1. Introduction to Game Theory

What is game theory?

- Important branch of applied mathematics / economics
  - Eight game theorists have won the Nobel prize, most notably John Nash (subject of “Beautiful mind” movie)
- Relevant to many other fields (biology, computer science, ...)
- Studies the behaviour of agents in strategic situations
- Uses games as mathematical model of these strategic situations
- Proposes ways of analyzing games to identify what choices rational agents will make
- Many different types of games have been proposed and studied
  - we consider simplest type: strategic (aka normal form) games

Formal definition of a strategic game

An strategic game consists of:

- a finite set $P = \{1, \ldots, n\}$ of players
- for each player i, a set A_i of possible actions
  - $A = A_1 \times \ldots \times A_n$ is the set of action profiles, also called outcomes
- for each player i, a utility function $u_i : A \to \mathbb{R}$, expressing the preferences of the player for the possible outcomes of the game
  - players aim to maximize their utility

Two player games: matrix form

When there are only two players, we can represent strategic games using matrices.

Each column represents an action for Player 2

Utility values for action profile (A,C)

u_1: player 1
u_2: player 2

	C	D	E
A	u_1, u_2		
B			

Each row represents an action for Player 1

Rock-paper-scissors

Rock beats scissors, scissors beats paper, paper beats rock

	Rock	Paper	Scissors
Rock	0 , 0	1 , -1	-1 , 1
Paper	1 , -1	0 , 0	1 , -1
Scissors	-1 , 1	1 , -1	0 , 0

Example of a zero-sum game: players utilities sum to zero.

Prisoner's dilemma

Two suspects in a major crime are held and interrogated separately.

- If both confess, they get 3 years of prison each.
- If only one confesses, the confessor goes free and the other is sentenced to 4 years of prison.
- If neither confesses, both get 1 year of prison.

	confess	don't confess
confess	-3 , -3	0 , -4
don't confess	-4 , 0	-1 , -1

Stag hunt

Two people are out hunting. If an individual hunts a stag, she must have the cooperation of her partner in order to succeed. Either hunter can get a hare by herself, but a hare is worth less than a stag.

	stag	hare
stag	4 , 4	0 , 3
hare	3 , 0	3 , 3

Bach or Stravinsky

Two people wish to go together to see a classical music concert of either Bach or Stravinsky. They have different preferences regarding which concert to attend: the first would prefer to see Bach, while the second prefers Stravinsky.

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

Coordination game

Assume that two drivers meet on a narrow dirt road. Both have to swerve in order to avoid a head-on collision. If both swerve to different sides they will pass each other, but if they choose the same side they will collide.

	Right	Left
Right	0 , 0	-50 , -50
Left	-50 ,-50	0 , 0

Matching pennies

Both players have the same possible actions: Heads and Tails.
If both players select the same action, then Player 1 wins,
else Player 2 wins.

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

Analyzing games

We now introduce different notions which will help us to formalize what it means for an action to be a good choice for a player, or for an action profile to be “stable”.

Pareto-optimality: cannot be uniformly improved

Dominant strategies: best choice for a player

Dominated strategies: strategies to be avoided

Nash equilibrium: players have no regrets about their choices

Some notation and terminology

We will first consider pure strategies, in which each player selects a single action to play, and later we will introduce mixed strategies, in which actions are selected probabilistically.

- S_i is the strategy set of player i
- $S = S_1 \times \ldots \times S_n$ is the set of strategy profiles
- S_{-i} is the set of strategy profiles for $P \setminus \{i\}$
- if s is a strategy profile, $s_i \in S_i$ is the strategy in s assigned to player i, and $s_{-i} \in S_{-i}$ is the strategy profile for the other players
- (s', s_{-i}) is a shorthand for $(s_1, \ldots, s_{i-1}, s', s_{i+1}, \ldots, s_n)$

For the moment, we consider only pure strategies, which means $S_i = A_i$ and $S = A$.

Pareto-optimality

A strategy profile s Pareto-dominates a strategy profile s' if for all $i \in P$, $u_i(s) \geq u_i(s')$, and there exists some $j \in P$ such that $u_j(s) > u_j(s')$.

A strategy profile s is Pareto-optimal a strategy if there does not exist another strategy profile s' such that s' Pareto-dominates s.

