Modelling a Simultaneous Game As an Extended Game of Imperfect Information

Consider the following game (a "Prisoner's Dilemma").

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

We think of "high" as the "nice, cooperate" strategy; and "low" as the "competitive" strategy

Player 2 doesn't know whether they are at α or β, since they don't observe Player 1's move. They only know that they are within "I"

Imagine Player 1 moves very slightly prior to Player 2, but Player 2 doesn't see their move.

We express this using an "information set" (the dotted line).

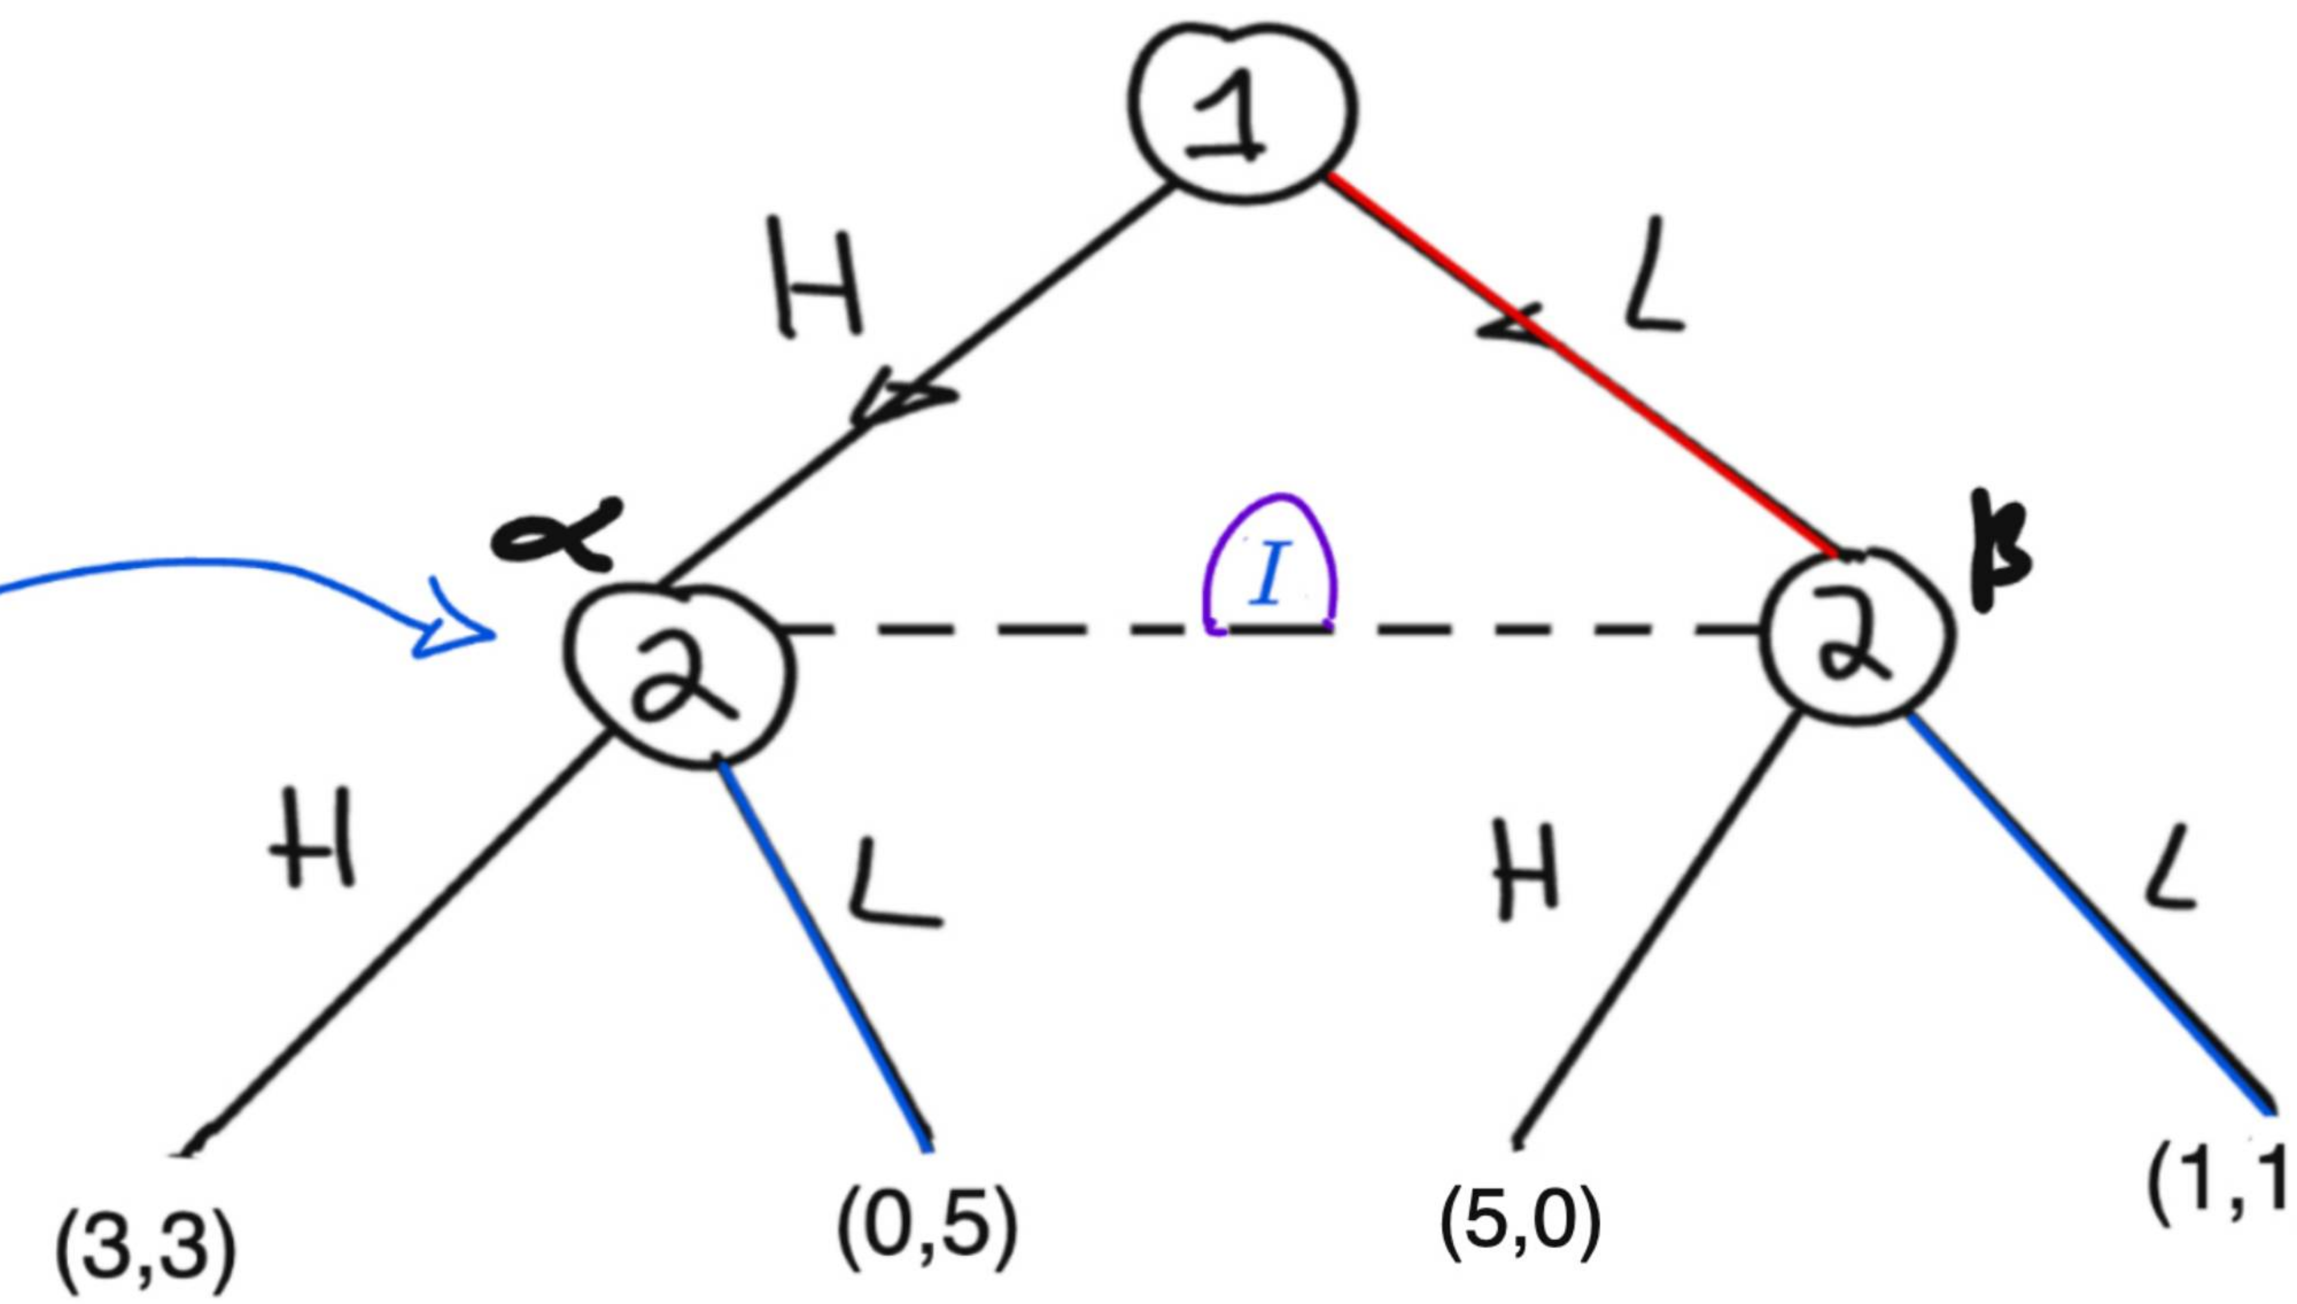

Usually, we'd be able to assign SEPARATE moves to Player 2, at α and β. But since they wouldn't know which of these to do, this wouldn't be feasible for Player 2. Instead we must assign the SAME move at both α and β; i.e. at the "information set" I.

This is a game of "imperfect information" -- the players know the whole tree, but not always where they are.

In such a game, a "pure strategy" is an action for each INFORMATION SET (not each node).

Subgames in a game of imperfect information

Consider the game as played out from α:

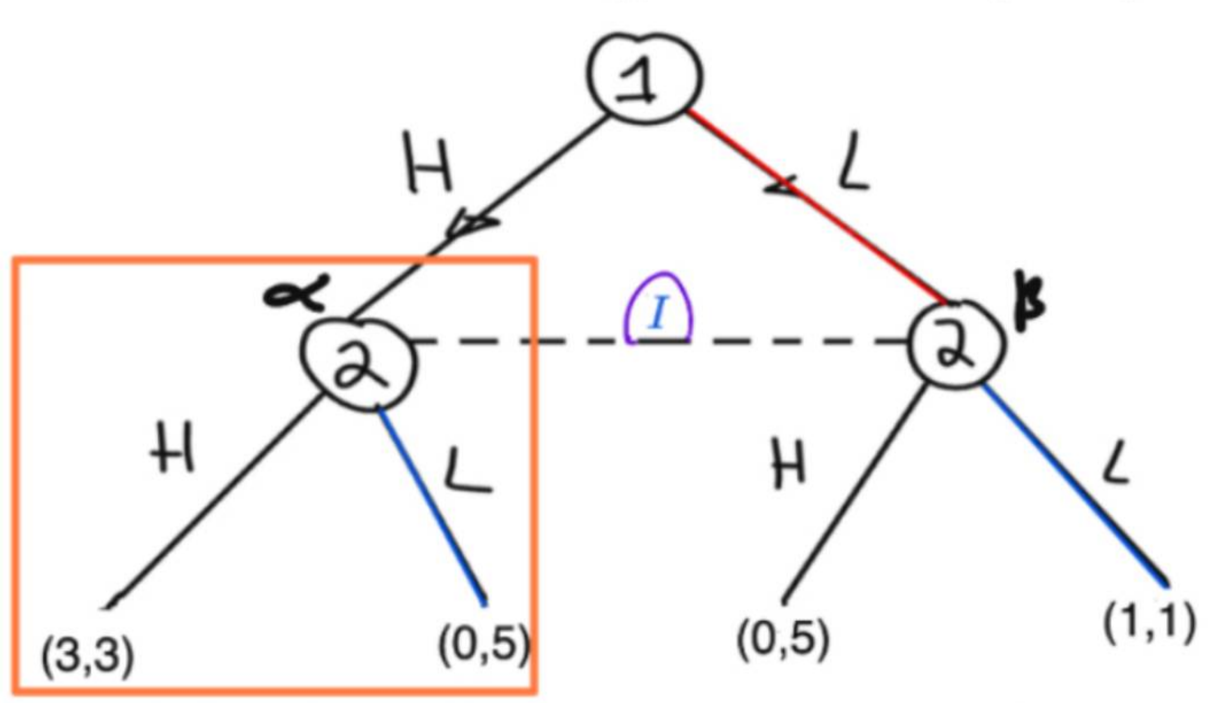

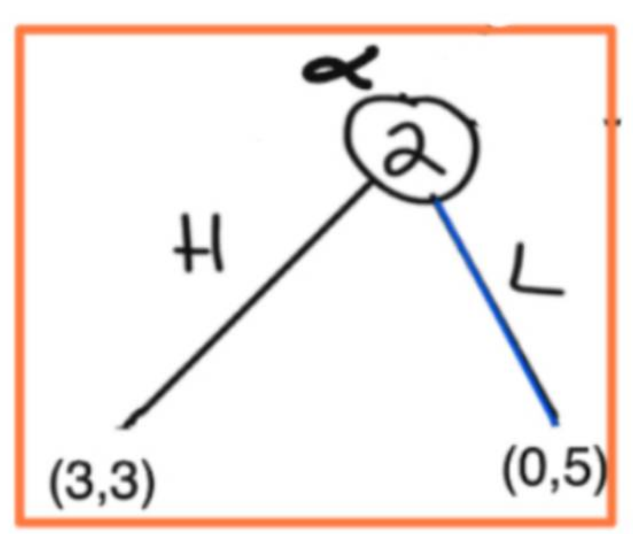

This is NOT a subgame (in sugames, we do NOT cut across information sets).
This is because, in this "minigame", Player 2 knows they are at α, with possible payoffs 3 and 5. But they're never in position in the main game.

