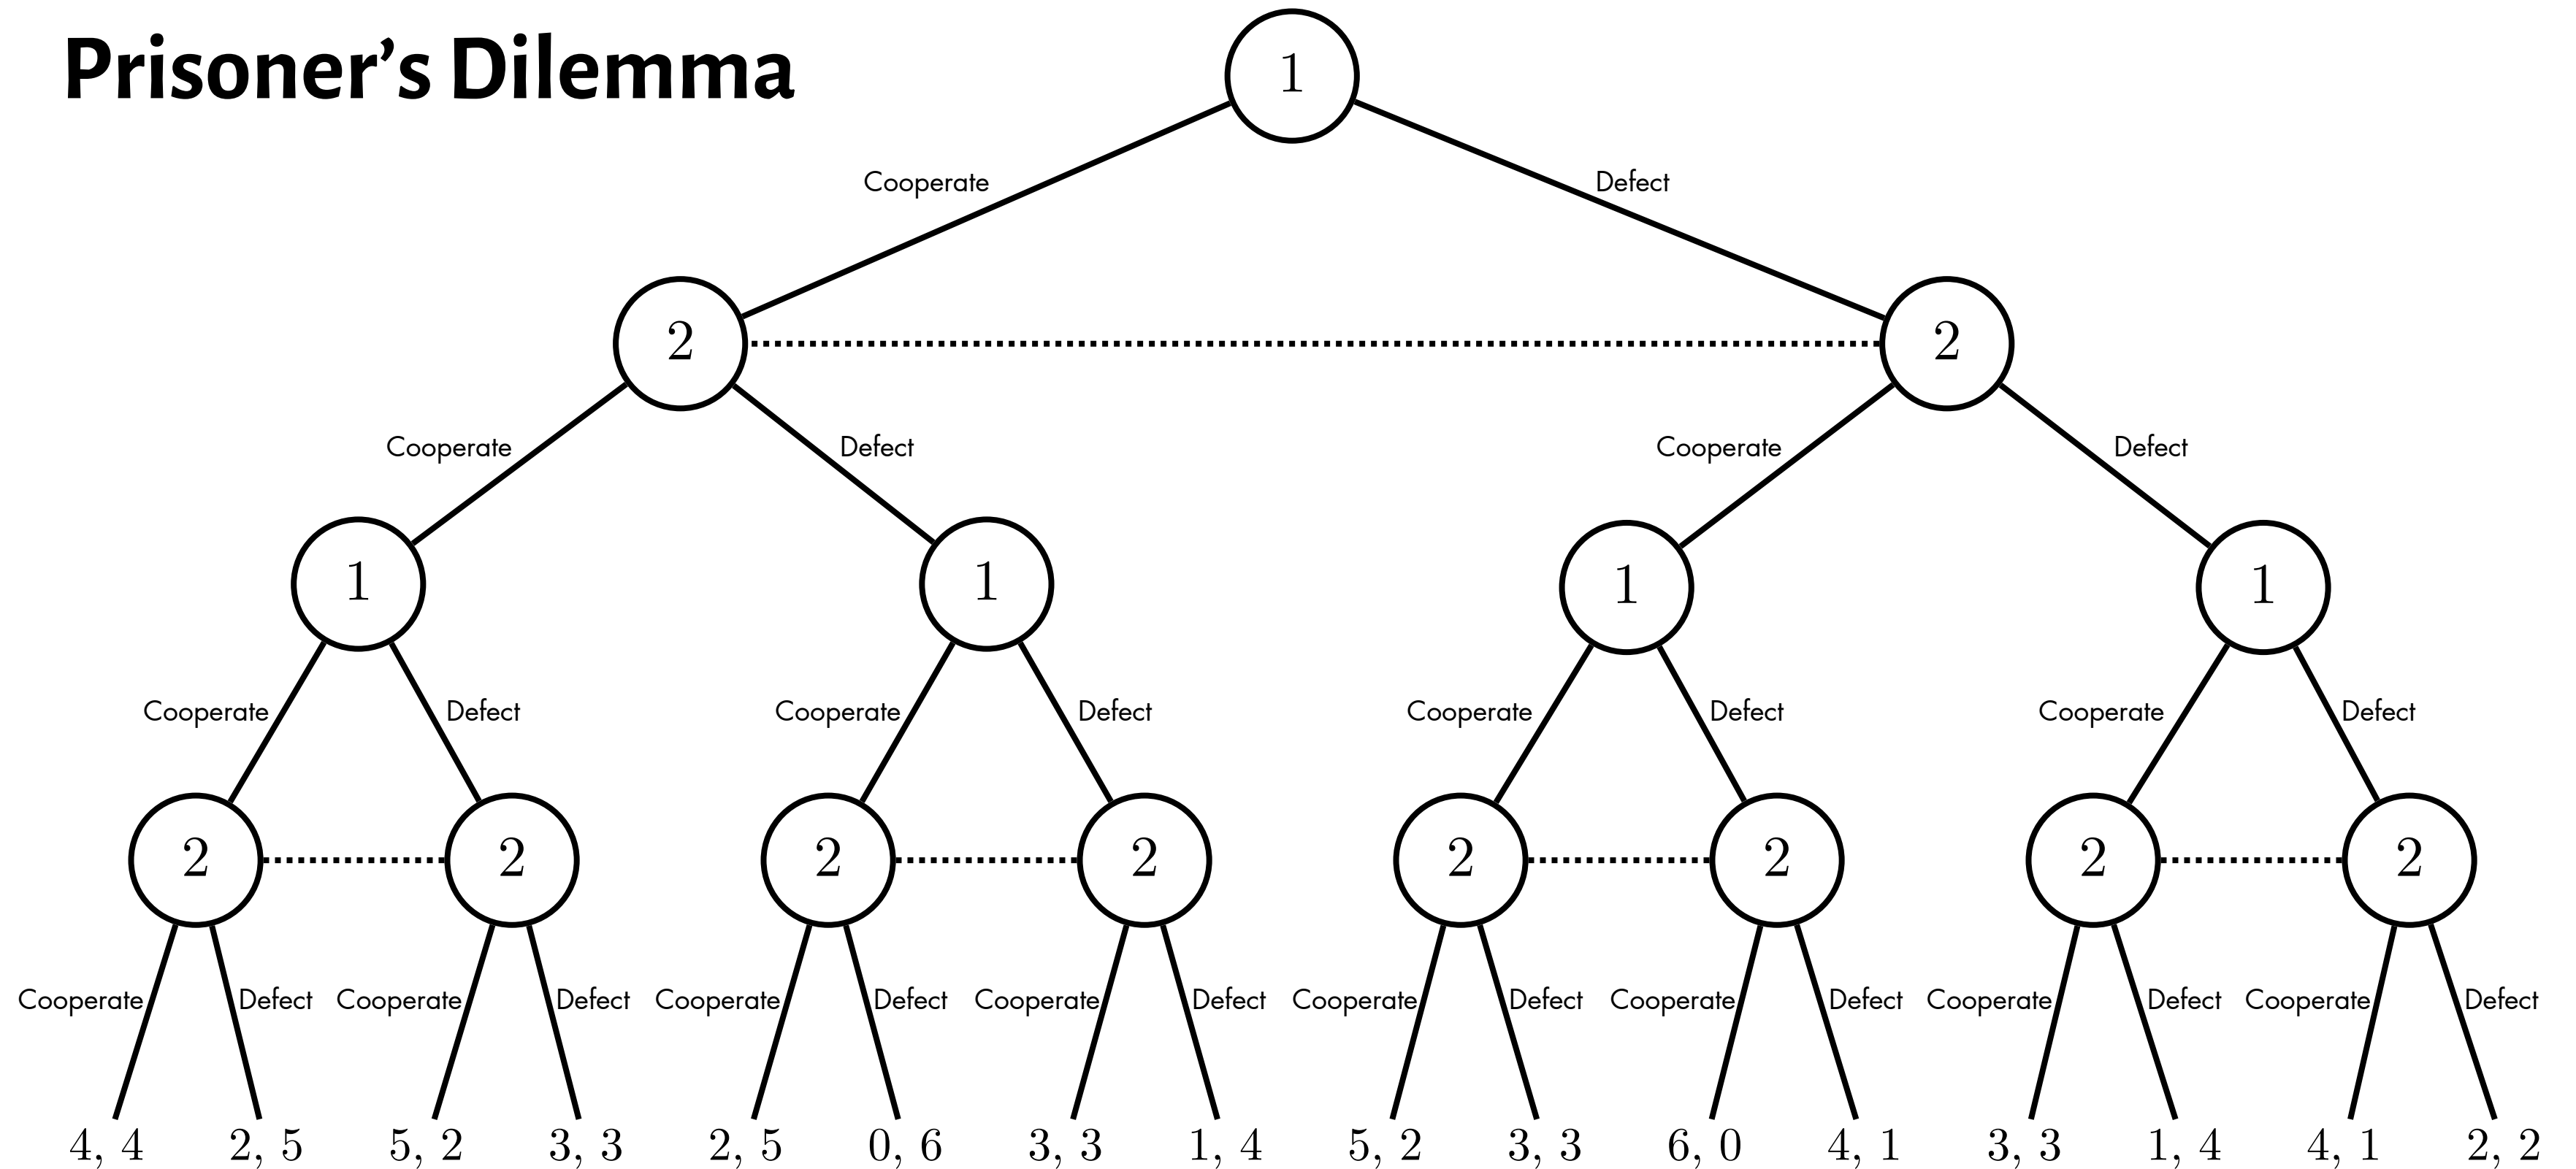
pure Nash equilibria

?

mixed Nash equilibria

?

Two Rounds of the Prisoner's Dilemma



Note that players know actions taken at *previous* rounds.

And thus can condition their strategies on what happened previously.