- Note that every game must have at least one Pareto-optimal strategy profile, but some games may have more than one.
  - In zero-sum games, all outcomes are Pareto-optimal.

Examples of Pareto-optimality

Pareto-optimal outcomes in yellow.

Stag Hunt:

	stag	hare
stag	4 , 4	0 , 3
hare	3 , 0	3 , 3

Prisoner's Dilemma:

	confess	don't confess
confess	-3 , -3	0 , -4
don't confess	-4 , 0	-1 , -1

Dominant strategies

For a player i, a strategy $s^* \in S_i$ strictly dominates $s' \in S_i$ if for all $s_{-i} \in S_{-i}$ we have

$$u_i(s^*, s_{-i}) > u_i(s', s_{-i})$$

A strategy $s^* \in S_i$ weakly dominates a strategy $s' \in S_i$ if for all $s_{-i} \in S_{-i}$ we have

$$u_i(s^*, s_{-i}) \geq u_i(s', s_{-i})$$

and for some $s_{-i} \in S_i$, we have $u_i(s^*, s_{-i}) > u_i(s', s_{-i})$.

If s^* dominates every other strategy in S_i, then s^* is said to be a dominant strategy for player i.

★ dominant strategy = unique best choice

Example: Dominant strategies

Dominant strategy solutions in yellow.

Stag Hunt:

	stag	hare
stag	4 , 4	0 , 3
hare	3 , 0	3 , 3

There are no dominant strategies in the Stag Hunt game.

Prisoner's Dilemma:

	confess	don't confess
confess	-3 , -3	0 , -4
don't confess	-4 , 0	-1 , -1

Confess is a dominant strategy for both players, and gives a **dominant strategy solution**.

Dominated strategies

A strategy s is said to be strictly dominated if there is another strategy which strictly dominates it.

Strictly dominated strategies are "bad moves": no reason for a rational player to use such strategies.

Note that strictly dominated strategies can exist even if there is no dominant strategy.

Iterated elimination of strictly dominated strategies (IESDS) is a procedure by which we remove strictly dominated strategies from a game until no such strategies remain.

It allows us to simplify a game, and sometimes identify a solution.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

Are any of player 1’s strategies strictly dominated ?

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

Strategies B and C are strictly dominated for player 1.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

Strategies B and C are strictly dominated for player 1.

... so we can simplify the game by assuming these strategies won't be played.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

In the simplified game, E is strictly dominated by F.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

In the simplified game, E is strictly dominated by F.

... so we can simplify the game by assuming strategy E won't be played.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

In the reduced game, strategy A is strictly dominated by D, so it also can be removed.

Example of IESDS

	E	F
A	5 , 2	4 , 3
B	3 , 6	3 , 2
C	2 , 1	4 , 1
D	4 , 3	5 , 4

In the reduced game, strategy A is strictly dominated by D, so it also can be removed.

Expected outcome is (D,F).

Nash equilibrium

A strategy profile is a Nash equilibrium if no player would benefit from changing his strategy unilaterally.

In other words, no player regrets his choice.

Formally:

A strategy profile $x^* \in S$ is a Nash equilibrium if for all players $i \in P$ and all $y_i \in S_i$ such that $y_i \neq x_i^*$, we have

$$u_i(x_i^*, x_{-i}^*) \geq u_i(y_i, x_{-i}^*)$$

To define Strict Nash equilibrium, we replace $\geq$ by $>$ above.

Weak Nash equilibrium: for some $y_i \in S_i$, equality holds.

Best responses

For $s_{-i} \in S_{-i}$, we define $\mathrm{BR}_i(s_{-i})$ as the set of player i's best strategies when the other players play s_{-i}:

$$\mathrm{BR}_i(s_{-i}) = \{s^* \in S_i : u_i(s^*, s_{-i}) \geq u_i(s', s_{-i}) \text{ for all } s' \in S_i\}$$

We call BR_i the best-response function of player i.

Another way to define Nash equilibria:

$$x^* \text{ is a Nash equilibrium iff } x_i^* \in \mathrm{BR}_i(x_{-i}^*) \text{ for all } i \in P$$

Every player's strategy is a best response to the strategies of the other players.