Playing a Game Twice

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

Two rounds of the "stage" game

The dotted line means an "information set":
P2 has to pick the SAME move here
(this is how we model moving at the same time)

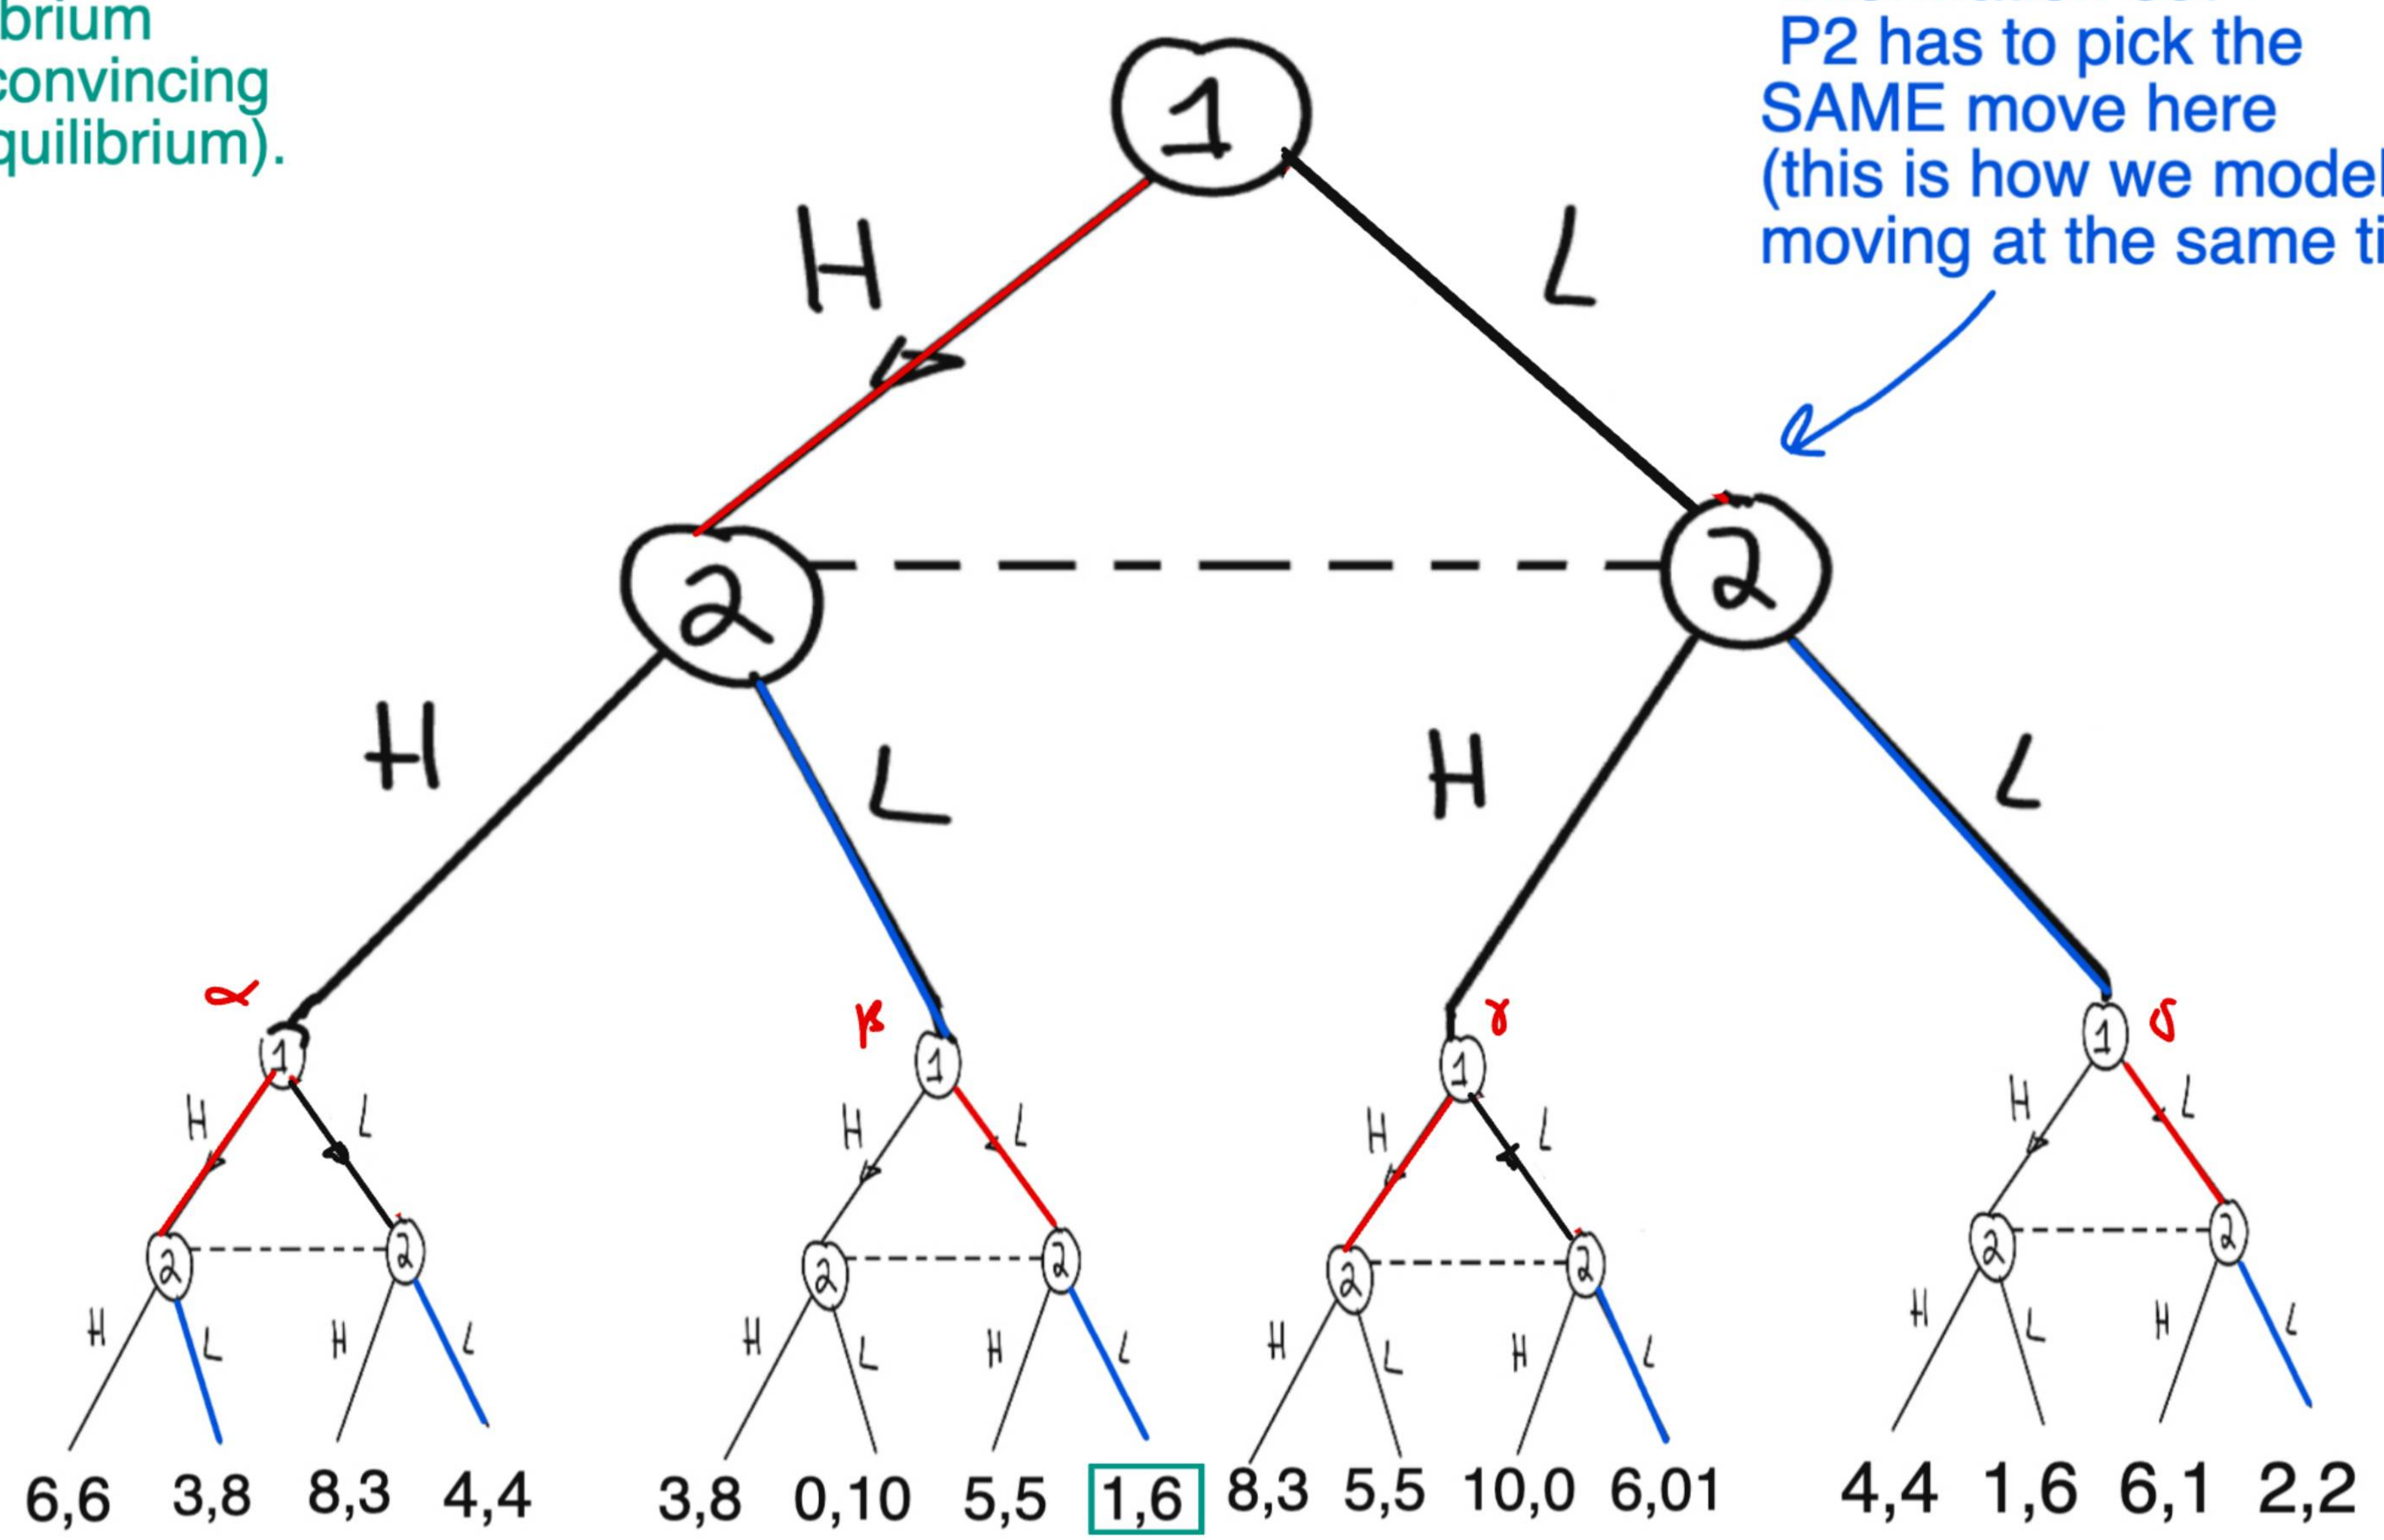

We play the "stage game" twice in a row.

In this case, we just add the payoffs; in other examples, we'll do something more complicated (with infinitely repeated games).

How many pure strategies does Player 1 have?

$$2 \times 2 \times 2 \times 2 \times 2 = 2^5 = 32$$

Likewise, since Player 2 has 5 information sets, they only have 5 decisions to make as well; and they also have 32 pure strategies.

In the strategy shown, Player 1 is choosing DIFFERENT moves in different "versions" of the second round.
We say their strategy is condition on HISTORY; i.e. what happens in the first round.

At node α, following HH in round 1, they play H
At node β, following HL in round 1, they play L
At node γ, following LH in round 1, they play H
At node δ, following LL in round 1, they play L

More concisely, they "copy" Player 2's move in the first round.

In complex repeated games, there are too many strategies to analyse all of them. We focus on nice "simply strategies" like this one -- called "tit for tat".

Player 2's strategy, however, is quite "boring" -- they always plays L.

Axelrod's Tournament

The repeated prisoner's dilemma is complicated enough that there was a whole tournament about it, where player's submitted strategies!

In 1980, Robert Axelrod, professor of political science at the University of Michigan, held a tournament of various strategies for the prisoner's dilemma. He invited a number of well-known game theorists to submit strategies to be run by computers. In the tournament, programs played games against each other and themselves repeatedly. Each strategy specified whether to cooperate or defect based on the previous moves of both the strategy and its opponent.

Some of the strategies submitted were:

- Always defect: This strategy defects on every turn. This is what game theory advocates. It is the safest strategy since it cannot be taken advantage of. However, it misses the chance to gain larger payoffs by cooperating with an opponent who is ready to cooperate.
- Always cooperate: This strategy does very well when matched against itself. However, if the opponent chooses to defect, then this strategy will do badly.
- Random: The strategy cooperates 50% of the time.

All of these strategies are prescribed in advance. Therefore, they cannot take advantage of knowing the opponent's previous moves and figuring out its strategy.

The winner of Axelrod's tournament was the TIT FOR TAT strategy. The strategy cooperates on the first move, and then does whatever its opponent has done on the previous move. Thus, when matched against the all-defect strategy, TIT FOR TAT strategy always defects after the first move. When matched against the all-cooperate strategy, TIT FOR TAT always cooperates. This strategy has the benefit of both cooperating with a friendly opponent, getting the full benefits of cooperation, and of defecting when matched against an opponent who defects. When matched against itself, the TIT FOR TAT strategy always cooperates.

Several variations to TIT FOR TAT have been proposed. TIT FOR TWO TATS is a forgiving strategy that defects only when the opponent has defected twice in a row. TWO TITS FOR TAT, on the other hand, is a strategy that punishes every defection with two of its own.

TIT FOR TAT relies on the assumption that its opponent is trying to maximize his score. When paired with a mindless strategy like RANDOM, TIT FOR TAT sinks to its opponent's level. For that reason, TIT FOR TAT cannot be called a "best" strategy.

It must be realized that there really is no "best" strategy for prisoner's dilemma. Each individual strategy will work best when matched against a "worse" strategy. In order to win, a player must figure out his opponent's strategy and then pick a strategy that is best suited for the situation.