Players

$$N = \{1, 2\}$$

Strategies of Player 1

(C, C), (C, D), (D, C), (D, D)

Strategies of Player 2

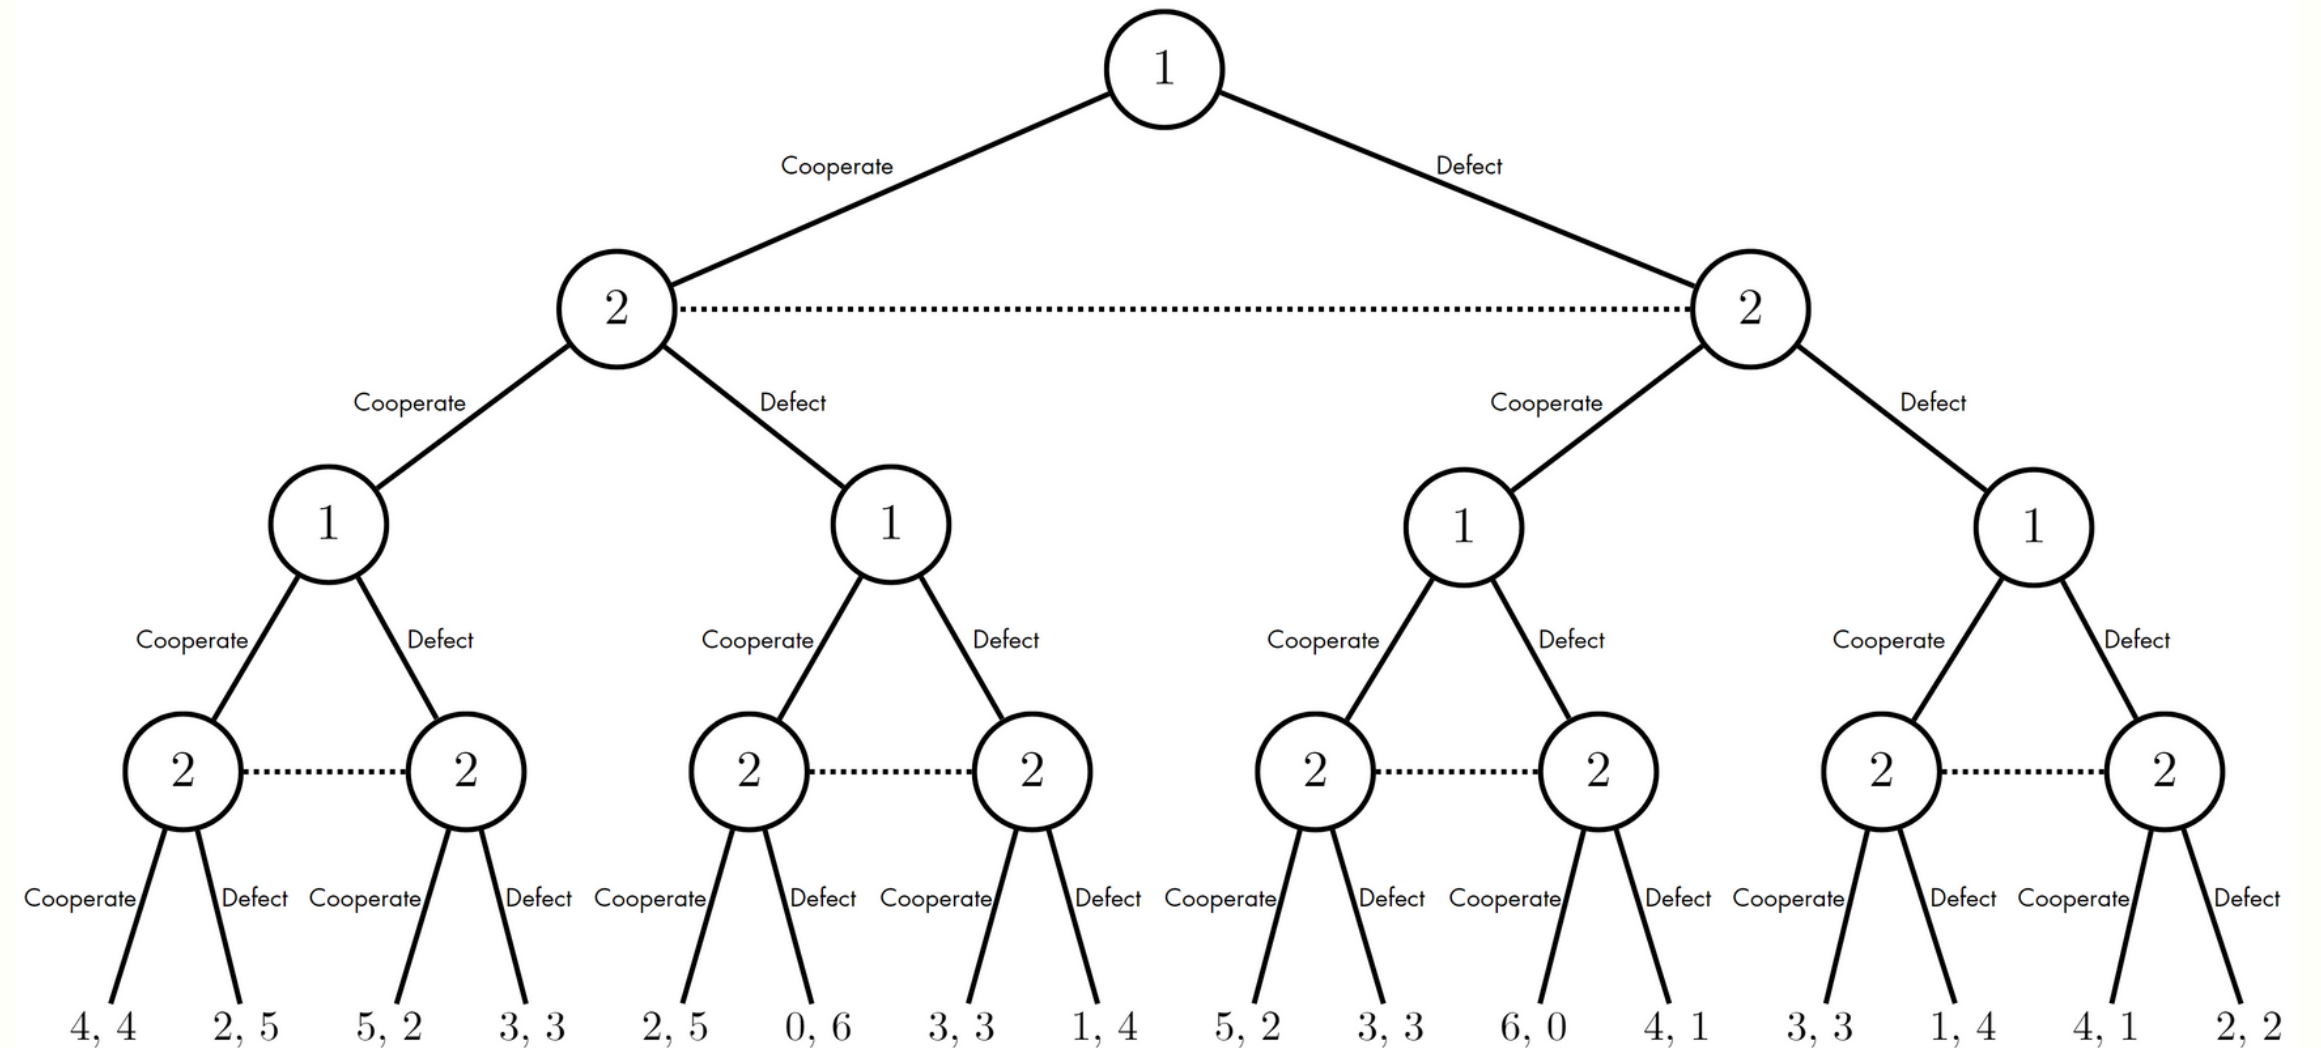
(C, C), (C, D), (D, C), (D, D)