Examples: Nash equilibrium

Nash equilibria in yellow.

Stag Hunt:

	stag	hare
stag	4 , 4	0 , 3
hare	3 , 0	3 , 3

Two Nash equilibria, only one of which is Pareto-optimal.

Prisoner's Dilemma:

	confess	don't confess
confess	-3 , -3	0 , -4
don't confess	-4 , 0	-1 , -1

The dominant strategy solution is the only Nash equilibrium.

In fact, **dominant strategy solutions always give unique Nash equilibria.**

Examples: Nash equilibrium

Nash equilibria in yellow.

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

Two Nash equilibria, capturing the two “stable outcomes”.

Matching Pennies:

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

No Nash equilibria, capturing the lack of “stable outcomes”.

Mixed strategies

In games like "Matching pennies" or "Rock-paper-scissors", it seems natural to allow players to make random moves.

A mixed strategy for a player i is a probability distribution over the set A_i, i.e. it is a function s_i such that: $0 \leq s_i(b) \leq 1$ for all $b \in A_i$, and $\Sigma_{b \in A_i} s_i(b) = 1$.

We can extend the players' utility functions to mixed strategy profiles as follows:

$$u_i(s) = \sum_{a \in A} \left(u_i(a) \cdot \prod_{j=1}^{n} s_j(a_j) \right)$$

here *s* is mixed strategy profile

consider all outcomes

utility of outcome

probability of this outcome

Mixed strategy Nash equilibrium

Nash equilibrium for mixed strategy profiles is defined the same way as for pure strategy profiles, except that we now allow players to use mixed strategies.

All pure strategy Nash equilibria are still Nash equilibria when we allow mixed strategies.

But what is interesting is that by considering mixed strategies, we can find new Nash equilibria.

Let's see this on an example.

Example: Mixed strategy Nash equilibrium

Matching Pennies:

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

Both players choosing the mixed strategy 1/2 Heads, 1/2 Tails is the unique Nash Equilibrium.

Example: Mixed strategy Nash equilibrium

Matching Pennies:

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

Both players choosing the mixed strategy 1/2 Heads, 1/2 Tails is the unique Nash Equilibrium.

Why a Nash equilibrium?

No matter what mixed strategy a player adopts, his expected utility is 0 if the other player uses 1/2 Heads, 1/2 Tails. So he has no incentive to change strategies.

Example: Mixed strategy Nash equilibrium

Matching Pennies:

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

Both players choosing the mixed strategy 1/2 Heads, 1/2 Tails is the unique Nash Equilibrium.

Why a Nash equilibrium?

No matter what mixed strategy a player adopts, his expected utility is 0 if the other player uses 1/2 Heads, 1/2 Tails. So he has no incentive to change strategies.

Why unique?

Suppose player 1 plays a mixed strategy $(\underset{\text{Heads}}{p}, \underset{\text{Tails}}{1-p})$, where $p > \frac{1}{2}$.

If player 2 does not play $(0, 1)$, then would prefer to play $(0, 1)$.

If player 2 does play $(0, 1)$, then player 1 should switch to $(1, 0)$.

Similar argument if $p < \frac{1}{2}$, or if player 2 is the one using a different strategy.

Example: Mixed strategy Nash equilibrium

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

This game has two pure Nash Equilibria (B,B) and (S,S), but it also has one mixed Nash equilibrium.

Example: Mixed strategy Nash equilibrium

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

This game has two pure Nash Equilibria (B,B) and (S,S), but it also has one mixed Nash equilibrium.

Finding the mixed Nash equilibrium:

Suppose (α_1, α_2) is a Nash equilibrium, where $\alpha_1 = (p, 1-p)$ and $\alpha_2 = (q, 1-q)$.

Example: Mixed strategy Nash equilibrium

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

This game has two pure Nash Equilibria (B,B) and (S,S), but it also has one mixed Nash equilibrium.

Finding the mixed Nash equilibrium:

Suppose (α_1, α_2) is a Nash equilibrium, where $\alpha_1 = (p, 1-p)$ and $\alpha_2 = (q, 1-q)$.

Need player 1 to be indifferent between B and S, i.e. the expected utilities should be the same:

$$u_1(B, \alpha_2) = 2q = 1 - q = u_1(S, \alpha_2)$$

So q must be equal to $\frac{1}{3}$.