"tit for tat"

Prog.	TFT	T&C	NY	GR	SH	S&R	FR	DA	GR	DO	FE	JO	TU	NA	RAN	Mean	Rank Point	No. of Wins	Rank Wins
TFT	600	595	600	600	600	595	600	600	597	597	280	225	279	359	441	**504**	1	0	15
T&C	600	596	600	601	600	596	600	600	310	601	271	213	291	455	573	**500**	2	11	2
NY	600	595	600	600	600	595	600	600	433	158	354	374	347	368	464	**486**	3	1	13.5
GR	600	595	600	600	600	594	600	600	376	309	289	236	305	426	507	**482**	4	4	6
SH	600	595	600	600	600	595	600	600	348	271	274	272	265	448	543	**481**	5	3	11.5
S&R	600	596	600	602	600	596	600	600	319	200	252	249	280	480	592	**478**	6	10	3.5
FR	600	595	600	600	600	595	600	600	307	207	235	213	263	489	598	**473**	7	6	8
DA	600	595	600	600	600	595	600	600	307	194	238	247	253	450	598	**472**	8	4	9.5
GR	597	305	462	375	348	314	302	302	588	625	268	238	274	466	548	**401**	9	5	9.5
DO	597	591	398	289	261	215	202	239	555	202	436	540	243	487	604	**391**	10	6	6
FE	285	271	426	286	297	255	235	239	274	704	246	236	272	420	467	**328**	11	12	3.5
JO	230	214	409	237	286	254	213	252	244	634	236	224	273	390	469	**304**	12	10	1
TU	284	287	415	293	318	271	243	229	278	193	271	260	273	426	478	**301**	13	6	6
NA	362	231	397	273	230	149	133	173	187	133	317	366	345	413	526	**282**	14	2	11.5
RAN	442	142	407	313	219	141	108	137	189	102	360	416	419	300	450	**276**	15	1	13.5

Program	Name	% Preliminary	% Wins
1	TFT	42.1	11.0
2	T&C	41.4	30.0
3	NY	41.5	10.2
4	GR	31.0	10.0
5	SH	24.4	6.7
6	S&R	39.9	24.8
7	FR	21.1	4.9
8	DA	13.8	2.1
9	GR	11.8	0.2
10	DO	20.8	0.1
11	FE	7.7	0
12	JO	2.5	0
13	TU	0.4	0
14	NA	1.2	0
15	RAND	0.5	0
		300.0%	100.0%

Subgames in the Previous Game

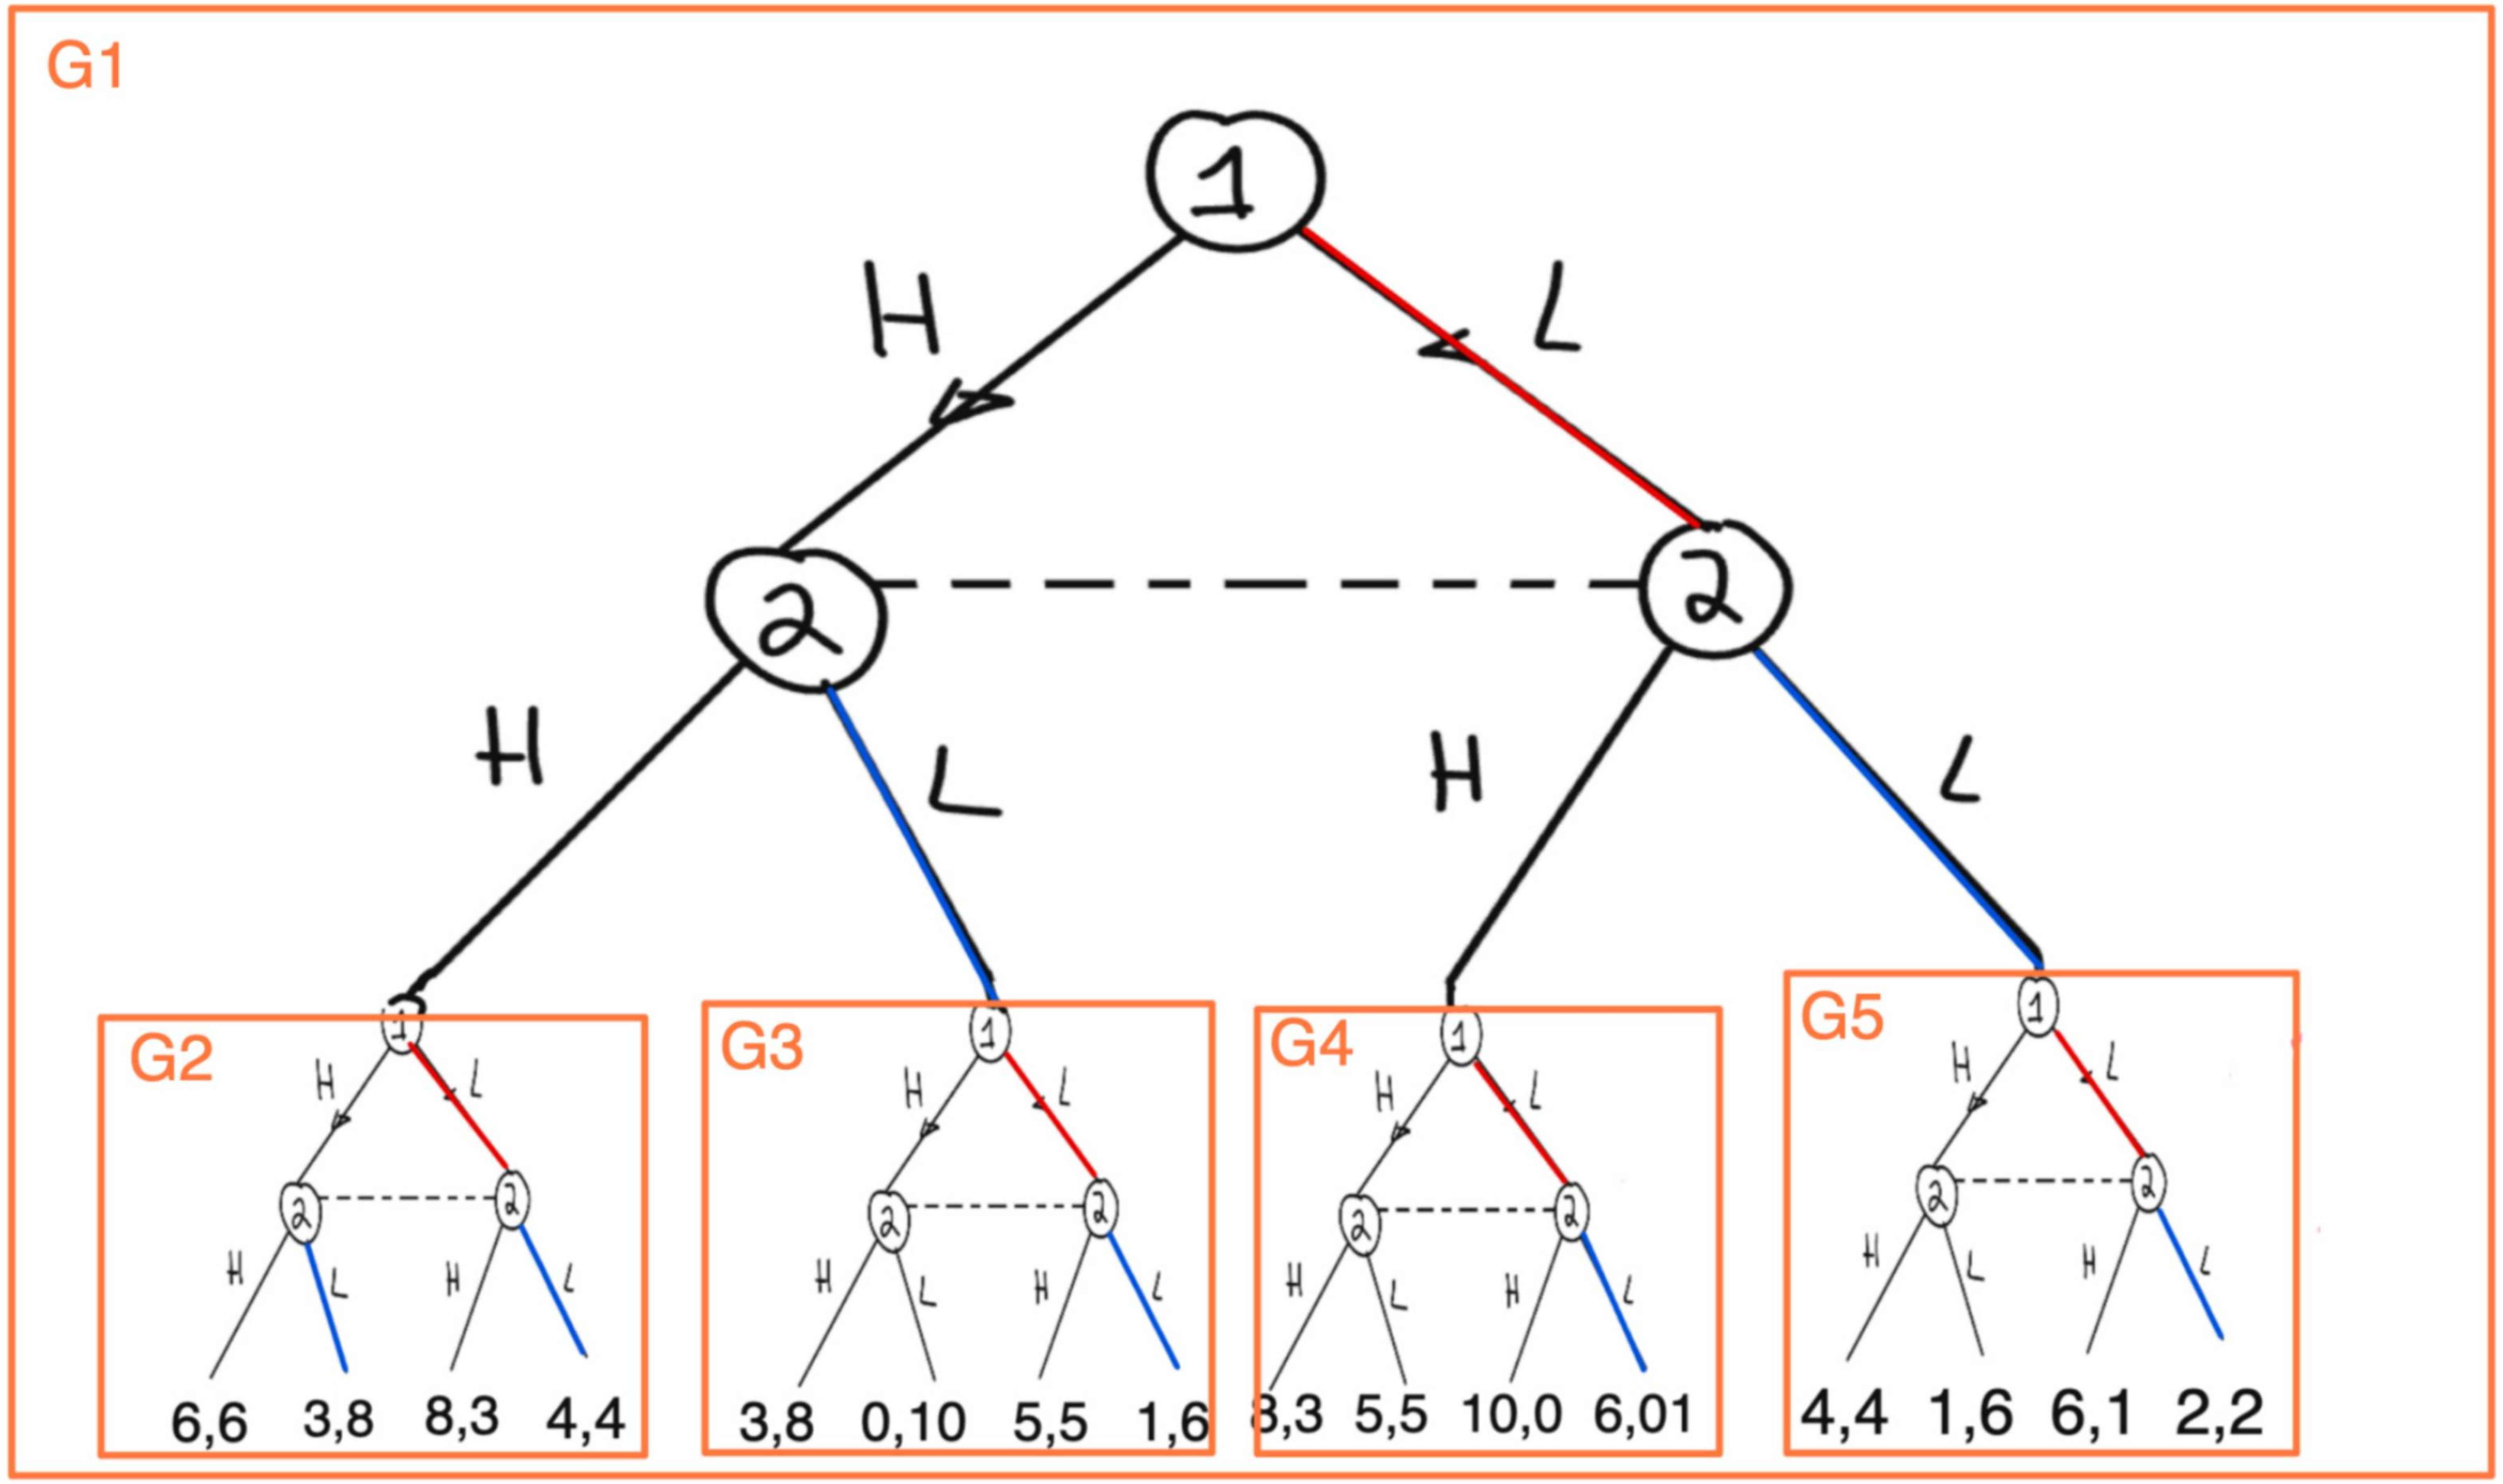

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

"Second rounds" are subgames, each a different one depending on history.

But, we must play (Low,Low) in all of them in ANY SPNE.

The previous game has FIVE subgames; each version of the second round, and the main game itself.

We end up playing the SAME NE in each version of the second round; so, although the first round is NOT a subgame and NOT analysed separately in general, in THIS CASE ONLY, it can be, since the second-round outcome is "pre-ordained".

Finitely Repeated Games -- Unique NE

This is basically a bit like playing the same game over and over. But, it's actually a lot more complicated!

Let's consider the following "stage" game. Suppose we play it twice in a row.

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

For the moment, let's just ADD the payoffs.

Even for this simple game, there are 32²=1024 strategy profiles.

BUT in fact, this game is quite boring to analysise.