Strategy profiles

((C, C), (C, C)), ((C, C), (C, D)), ...

Payoffs (aka utilities)

hopefully clear



Straightforward to get a
table now.

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|--------------|--------------|--------------|--------------|
| C, C | 2 + 2, 2 + 2 | 2 + 0, 2 + 3 | 0 + 2, 3 + 2 | 0 + 0, 3 + 3 |
| C, D | 2 + 3, 2 + 0 | 2 + 1, 2 + 1 | 0 + 3, 3 + 0 | 0 + 1, 3 + 1 |
| D, C | 3 + 2, 0 + 2 | 3 + 0, 0 + 3 | 1 + 2, 1 + 2 | 1 + 0, 1 + 3 |
| D, D | 3 + 3, 0 + 0 | 3 + 1, 0 + 1 | 1 + 3, 1 + 0 | 1 + 1, 1 + 1 |

strictly dominant strategies

?

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies

?

Pareto optimal strategy profiles

?

pure Nash equilibria

?

mixed Nash equilibria

?

In general, every game of imperfect information corresponds to a normal-form game, and vice-versa.

Thus, Nash equilibria and everything else are defined as for normal-form games.

So how do we analyze the 2-
round Prisoner's Dilemma?

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies
?

Pareto optimal strategy profiles

pure Nash equilibria

mixed Nash equilibria

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies

$((D, D), (D, D))$

Pareto optimal strategy profiles

?

pure Nash equilibria

mixed Nash equilibria

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies

$((D, D), (D, D))$

Pareto optimal strategy profiles

see above

pure Nash equilibria

?

mixed Nash equilibria

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies

$((D, D), (D, D))$

Pareto optimal strategy profiles

see previous

pure Nash equilibria

$((D, D), (D, D))$

mixed Nash equilibria

?

Iterated Prisoner's Dilemma



2 iterations

Two players play the Prisoner's Dilemma over $k = 2$ rounds.

The final payoffs are the sum of the payoffs from each round.

payoffs

| | C, C | C, D | D, C | D, D |
|------|------|------|------|------|
| C, C | 4, 4 | 2, 5 | 2, 5 | 0, 6 |
| C, D | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, C | 5, 2 | 3, 3 | 3, 3 | 1, 4 |
| D, D | 6, 0 | 4, 1 | 4, 1 | 2, 2 |

strictly dominant strategies

$((D, D), (D, D))$

Pareto optimal strategy profiles

see previous

pure Nash equilibria

$((D, D), (D, D))$

mixed Nash equilibria

none

Again, the only Nash equilibrium is to always defect, for both players.