Sequential forms

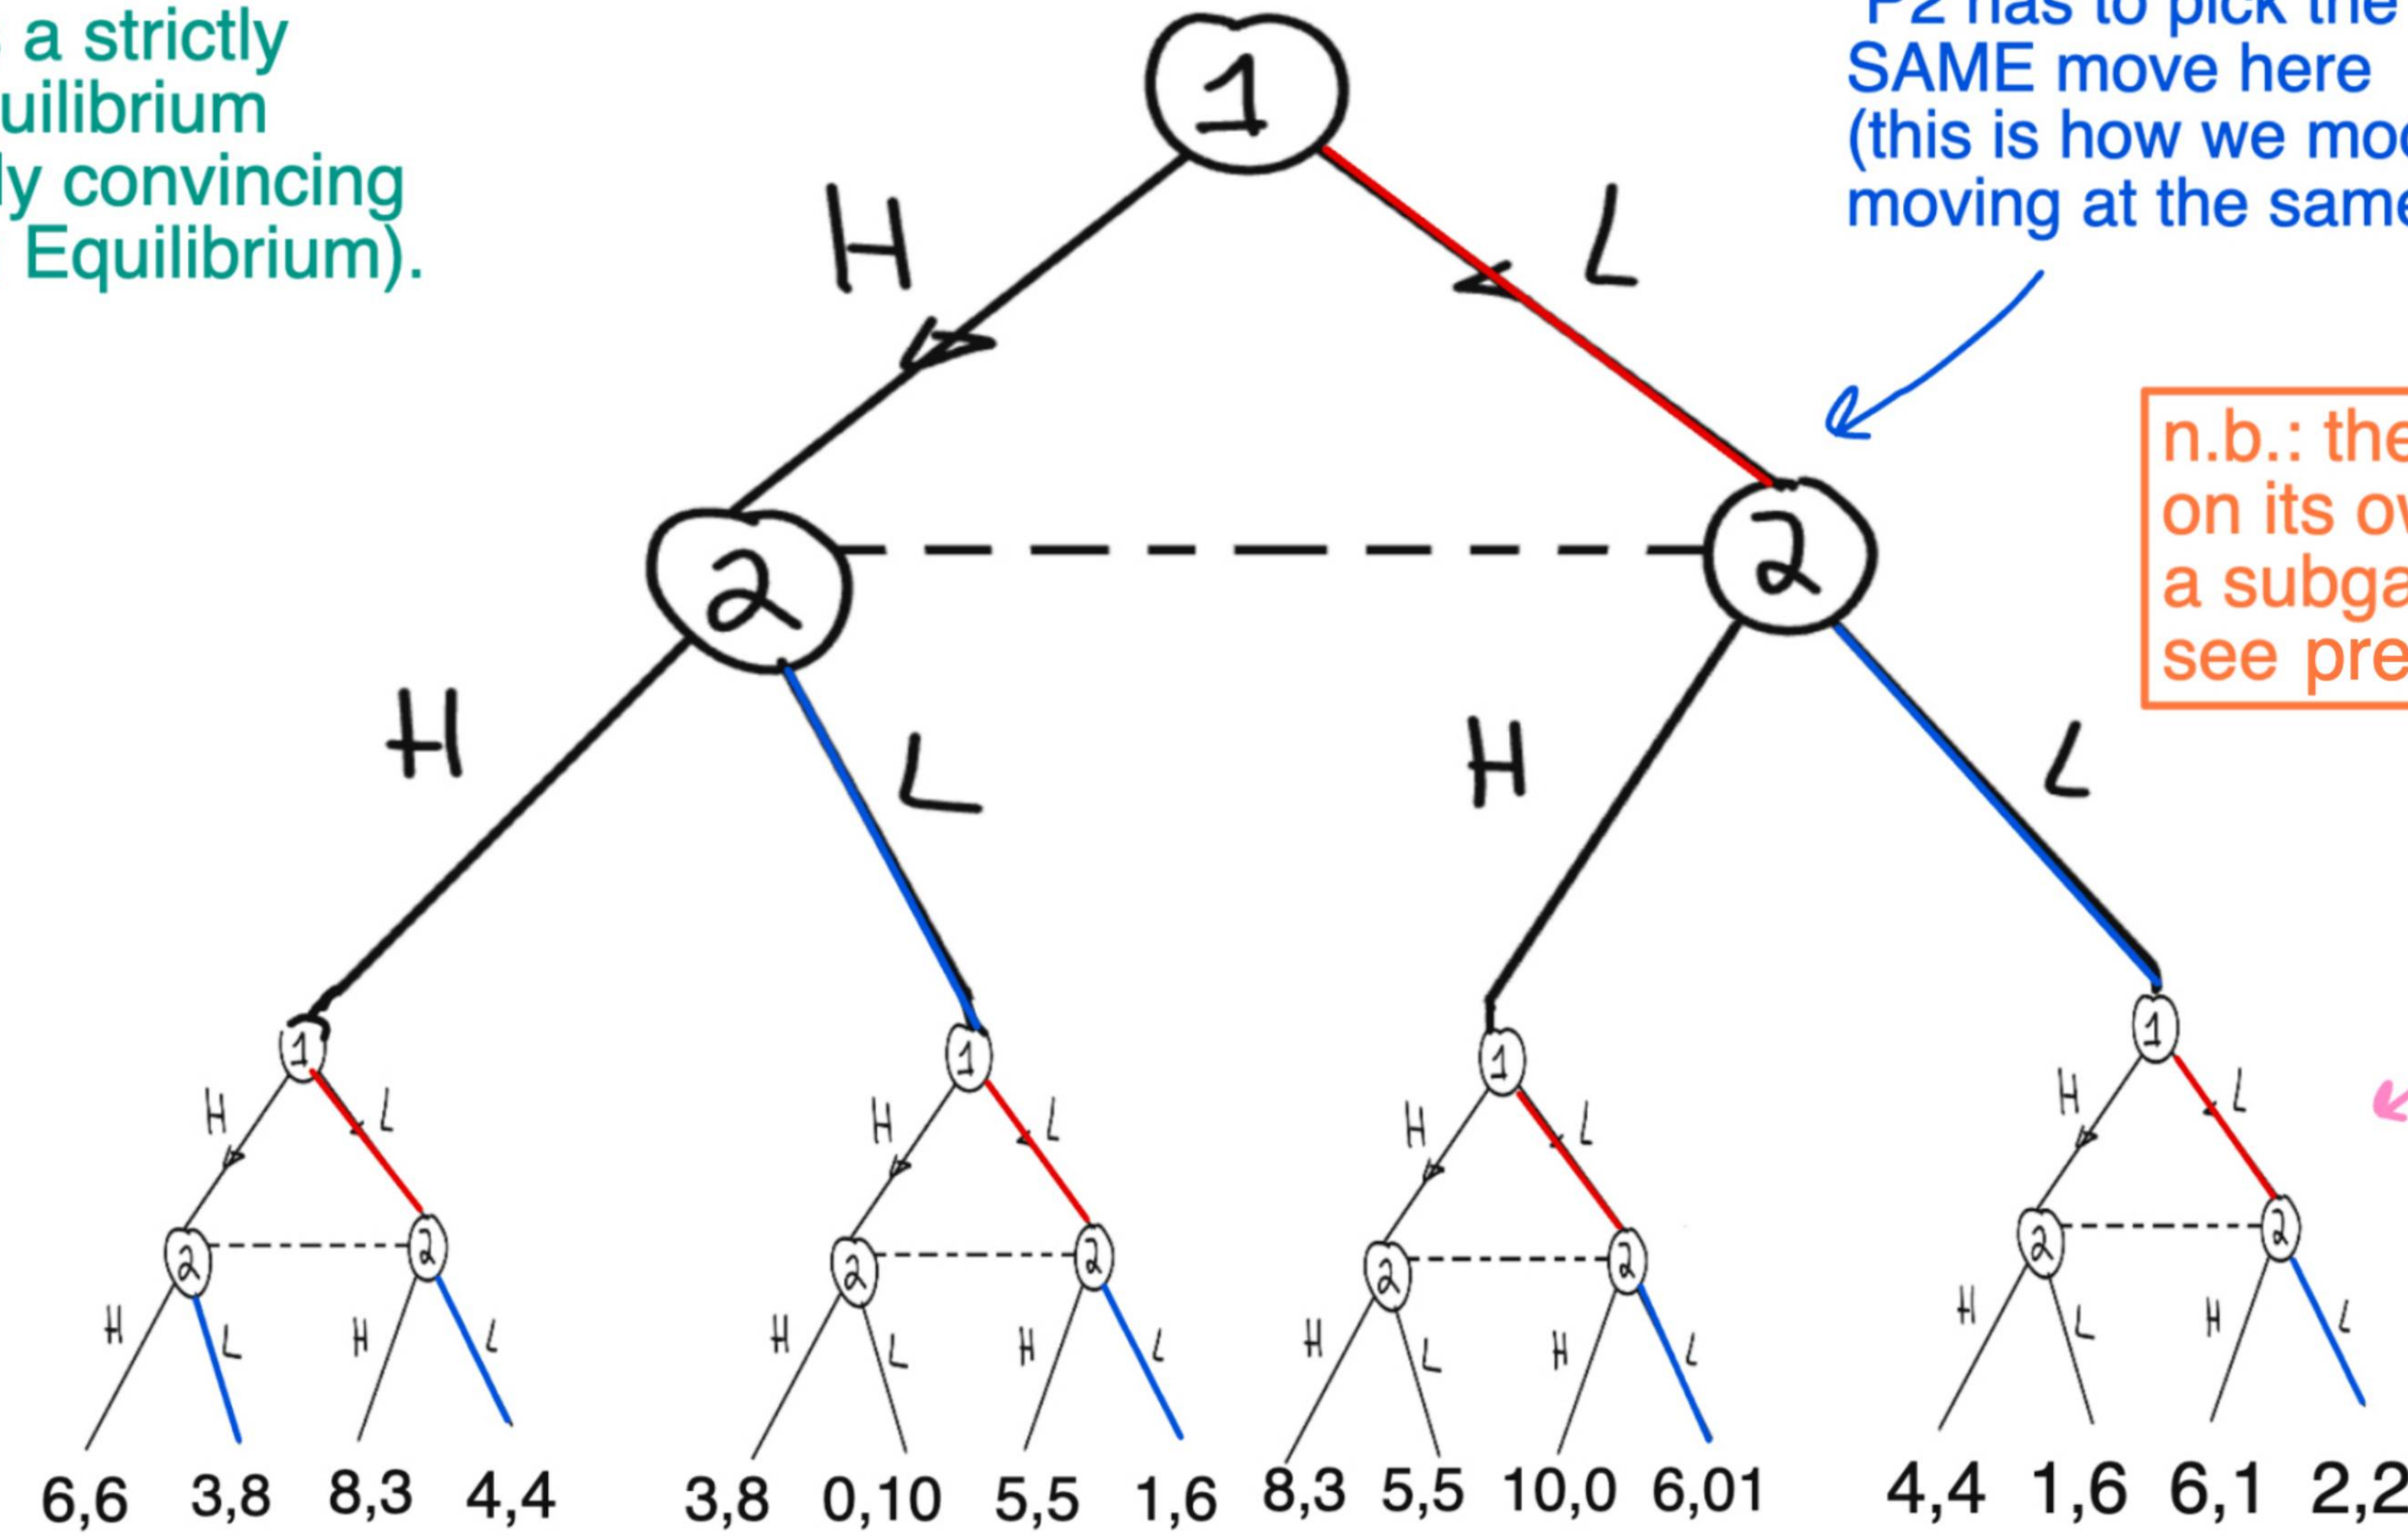

The dotted line means an "information set": P2 has to pick the SAME move here (this is how we model moving at the same time)

n.b.: the first round on its own is NOT a subgame. see previous slide

Every version of the second round plays out the same! just (Low,Low)

Second round play is "inevtiable"

Note that the second-stage subgames are all just the same as the stage game.

So, in ANY SPNE, players play "L" in ALL the second-stage subgames (regardless of history).
Hence, in the FIRST round, there is "no need to worry about the second round". So, players choose L in the first round as well. We have the UNIQUE SPNE ((L,(L,L,L,L)),(L,(L,L,L,L)).

Theorem. For any finitely repeated games (i.e. the stage game is repeated N times), if the stage game has a unique Nash Equilibrium, the only SPNE is just to play the NE strategies at every point in the game.

reasoning backwards from the end, with any number of rounds ...

Finitely Repeated Games with Multiple NE

Consider the following game, with TWO Nash Equilibria (one good; one bad).

	Left	Middle	Right
Top	3 3 🏆	0 4	−2 0
Middle	4 0	1 1	−2 0
Bottom	0 −2	0 −2	−1 −1

If there is more than one NE, then players may play a DIFFERENT NE in the second round, depending on what happens in the first round.

So, first-round play can now have consequences for the second round.

In particular, the strategies shown ARE a SPNE!

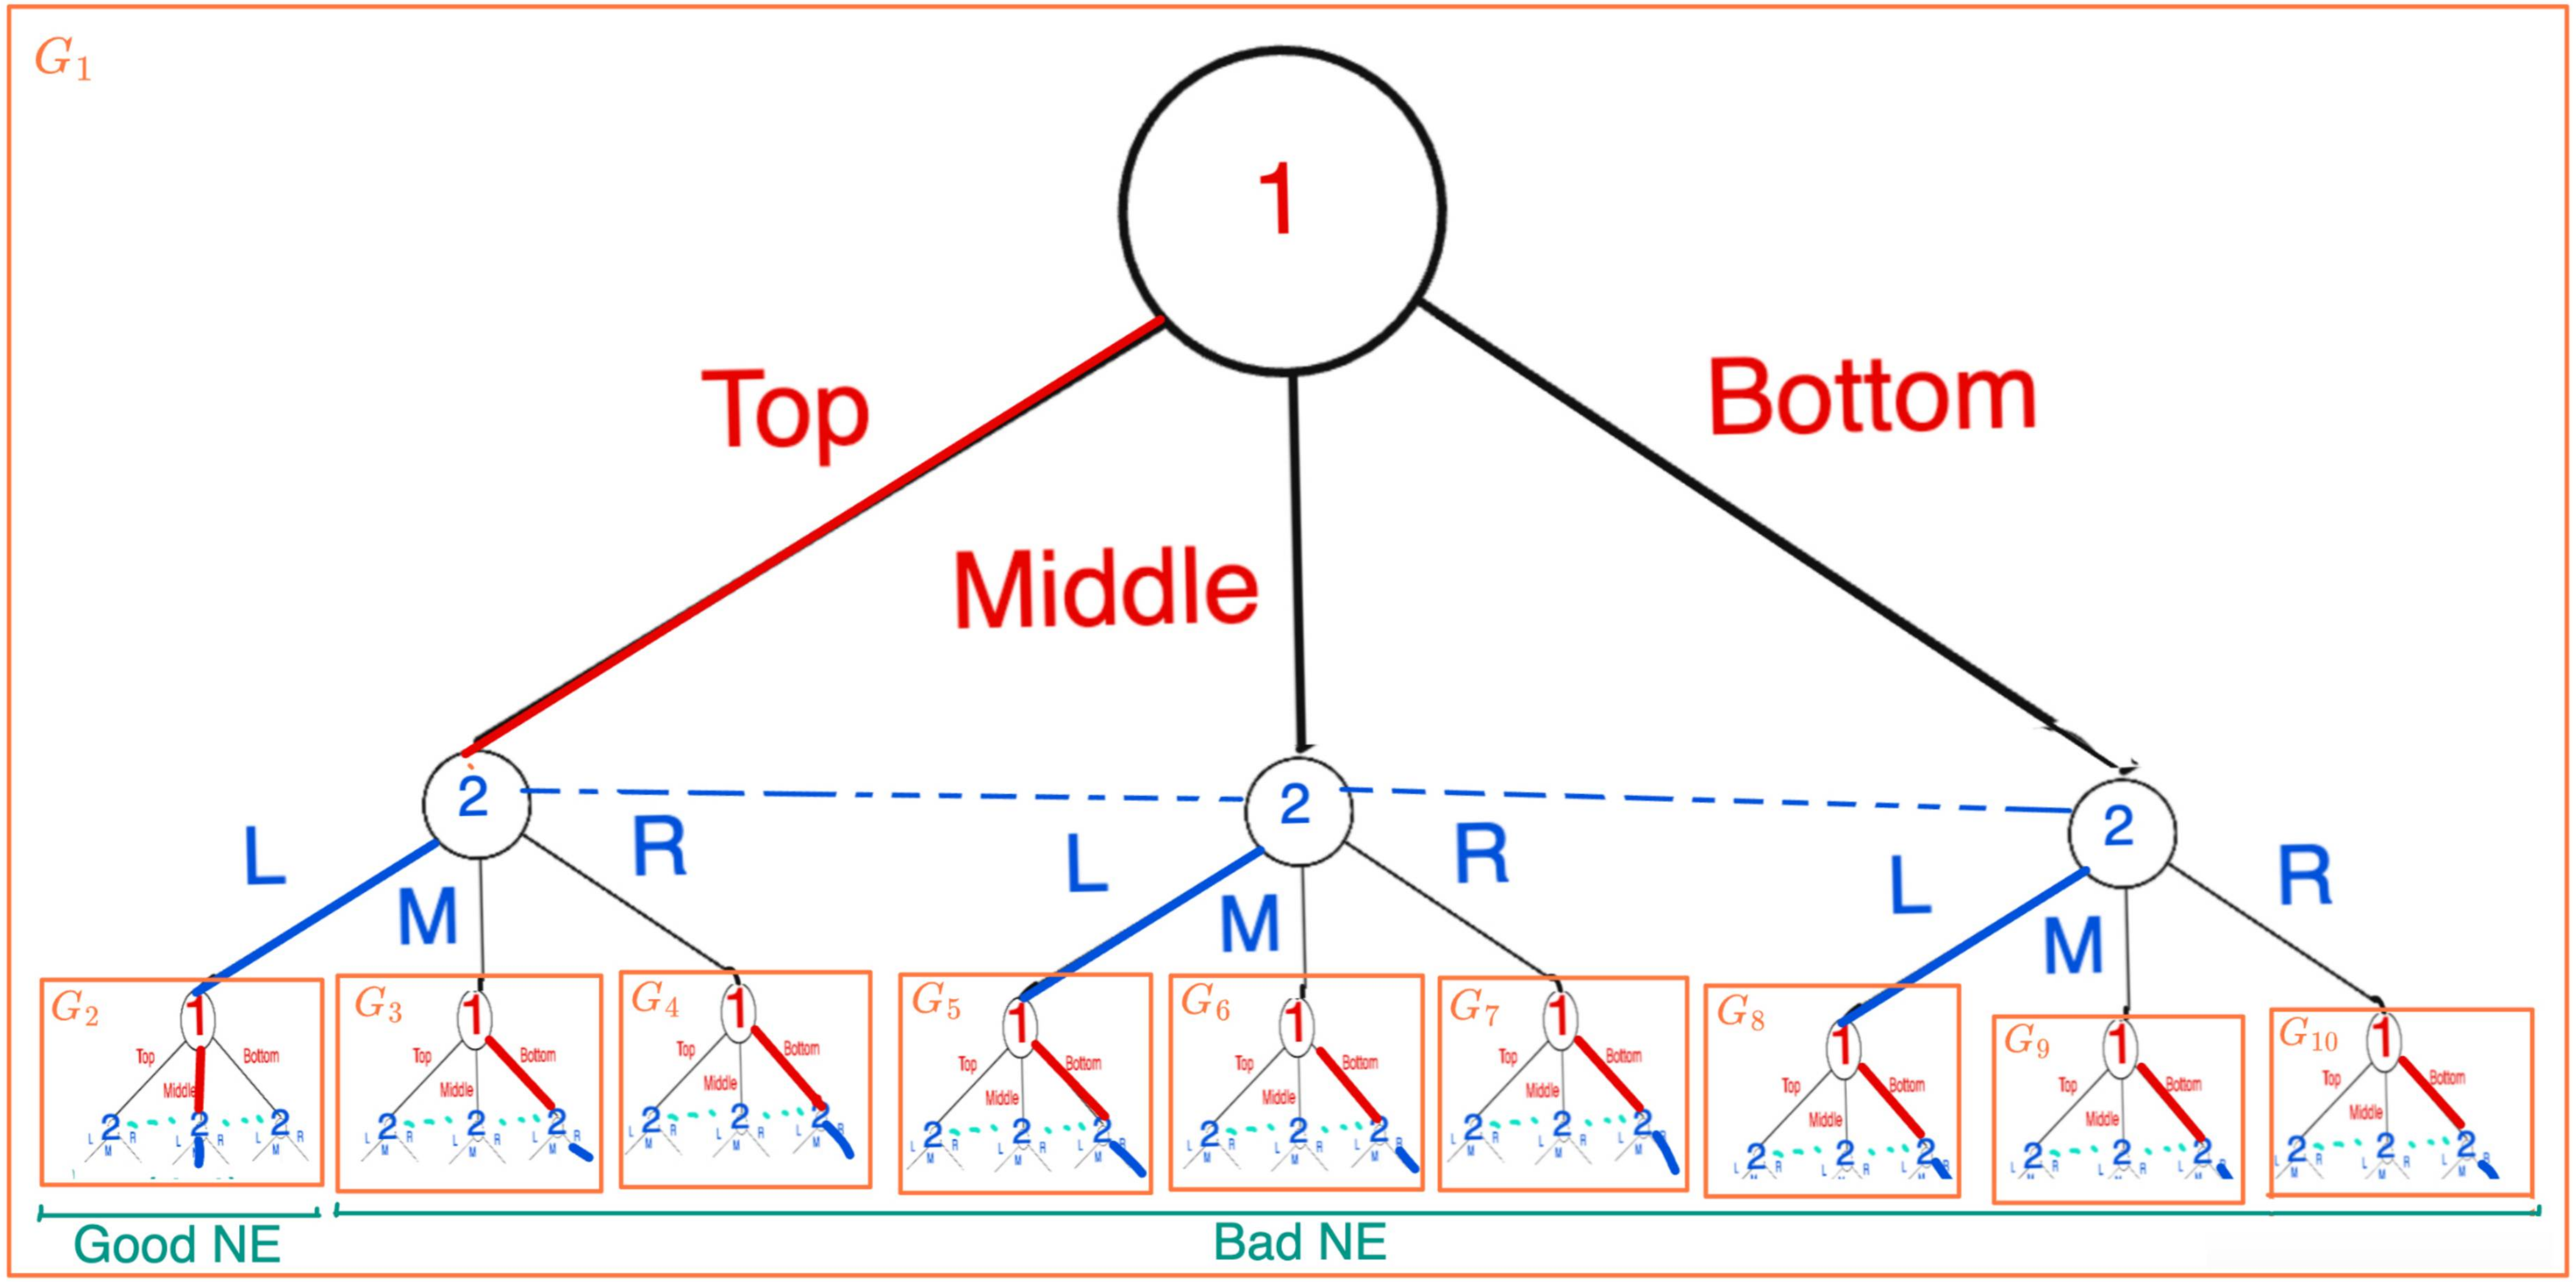

We have 10 subgames: the whole game, and 9 versions of the second round.

Player 1 plays "Top" in the first round, even through in the stage game, "Middle" is the best response to "Left".

They do this to secure the "good" NE in the second round.

The desirable payoff (3,3) can be secured in the first round, even though (Top,Left) is NOT a NE of the stage game.

This is possible since the first round is NOT a subgame by itself.

Finitely Repeated Games -- Multiple NE

Let's do the analysis more careully.

	Left	Middle	Right
Top	3 3 🏆	0 4	−2 0
Middle	4 0	$\boxed{1 \quad 1}$	−2 0
Bottom	0 −2	0 −2	$\boxed{-1 \quad -1}$

Suppose we repeat the "stage game" twice.

Now we have TWO Nash equilibria. So, the argument of the previous slide does NOT apply. That is, it's now NOT CLEAR what will happen in the final round; so, the first round is NOT just like another copy of the second round (it can have consequences).

In the stage game, we have a "good" NE, (Middle,Middle), and a "bad" NE, (Bottom Right).

We must be playing ONE of these in all 9 versions of the second round in any SPNE, because of each of these is a "subgame".

BUT -- crucially, we may see a DIFFERENT NE in the second round, DEPENDING ON what happens in the first round.

That is -- the first round matters! (i.e. for what happens in the second round).

Players may then play a NON-Nash Equilibrium strategy in the first round to AVOID ending up at the "bad" NE in round 2.

Consider the following strategy profile:

With these strategies, play continues (Top,Left), (Middle,Middle), with payoffs 3+1 = 4 for each player.

i.e. in the first round

s_1 : -- Play "Top" in round one
-- In round two, play "Middle" following (Top,Left), but "Bottom" otherwise.

s_2 : -- Play "Left" in round one
-- In round two, play "Middle" following (Top,Left), but "Right" otherwise.

Notice that players do NOT play a Nash Equilibrium of the STAGE game in the first round. But, the first round is NOT a subgame! In fact, this strategy profile is a SPNE. Let's check this.

The subgames are: 1.) each version of the second round (following any history), and 2.) the whole game itself.

In subgames of type 1.), players either play (Middle,Middle), following (Top,Left), or (Bottom,Right) otherwise. These are both NE. So, a NE is played in every subgame of type 1.). ✓

It remains to be check that players play a NE in the main game (both rounds together). We consider a potential deviation for P1. Of course, P2 continues to play s_2.

If P1 deviates, P2 plays "Right" in the second round. P1 may as well play "Bottom" and get -1. So, P1 goes for "Middle" in the first round, getting their maximum payoff of 4 for that round. But, their total payoff is 4+(-1)=3, so it's best NOT to deviate. ✓

Recap: Geometric Series

This theory is needed to find payoffs in infinitely repeated games.

Geometric sequences

The defining feature of a geometric sequence is that you must multiply by a common ratio, r, to get from one term to the next. Geometric sequences are of the form

$$a, \quad ar, \quad ar^2, \quad ar^3, \quad ar^4, \quad \ldots$$

first term "a"
common ratio "r"

where a is the first term in the sequence and r is the common ratio.

- The n[th] term of a geometric sequence is given by: $\boldsymbol{u_n = ar^{n-1}}$

Sum to infinity

The sum to infinity of a geometric sequence only exists for convergent sequences, and is given by:

$$\boldsymbol{S_\infty = \frac{a}{1-r}}$$

Infinitely Repeated Games

Now the stage game is repeated an INFINITE number of times.