Again, the only Nash equilibrium is to always defect, for both players.

Note that we'd get the same equilibrium by backward induction.

Again, the only Nash equilibrium is to always defect, for both players.

Note that we'd get the same equilibrium by backward induction.

Note, as well, that we'd get the same conclusion for $k > 2$ rounds.

Well that was pointless.

Well that was pointless.

Let's do a recap of where we are.

In the Prisoner's Dilemma, the unique Nash equilibrium (in strictly dominant strategies even) requires both players to defect.

In the Prisoner's Dilemma, the unique Nash equilibrium (in strictly dominant strategies even) requires both players to defect.

We often observe cooperation in the real world.

In the Prisoner's Dilemma, the unique Nash equilibrium (in strictly dominant strategies even) requires both players to defect.

We often observe cooperation in the real world.

What should we add to our model to make cooperation rational?

In the Prisoner's Dilemma, the unique Nash equilibrium (in strictly dominant strategies even) requires both players to defect.

We often observe cooperation in the real world.

What should we add to our model to make cooperation rational?

Maybe if players acknowledge they are in a repeated relationship.

In the Prisoner's Dilemma, the unique Nash equilibrium (in strictly dominant strategies even) requires both players to defect.

We often observe cooperation in the real world.

What should we add to our model to make cooperation rational?

Maybe if players acknowledge they are in a repeated relationship.

Unfortunately, if the Prisoner's Dilemma is repeated a commonly known finite number of times, backwards induction implies that players still defect at every round.


ROBERT AUMANN

What if the game is played for an infinite number of times?

As in, we don't have a fixed number k of rounds at which the
game ends.



Iterated Prisoner's Dilemma



infinitely iterated

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round.

Players

$$N = \{1, 2\}$$

Strategies of Player 1

$(C, C, \dots), (C, D, \dots), \dots$

Strategies of Player 2

$(C, C, \dots), (C, D, \dots), \dots$

Payoffs (aka utilities)

In general, infinite sums.

For instance, if both players always cooperate, payoffs are infinite series: $(2, 2, \dots)$, and the final payoff is:

$$2 + 2 + \dots = \infty$$

ROBERT AUMANN

Let's also add a *discount factor* δ , with $0 < \delta < 1$, which works as follows.



At every new round, the payoffs are multiplied by δ .

ROBERT AUMANN


Let's also add a *discount factor* δ , with $0 < \delta < 1$, which works as follows.



At every new round, the payoffs are multiplied by δ .

So for $\delta = 0.8$, \$100 today is worth $0.8 \cdot \$100 = \80 tomorrow,
and $0.8 \cdot \$80 = \64 in two days.

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor, $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Players

$$N = \{1, 2\}$$

Strategies of Player 1

$(C, C, \dots), (C, D, \dots), \dots$

Strategies of Player 2

$(C, C, \dots), (C, D, \dots), \dots$

Payoffs (aka utilities)

In general, infinite sums.

For instance, if both players always cooperate, payoffs are infinite series: $(2, 2\delta, 2\delta^2, \dots)$, and the final payoff is:


$$2 + 2\delta + 2\delta^2 + \dots$$

In general, for infinite sums we can use the following identity, for $0 < x < 1$:

$$1 + x + x^2 + \dots = \frac{1}{1-x}$$



Iterated Prisoner's Dilemma



infinitely iterated, with discount factor, $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Players

$$N = \{1, 2\}$$

Strategies of Player 1

$(C, C, \dots), (C, D, \dots), \dots$

Strategies of Player 2

$(C, C, \dots), (C, D, \dots), \dots$

Payoffs (aka utilities)

In general, infinite sums.

For instance, if both players always cooperate, payoffs are infinite series: $(2, 2\delta, 2\delta^2, \dots)$, and the final payoff is:

$$\begin{aligned} 2 + 2\delta + 2\delta^2 + \dots &= 2(1 + \delta + \delta^2 + \dots) \\ &= 2 \cdot \frac{1}{1 - \delta} \end{aligned}$$

What does the discount factor
 δ stand for?

Interpreting the discount factor

Patience

You're more patient the less you mind waiting for something valuable, rather than receiving it immediately.

For a discount factor δ you value \$1, received t rounds from now, at $\$1 \cdot \delta^t$.

This is less than \$1, because $0 < \delta < 1$.

As δ gets closer to 1, the agent is more patient.

Interpreting the discount factor

Patience

You're more patient the less you mind waiting for something valuable, rather than receiving it immediately.

For a discount factor δ you value \$1, received t rounds from now, at $\$1 \cdot \delta^t$.

This is less than \$1, because $0 < \delta < 1$.

As δ gets closer to 1, the agent is more patient.

Uncertainty about the future

You might prefer \$1 today to \$1 tomorrow because you're not sure tomorrow will even come.

δ can be the probability that there is a round $t + 1$, if round t has happened.