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

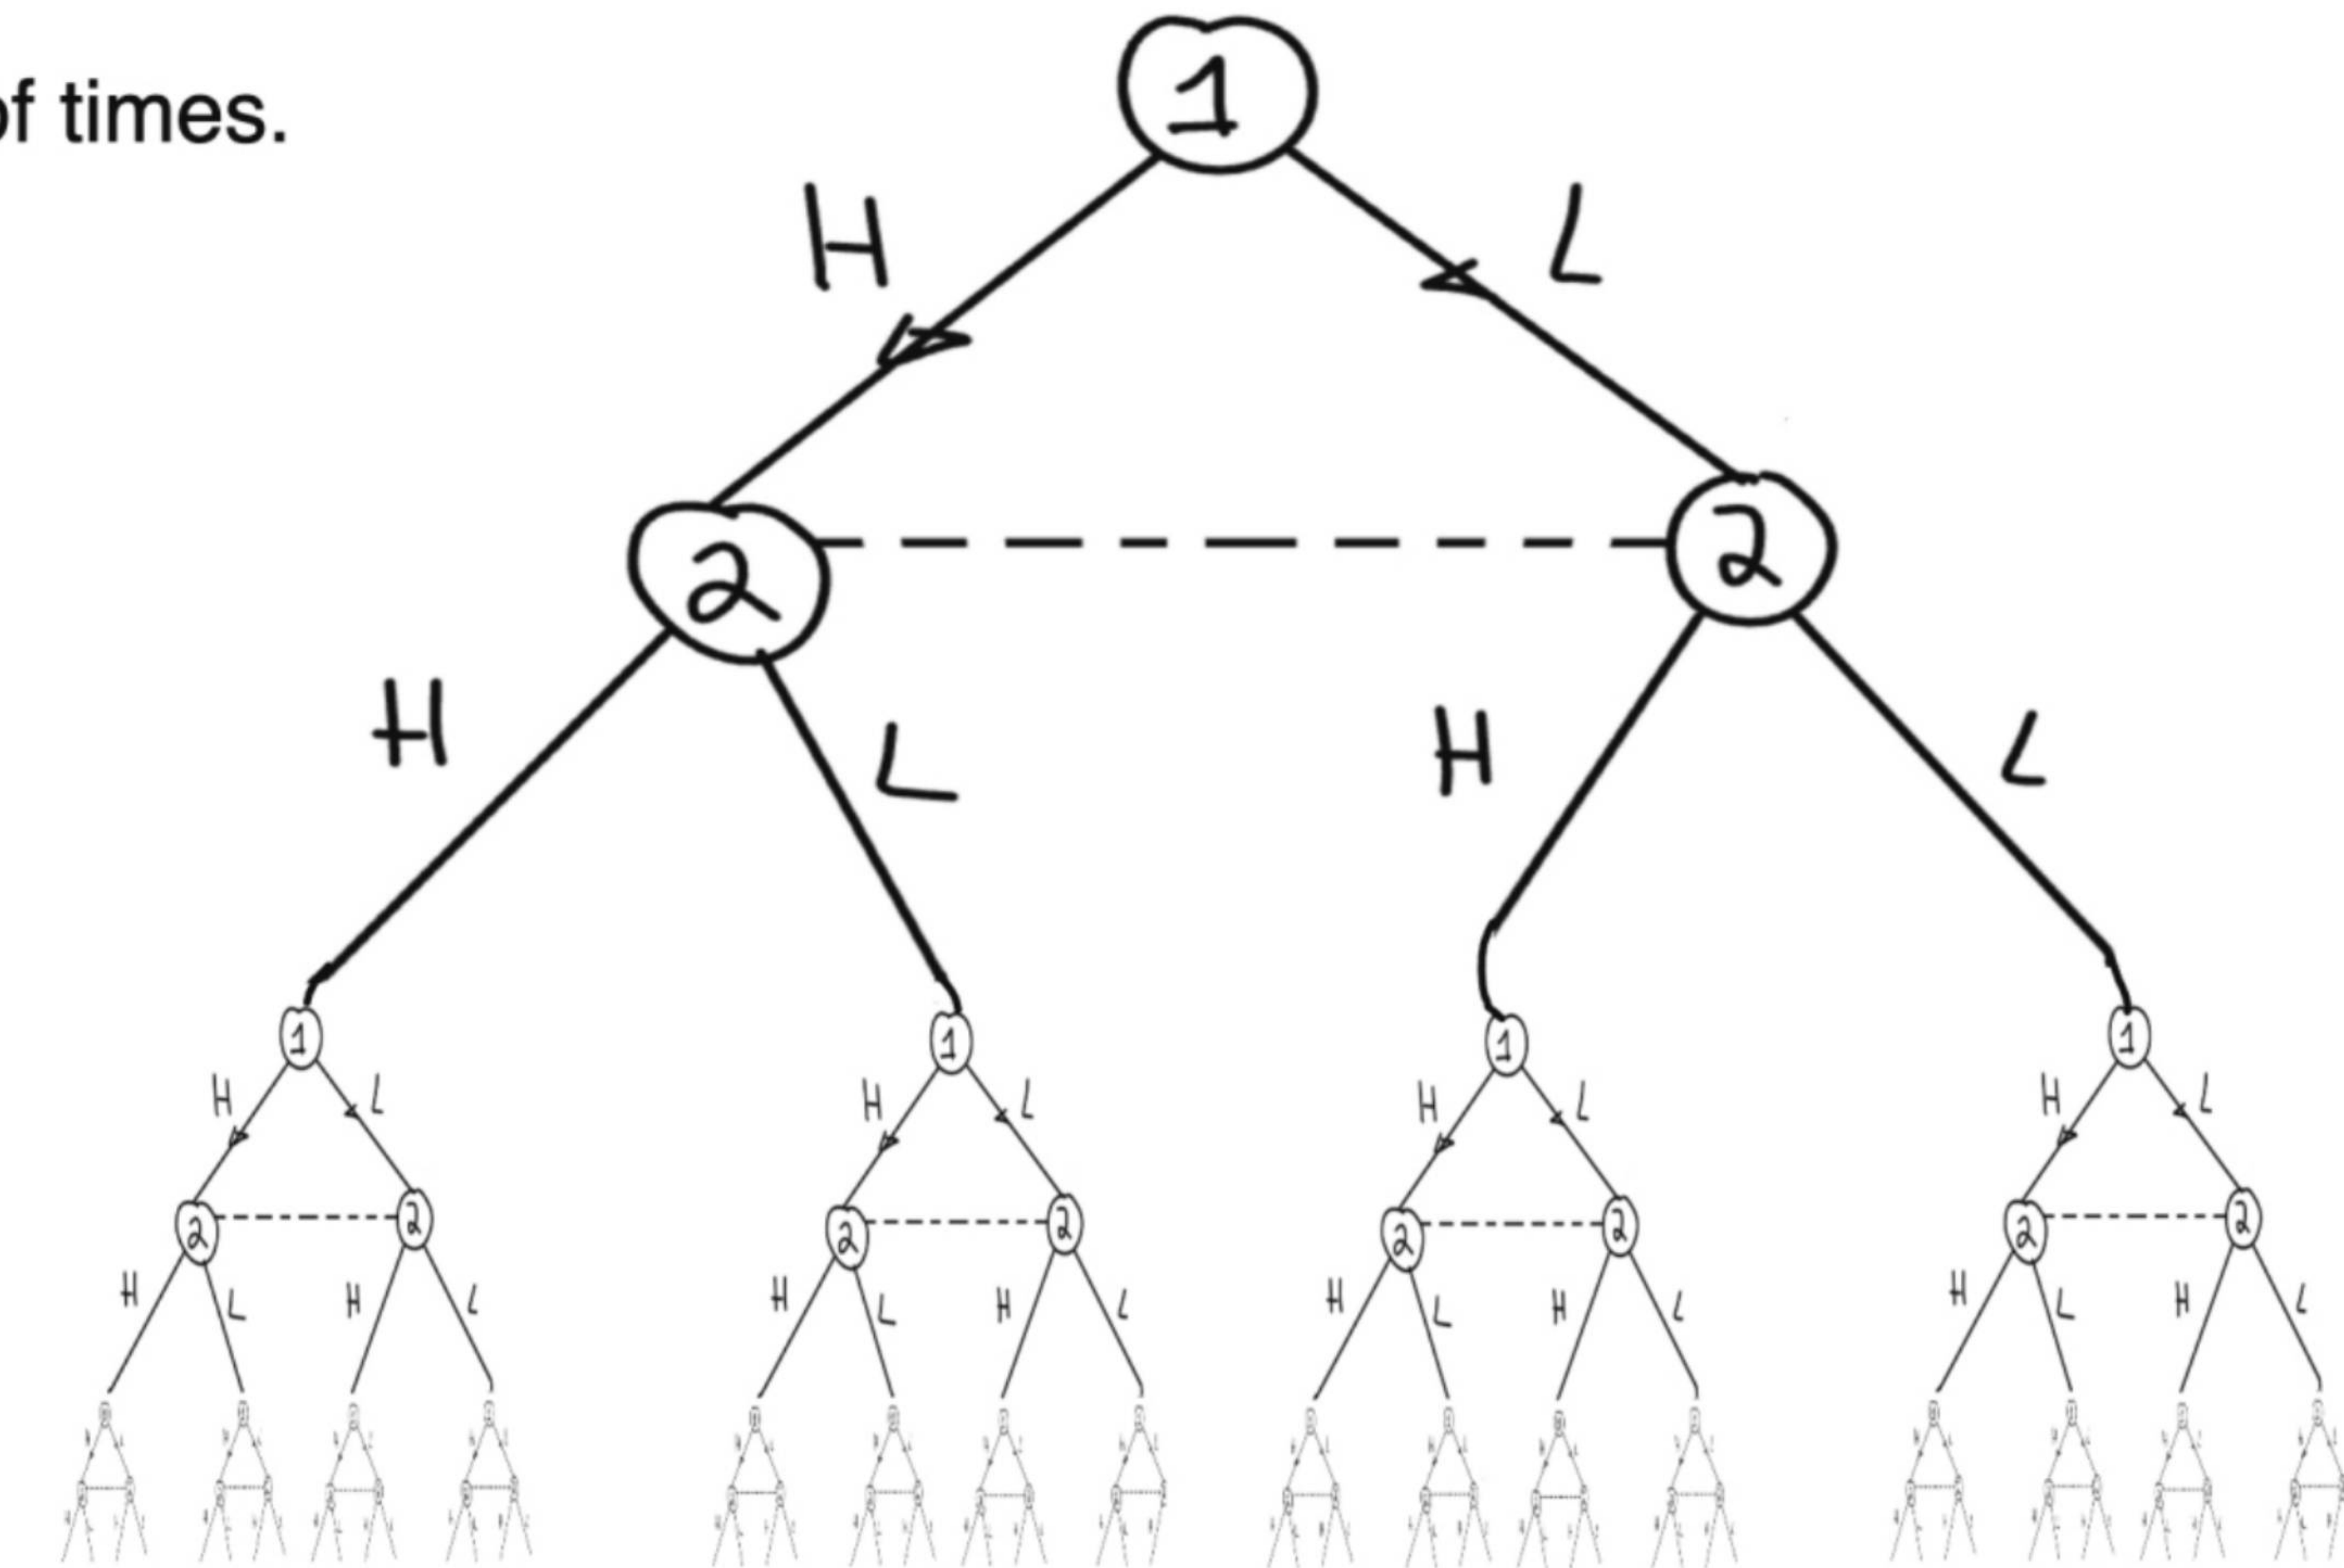

We can't apply the same "backwards induction" argument, because there is NO end to the game.

There is also a problem with the payoffs.
E.g.: suppose again, everyone plays "L" whenever they move.

The payoffs are 1+1+1+.... = ∞.
This is the same as 3+3+3+... = ∞.

This seems a bit silly!

Instead, we use "temporal discounting". We think of the stages as taking place over time, and each period of time later, the payoff is multipled by a number δ, where δ<1.

Now suppose that everyone always plays "L".

The payoffs are now $1 + \delta + \delta^2 + \delta^3 + \ldots = \dfrac{1}{1-\delta}$

It turns out that, in this kind of game, there ARE more interesting subgame perfect equilibria, than the "boring" one where we just play "L" in every period. We think about playing "H" to "be on good terms" with the other player.

Formally, we can show that, if δ is high enough, both players playing the following "grim trigger" strategy counts as a Subgame Perfect Nash Equilibrium:

Play "H" in the first round, and in later rounds if everyone has played (H,H) in all previous rounds.
If anyone has played "L" in any previous round, then we play "L".

The players start off "friends", and cooperate

If anyone "defects", they are enemies forever

Infinitely Repeated Games

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

We first show that both players playing the following strategies is a Nash Equilibrium (no incentive to deviate).

Play "H" in the first round, and in later rounds if everyone has played (H,H) in all previous rounds. If anyone has played "L" in any previous round, then we play "L".

Assuming both players play this strategy, play continues:

(H,H), (H,H), (H,H), ...

The players' payoffs are

$$3 + 3\delta + 3\delta^2 + 3\delta^3 + \ldots = \frac{3}{1-\delta}$$

Now let's consider a possible deviation for P1. Let's assume they begin deviating in the first round. P2 continues with the same strategy.

Play continues: (L,H),(L,L),(L,L),...