$\$1 \cdot \delta^t$ is then the expected payoff at round t .

ROBERT AUMANN

Consider, now, the following strategy, called *Grim Trigger*.



Start by cooperating. If the other player defects at some round t , switch to defecting forever, i.e., at every round $t' > t$.

Let's look at a run of the game
when one player plays Grim
Trigger.

Example Runs with Grim Trigger

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Start by cooperating; defect once at some random round $t > 1$

Sample run

| | actions taken |
|----------|-----------------------|
| Player 1 | C, C, C, D, D, D, ... |
| Player 2 | C, C, D, C, C, C, ... |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Example Runs with Grim Trigger

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Start by cooperating; defect once at some random round $t > 1$


Sample run

| | actions taken |
|----------|-----------------------|
| Player 1 | C, C, C, D, D, D, ... |
| Player 2 | C, C, D, C, C, C, ... |

betrayal

revenge forever after

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Example Runs with Grim Trigger

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Start by cooperating; defect once at some random round $t > 1$


Sample run

| | actions taken | payoffs | total payoff |
|----------|---|---|------------------|
| Player 1 | C, C, C, D , D , D , ... | $2, 2\delta, 0\delta^2, 3\delta^3, 3\delta^4, 3\delta^5, \dots$ | the infinite sum |
| Player 2 | C, C, D , C, C, C, ... | $2, 2\delta, 3\delta^2, 0\delta^3, 0\delta^4, 0\delta^5, \dots$ | the infinite sum |

revenge forever after

betrayal

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

And when both players use
Grim Trigger?

Example Runs with Grim Trigger

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Grim Trigger

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--------------------------------|------------------------|
| Player 1 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/1-\delta)$ |
| Player 2 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/1-\delta)$ |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Does any agent have an
incentive to deviate from Grim
Trigger?

Player 2 deviates by always defecting

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Deviate, starting at first round

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--|---------------------|
| Player 1 | C, D, D, D, D, D, ... | $0, \delta, \delta^2, \delta^3, \dots$ | $\delta/(1-\delta)$ |
| Player 2 | D, D, D, D, D, D, ... | $3, \delta, \delta^2, \delta^3, \dots$ | $2 + 1/(1-\delta)$ |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by always defecting

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Deviate, starting at first round

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--|-------------------------|
| Player 1 | C, D, D, D, D, D, ... | $0, \delta, \delta^2, \delta^3, \dots$ | $\delta/(1-\delta)$ |
| Player 2 | D, D, D, D, D, D, ... | $3, \delta, \delta^2, \delta^3, \dots$ | $2 + \delta/(1-\delta)$ |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by always defecting

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Deviate, starting at first round

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--|---------------------|
| Player 1 | C, D, D, D, D, D, ... | $0, \delta, \delta^2, \delta^3, \dots$ | $\delta/(1-\delta)$ |
| Player 2 | D, D, D, D, D, D, ... | $3, \delta, \delta^2, \delta^3, \dots$ | $2 + 1/(1-\delta)$ |

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Grim Trigger

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--------------------------------|--------------------------|
| Player 1 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |
| Player 2 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by always defecting

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Deviate, starting at first round

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--|---------------------|
| Player 1 | C, D, D, D, D, D, ... | $0, \delta, \delta^2, \delta^3, \dots$ | $\delta/(1-\delta)$ |
| Player 2 | D, D, D, D, D, D, ... | $3, \delta, \delta^2, \delta^3, \dots$ | $2 + 1/(1-\delta)$ |

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Grim Trigger

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--------------------------------|--------------------------|
| Player 1 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |
| Player 2 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |

Profitable?


Not a profitable deviation for Player 2 as long as:

$$2 + \frac{1}{1-\delta} \leq 2 \cdot \frac{1}{1-\delta},$$

which happens if and only if:

$$\delta \geq \frac{1}{2}$$

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

What if Player 2 defects later?

Player 2 deviates by defecting later

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Deviate, starting at round $k > 1$

Sample run

| | actions taken | payoffs | total payoff |
|----------|-------------------------|--|------------------|
| Player 1 | C, ..., C, C, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 0\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |
| Player 2 | C, ..., C, D, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 3\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by defecting later

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Deviate, starting at round $k > 1$

Sample run

| | actions taken | payoffs | total payoff |
|----------|-------------------------|--|------------------|
| Player 1 | C, ..., C, C, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 0\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |
| Player 2 | C, ..., C, D, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 3\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by defecting later

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Deviate, starting at round $k > 1$

Sample run

| | actions taken | payoffs | total payoff |
|----------|-------------------------|--|------------------|
| Player 1 | C, ..., C, C, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 0\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |
| Player 2 | C, ..., C, D, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 3\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |

Strategy of Player 1

Grim Trigger


Strategy of Player 2

Grim Trigger

Sample run

| | actions taken | payoffs | total payoff |
|----------|-----------------------|--------------------------------|------------------------|
| Player 1 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/1-\delta)$ |
| Player 2 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/1-\delta)$ |

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Player 2 deviates by defecting later

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Deviate, starting at round $k > 1$

Sample run

| | actions taken | payoffs | total payoff |
|----------|-------------------------|--|------------------|
| Player 1 | C, ..., C, C, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 0\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |
| Player 2 | C, ..., C, D, D, D, ... | $2, 2\delta, \dots, 2\delta^{k-1}, 3\delta^k, \delta^{k+1}, \delta^{k+2}, \dots$ | the infinite sum |

Strategy of Player 1

Grim Trigger

Strategy of Player 2

Grim Trigger

Sample run


| | actions taken | payoffs | total payoff |
|----------|-----------------------|--------------------------------|--------------------------|
| Player 1 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |
| Player 2 | C, C, C, C, C, C, ... | $2, 2\delta, 2\delta^2, \dots$ | $2 \cdot (1/(1-\delta))$ |

Profitable?

Not a profitable deviation for Player 2 as long as:

$$\begin{aligned}
 2 + 2\delta + \dots + 2\delta^{k-1} + 3\delta^k + \delta^{k+1} + \dots &\leq 2 + 2\delta + \dots + 2\delta^{k-1} + 2\delta^k + 2\delta^{k+1} + \dots && \text{iff} \\
 3\delta^k + \delta^{k+1} + \dots &\leq 2\delta^k + 2\delta^{k+1} + \dots && \text{iff} \\
 3 + \delta + \delta^2 + \dots &\leq 2 + 2\delta + 2\delta^2 + \dots && \text{iff} \\
 &\delta \geq 1/2.
 \end{aligned}$$

Iterated Prisoner's Dilemma



infinitely iterated, with discount factor $0 < \delta < 1$

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

but an infinite number of times.

The final payoffs are the sum of the payoffs from each round, taking into account the discount factor δ .

Note that once Player 2 triggers Player 1 by defecting, Player 2 has no incentive to start cooperating again if all-defection is not profitable.

ROBERT AUMANN

We've just shown that if $\delta \geq 0.5$, no agent has an incentive to deviate.



In other words, both players playing Grim Trigger is a Nash equilibrium!

Finally, a positive result!

Infinite games (with sufficiently large discount factor) admit equilibria where players cooperate!

The moral?

If players send out a clear signal that they cannot be pushed around, it makes sense to cooperate.

ROBERT AUMANN



There's many other ways of analyzing repeated games.

With or without discounting, with different ways of computing total payoffs, with different types of equilibria (Nash, subgame-perfect).

When these equilibria can be achieved is the subject of intense research.

Results here usually go under the name of *folk theorems*.

At the same time, Grim Trigger strategies are just one drop in the vast sea of possible strategies.

They are especially unforgiving, and do not match what we see in real life.

What else can we do?

ROBERT AXELROD
How about this.



Axelrod, R. (1984), *The Evolution of Cooperation*. Basic Books

ROBERT AXELROD
How about this.



Take a bunch of strategies, whatever sounds plausible, and
pit them against each other.

Tournament style!

Axelrod, R. (1984), The Evolution of Cooperation. Basic Books

ROBERT AXELROD
How about this.



Take a bunch of strategies, whatever sounds plausible, and
pit them against each other.

Tournament style!

Strategies with highest average payoffs are declared the
winners.

Axelrod, R. (1984), The Evolution of Cooperation. Basic Books

Axelrod organized such a tournament in the late '70s.

Axelrod organized such a tournament in the late '70s.

He invited researchers from across the world to submit strategies for the repeated Prisoner's Dilemma.

Axelrod organized such a tournament in the late '70s.

He invited researchers from across the world to submit strategies for the repeated Prisoner's Dilemma.

Strategies could take into account previous moves, and could be as complex as their authors wanted.

Axelrod organized such a tournament in the late '70s.

He invited researchers from across the world to submit strategies for the repeated Prisoner's Dilemma.

Strategies could take into account previous moves, and could be as complex as their authors wanted.

They were then pitted against each other in a round-robin tournament, played on the computer.

Axelrod organized such a tournament in the late '70s.

He invited researchers from across the world to submit strategies for the repeated Prisoner's Dilemma (200 rounds).

Strategies could take into account previous moves, and could be as complex as their authors wanted.

They were then pitted against each other in a round-robin tournament, played on the computer.

Fourteen strategies were submitted, and Axelrod added one extra.

How would you play?

How would you play?

Here's some of the strategies
that were submitted.

ROBERT AXELROD

Random: cooperate and defect randomly.



ROBERT AXELROD

Random: cooperate and defect randomly.



ANATOL RAPOPORT

Tit-for-tat: start by cooperating, then copy opponent's previous move.

ROBERT AXELROD

Random: cooperate and defect randomly.



ANATOL RAPOPORT


Tit-for-tat: start by cooperating, then copy opponent's previous move.



JOHANN JOSS

Defect after other player defects (like tit-for-tat).

When the other player cooperates, cooperate 90% of the time. And, hence, defect 10% of the time.



ROBERT AXELROD

Random: cooperate and defect randomly.



ANATOL RAPOPORT

Tit-for-tat: start by cooperating, then copy opponent's previous move.



JOHANN JOSS

Defect after other player defects (like tit-for-tat).