$$5 + \delta + \delta^2 + \delta^3 + \ldots = 5 + \frac{\delta}{1-\delta}$$

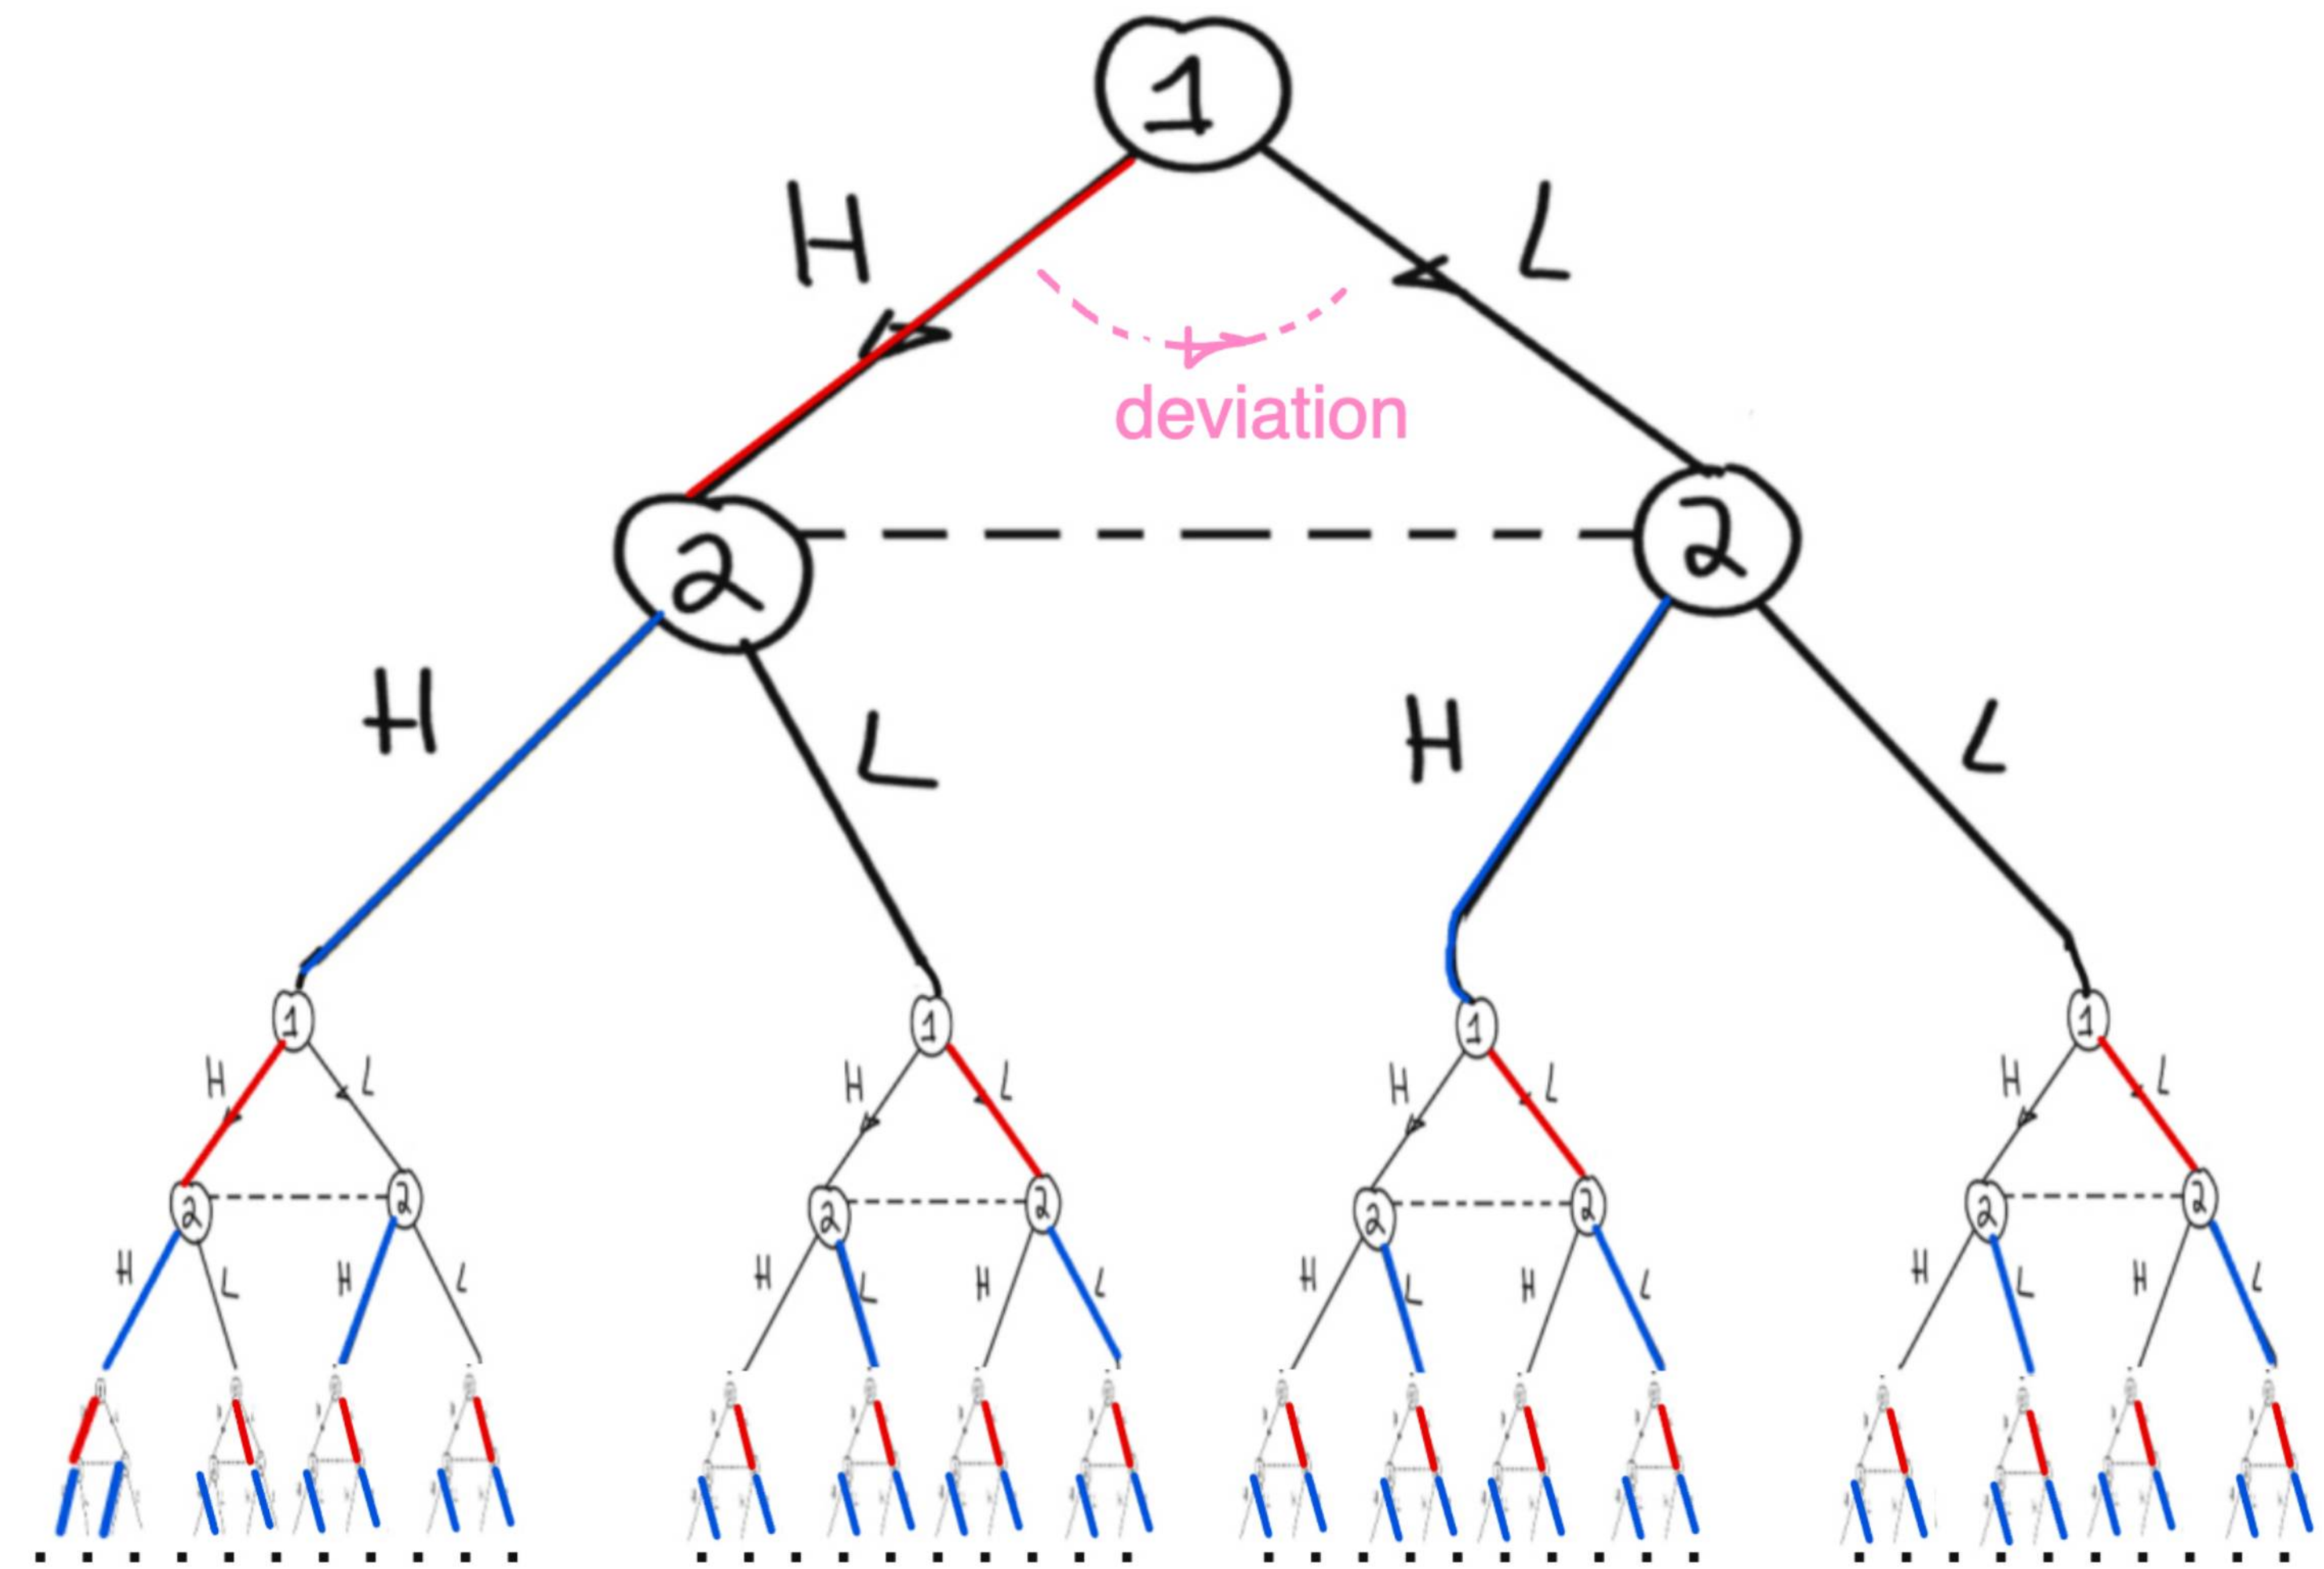

These strategies form a Nash Equilibrium as long as

$$\frac{3}{1-\delta} \geq 5 + \frac{\delta}{1-\delta} \Leftrightarrow 3 \geq 5(1-\delta) + \delta = 5 - 4\delta$$

$$\Leftrightarrow 4\delta \geq 2 \Leftrightarrow \delta \geq 1/2.$$ i.e. we have to be "patient" enough.

In fact, these strategies are Subgame Perfect Equilibria too.

Seeing this requires a bit more theory (see the next page).

$$\underbrace{5}_{\text{initial "bonus"}} + \underbrace{\delta + \delta^2 + \delta^3 + \ldots}_{\text{"punishment"}} = 5 + \frac{\delta}{1-\delta}$$

Infinitely Repeated Games

Now we discuss subgame perfect NE in infinitely repeated games.

What are the subgames?

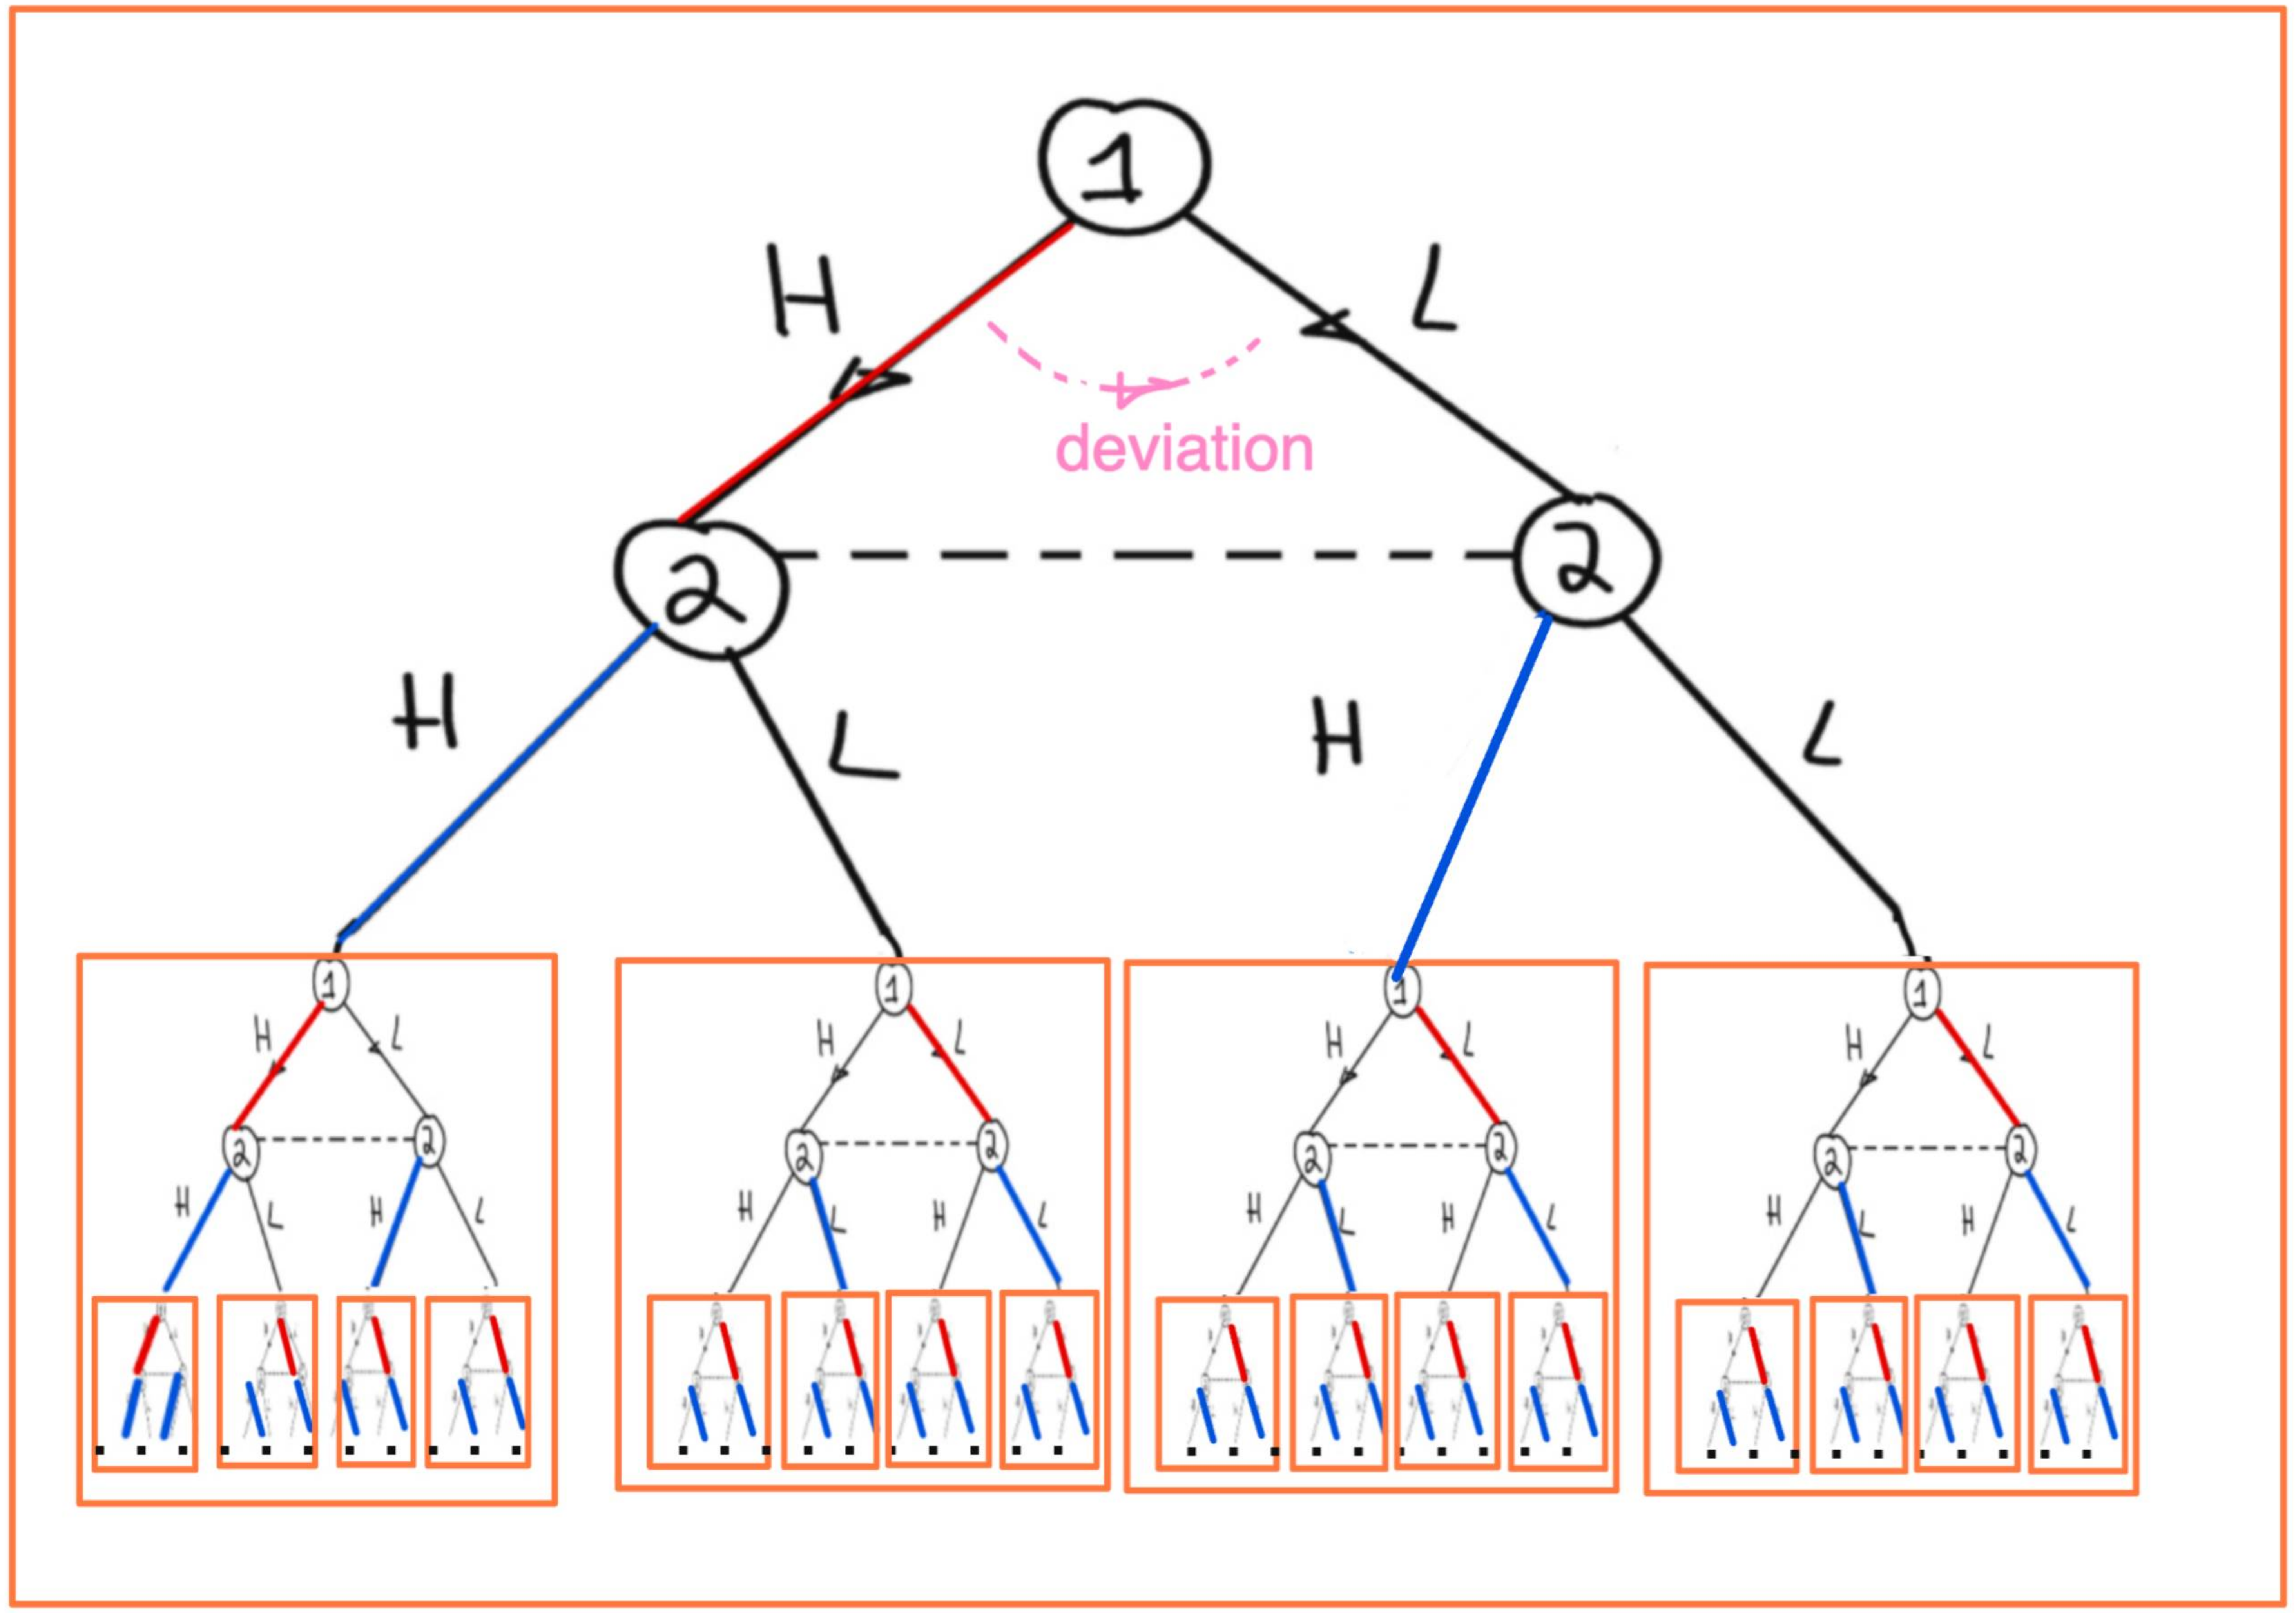

This looks complicated. The "subgames" are just exact copies of the original game! i.e. they go on forever, like a fractal.

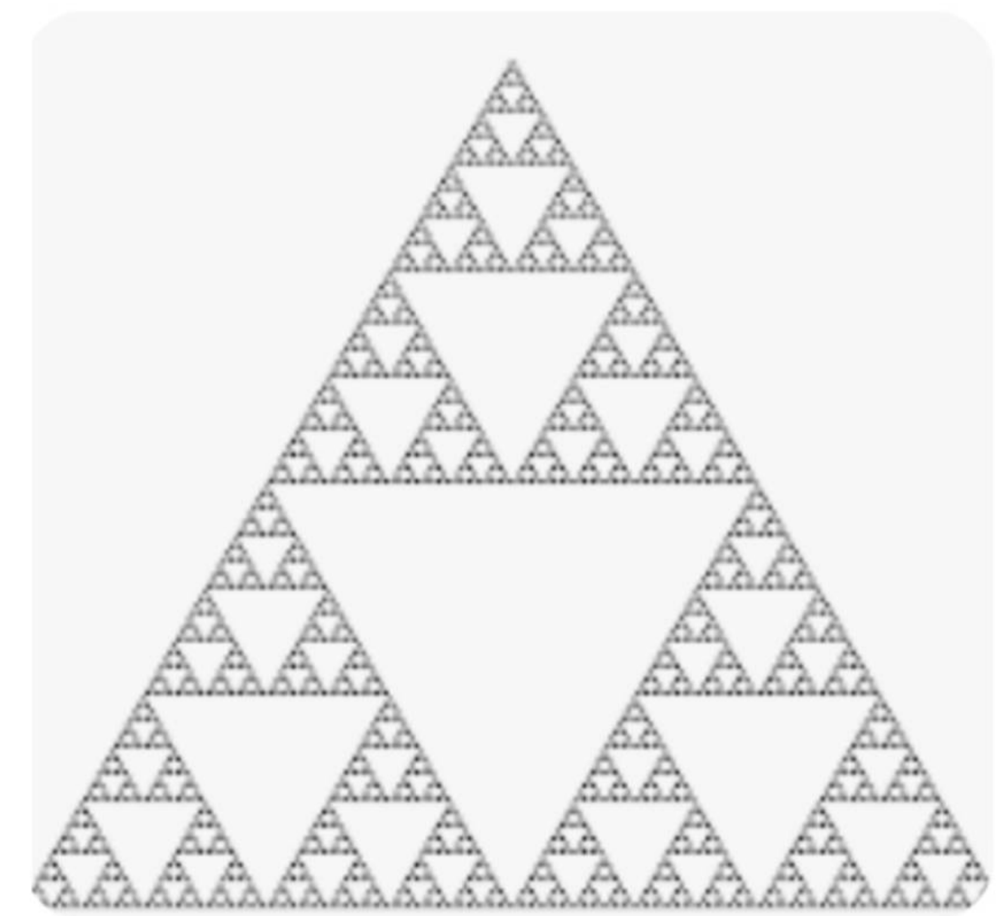

They also follow different histories.

Fortunately, there is a theorem which will enable us to analyse strategy profiles in such games fairly easily, despite this complexity.

Moreover, for the "grim trigger" strategies, there also are basically only "two types" of subgame:

-- those where we're still friends; "L" has never been played
-- those where we've become enemies, "L" has been played.

Infinitely Repeated Games -- One-Deviation Principle

Consider again the following stage game:

Normal/Strategic form (one stage)

	High	Low
High	3, 3	0, 5
Low	5, 0	1, 1

(Low,Low) is a strictly dominant equilibrium (an especially convincing kind of Nash Equilibrium).

And the following strategies for each player:

Play "H" in the first round, and in later rounds if everyone has played (H,H) in all previous rounds.
If anyone has played "L" in any previous round, then we play "L".

("Grim trigger" strategy -- following any deviation from (H,H) we are "no longer friends" and enter the "punishment phase".)

We now check these strategies form a SPNE.

Infinitely-repeated games are COMPLICATED. There are infinite number of subgames to check! BUT, we are helped greatly by the following principle.

One-Deviation Principle.

A pair of strategies form a SPNE in a repeated game, if neither player has an incentive to change ONE ACTION at any stage, holding fixed:

1.) The other player's strategy (as usual), and
2.) The REST OF THEIR OWN strategy.

We call such a deviation a "one-deviation" or "one-stage deviation".

Essentially, we don't need to check complicated deviations; only simple ones.

Moreover, although there are many subgames, from the perspective of one-deviations, there are essentially only TWO TYPES of subgame:

Type 1. Subgames where no-one has played L before (including the first round).
With no deviation, play continues (H,H),(H,H), ... with payoffs $\frac{3}{1-\delta}$ (see previous page)
If P1 one-deviates in the first round of such a subgame, play continues (L,H),(L,L),(L,L),... with payoffs $5+\frac{\delta}{1-\delta}$. This is not profitable when $\delta \geq \frac{1}{2}$
This is the ONLY deviation for P1. After this, P1 "follows the box" (see previous page)

Type 2. Subgames where someone has played L before.
Players play (L,L) forever; no profitable one-devation (or any kind of deviation) as (L,L) is a NE.
"Punishment is credible", as 1>0.

Another Worked Example

Consider the following game which is repeated infinitely many times:

	C	D
C	7, 7	5, 8
D	8, 5	6, 6

Calculate the value of discount factor δ, which helps the GRIM strategy to sustain the cooperative outcome as the Nash equilibrium of the game. (25 marks)

This is the "grim trigger" strategy:

> Play "C" in the first round, and in later rounds if everyone has played (,) in all previous rounds.
> If anyone has played "D" in any previous round, then we play "D".

Consider whether P1 wants to deviate (the analysis for p2 is the same).

Of course, we keep P2's strategy fixed (as the "grim trigger" strategy).
P2 does "what is in the box".

If player 1 continues with the strategy as written, play continues:

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

P1's payoff is $7 + 7\delta + 7\delta^2 + \ldots = \frac{7}{1-\delta}$

overall payoffs from not deviating and deviating

Suppose player 1 deviates. Play continues:

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

P1's payoff is $8 + \underbrace{6\delta + 6\delta^2 + \ldots}_{} = 8 + \frac{6\delta}{1-\delta}$

geometric series;
first term 6δ,
common ratio δ

For a NE, we need to check whether

$$\frac{7}{1-\delta} \geqslant 8 + \frac{6\delta}{1-\delta}$$

i.e. no incentive to deviate

$$\Leftrightarrow \quad 7 \geq 8(1-\delta) + 6\delta = 8 - 2\delta$$

$$\Leftrightarrow \quad 2\delta \geq 1 \qquad \Leftrightarrow \quad \delta \geq 1/2$$

clearing fractions and re-arranging

co-operative outcome

	C	D
C	7, 7	5, 8
D	8, 5	$\boxed{6, 6}$

NE outcome

Further Problems

1.) Consider the following game which is repeated infinitely many times:

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

Calculate the value of discount factor δ, which helps the GRIM strategy to sustain the cooperative outcome as the Nash equilibrium of the game.

2.)

The following question is taken from Tadelis (it is Q10.2 on page 215). Consider the infinitely repeated game with discount factor δ<1 shown below (player 2 is the column player):

	L	*C*	*R*
T	6,6	−1,7	−2,8
M	7,−1	4,4	−1,5
B	8,−2	5,−1	0,0

a. For which values of the discount factor δ can the players support the pair of actions (M,C) played in every period?

b. For which values of the discount factor δ can the players support the pair of actions (T,L) played in every period? Why is you answer different from that for (a)?

3.) Consider an infinitely repeated game with the following stage game.

		Player 2	
		Chicken	Lion
Player 1	Chicken	3, 3	1, 4
	Lion	4, 1	0, 0

Each player maximises the discounted sum of her stage payoff with discount factor δ.

a) For each of the strategy profiles below determine values of δ for which it is a Nash equilibrium of the infinitely repeated game.

   i. Choose Chicken in the first period. Continue to choose Chicken as long as both players have chosen Chicken in all previous periods. If in any period any player chooses Lion, then choose Lion in every subsequent period.

   ii. Choose Chicken in the first period. In subsequent periods, choose the action that the other player has chosen in the previous period.

b) Suppose $\delta = 0.5$. Determine if the following strategy profile is a Subgame Perfect Nash equilibrium of the infinitely repeated game. Verify your answer.

   There are two modes: Cooperation and Fight. The game starts in the Cooperation mode. In the Cooperation mode, each player plays Chicken. If both players play Chicken, then they remain in the Cooperation mode in the next period; otherwise they go to the Fight mode in the next period. In the Fight mode, each player plays Lion. Regardless of the actions that they choose they go back to the Cooperation mode in the following period.

c) Consider the following stage game.

		Player 2		
		A	B	C
Player 1	X	6,0	0,0	0,1
	Y	0,0	2, -1	9,1

Find all Subgame Perfect Nash equilibria of the 8-times repeated version of this game. Explain your answer.

Further Problems

4.)

A manager and an employee are playing a game where actions and payoffs are given in Figure 1. For each profile of actions, the first number is the payoff to the employee, while the second is the payoff to the manager.

		Manager: Monitor	Manager: No Monitor
Employee	Work	10 , 5	10 , 10
	Shirk	0 , -1	20 , -10

Figure 1

a. Suppose that the game is played once and players choose actions simultaneously. Describe the employee's Best Response to every strategy of the manager.

b. Suppose that the game is played once and players choose actions simultaneously. Find ALL pure-strategy Nash equilibria and ALL mixed-strategy Nash equilibria.

c. Suppose that the employee and the manager repeat the game four times, where each time they choose actions simultaneously and both observe the outcome. How many different pure strategies does the employee have?

d. Refer to part c. If players look forward and reason back, what would you expect them to play in each period? Compare your answer to part b where they play just once.

5.)

Consider the following Tournament between Aditi and Keiko with payoffs given in Figure 1 below. Each action 1 through 5 indicates the level of engagement by each player, while for each profile of actions the first number is the payoff to Aditi, while the second is the payoff to Keiko.

		Keiko: 1	2	3	4	5
Aditi	1	5, 5	2, 7	1, 9	0,10	0,9
	2	7, 4	4, 4	2, 5	1,7	0,8
	3	9,1	5,2	3, 3	2,5	2,6
	4	10,0	7,1	5,2	2,2	1,5
	5	9,0	8,0	6,2	5,1	1,1

Figure 1

a. Suppose that the game is played once, and players choose actions simultaneously. Derive all pure strategy Nash Equilibria of this game.

b. Suppose that players have available all 5 levels of engagement and choose sequentially, with the second mover observing the choice of the first mover. Focusing on equilibria in which players look forward and reason back, what would Aditi play if she moves first? What would Keiko play if she moves first? How many different strategies does Keiko and Aditi have if Keiko moves first?

c. Suppose that players play the simultaneous game in two rounds, but they have less options in the second round. Namely, Aditi and Keiko simultaneously choose an engagement level in the first round, after which they choose simultaneously in the second round but engagement levels 4 and 5 are not available to either player in the second round. How many different strategies does Aditi have in this two-period game?