When the other player cooperates, cooperate 90% of the time. And, hence, defect 10% of the time.



JAMES GRAASKAMP

Play tit-for-tat 50 moves, defect once, play tit for tat for another 5 moves, and then examine the history of the game so far.



Try to guess the opponent, and adjust moves accordingly.

Random vs tit-for-tat

Random, $p = 0.7$ (Axelrod)

Cooperate with probability p , defect with probability $1 - p$.

Tit-for-tat (Rapoport)

Start by cooperating; thereafter copy opponent's last move.

Sample run

| | strategy | actions taken (8 rounds) | payoffs | total payoff |
|----------|-------------|--------------------------|------------------------|--------------|
| Player 1 | random | C, C, D, C, C, D, C, C | 2, 2, 3, 0, 2, 3, 0, 2 | 14 |
| Player 2 | tit-for-tat | C, C, C, D, C, C, D, C | 2, 2, 0, 3, 2, 0, 3, 2 | 14 |

Iterated Prisoner's Dilemma

Iterated for 8 rounds

Two players play the regular Prisoner's Dilemma:

| | C | D |
|---|------|------|
| C | 2, 2 | 0, 3 |
| D | 3, 0 | 1, 1 |

Game is played 8 across rounds.

The final payoffs are the sum of the payoffs from each round.

Who won?

ROBERT AXELROD

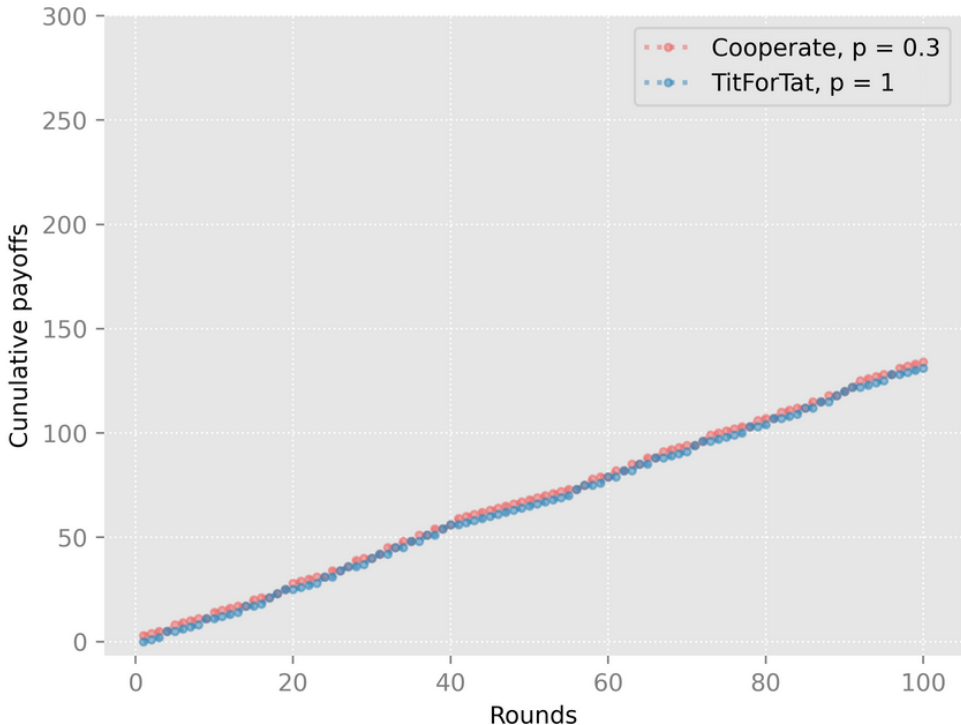


One would expect the most complex, sophisticated program
would win.



But in fact, tit-for-tat won.

This was pretty much the simplest strategy submitted: the
code for it was four lines.

Tit-for-tat versus an occasional cooperators



ROBERT AXELROD

A couple of years later, I organized another tournament, this time with 63 entries submitted.



ROBERT AXELROD

A couple of years later, I organized another tournament, this time with 63 entries submitted.



JOHN MAYNARD SMITH

One was mine, called tit-for-two-tats: cooperate unless opponent defects twice in a row.

ROBERT AXELROD

A couple of years later, I organized another tournament, this time with 63 entries submitted.



JOHN MAYNARD SMITH

One was mine, called tit-for-two-tats: cooperate unless opponent defects twice in a row.



ROBERT AXELROD

Tit-for-tat won again.



Morals?

ROBERT AXELROD

One property shared by the highest-scoring strategies.



Being *nice*.

That is, never being the first to defect.

Tit-for-tat does not bear a grudge beyond the immediate retaliation, and provides the opportunity to establish 'trust' between opponents.

Does Axelrod's tournament say anything about the real world?

ROBERT L. TRIVERS



Yes! We can see reciprocity throughout the natural world.

Trivers, R.L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*. 46: 35–57.



MANFRED MILINSKI

Stickleback fish rely on tit-for-tat to inspect potential predators.

Milinski, M. (1987). TIT FOR TAT in sticklebacks and the evolution of cooperation. *Nature*, 325(6103), 433–435.