Example: Mixed strategy Nash equilibrium

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

This game has two pure Nash Equilibria (B,B) and (S,S), but it also has one mixed Nash equilibrium.

Finding the mixed Nash equilibrium:

Suppose (α_1, α_2) is a Nash equilibrium, where $\alpha_1 = (p, 1-p)$ and $\alpha_2 = (q, 1-q)$.

Need player 1 to be indifferent between B and S, i.e. the expected utilities should be the same:

$$u_1(B, \alpha_2) = 2q = 1 - q = u_1(S, \alpha_2)$$

So q must be equal to $\frac{1}{3}$.

Player 2 must also be indifferent:

$$u_2(\alpha_1, B) = p = 2(1-p) = u_2(\alpha_1, S)$$

So we get $p = \frac{2}{3}$.

Example: Mixed strategy Nash equilibrium

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

This game has two pure Nash Equilibria (B,B) and (S,S), but it also has one mixed Nash equilibrium.

Finding the mixed Nash equilibrium:

Suppose (α_1, α_2) is a Nash equilibrium, where $\alpha_1 = (p, 1-p)$ and $\alpha_2 = (q, 1-q)$.

Need player 1 to be indifferent between B and S, i.e. the expected utilities should be the same:

$$u_1(B, \alpha_2) = 2q = 1 - q = u_1(S, \alpha_2)$$

So q must be equal to $\frac{1}{3}$.

Player 2 must also be indifferent:

$$u_2(\alpha_1, B) = p = 2(1-p) = u_2(\alpha_1, S)$$

So we get $p = \frac{2}{3}$.

$$\alpha_1 = (\tfrac{2}{3}, \tfrac{1}{3}) \qquad \alpha_2 = (\tfrac{1}{3}, \tfrac{2}{3})$$

Existence of Nash equilibria

We know that some games do not admit any pure strategy Nash equilibria, but what if we allow mixed strategies?

Nash's Theorem:
Every finite strategic game has a mixed strategy Nash equilibrium.

Nash's Theorem is one of the most important results in game theory. Nash won a Nobel prize in economics for his work.

In the statement of the theorem, "finite" means that each player has only finitely many actions.

The finiteness condition is important, since games with infinite action sets may have no Nash equilibria.

Correlated equilibrium

Suppose two drivers arrive at an intersection. They both have the option to stop or to continue. If both continue, they collide, and if both wait, they don't get anywhere.

Solution for this coordination problem in real life? Traffic lights.

Note that nothing forces the drivers to follow the light, but it is generally in their best interest to do so.

Idea: an independent third-party selects (probabilistically) an action profile and tells each player what action she should play.

Correlated equilibrium if no player wishes to use a different action assuming the other players all follow the recommendation.

Correlated equilibrium

A correlated equilibrium is a probability distribution p over the set A of action profiles such that for every player $i \in P$ and every pair of actions $a_i, a_i' \in A_i$, we have:

$$\sum_{a_{-i} \in A_{-i}} p(a_i, a_{-i}) \cdot u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} p(a_i, a_{-i}) \cdot u_i(a_i', a_{-i})$$

Correlated equilibria are a generalization of Nash equilibria: All Nash equilibria are correlated equilibria, but correlated equilibria do not necessarily correspond to Nash equilibria.

Example: Bach or Stravinsky

Bach or Stravinsky:

	Bach	Stravinsky
Bach	2 , 1	0 , 0
Stravinsky	0 , 0	1 , 2

Claim: $p(BB) = p(SS) = \frac{1}{2}$ is a correlated equilibrium.

Utility: $(\frac{3}{2}, \frac{3}{2})$

- if 1 is told to play B, knows 2 will play B, so B is best option.
- if 1 is told to play S, knows 2 will play S, so S is best option.
- if 2 is told to play B, knows 1 will play B, so B is best option.
- if 2 is told to play S, knows 1 will play S, so S is best option.

Example: Traffic game

Suppose we model the traffic game as follows, with actions S (for stop) and C (for continue).

	S	C
S	4 , 4	1 , 5
C	5 , 1	0 , 0

Three Nash equilibria:

- pure strategy profile (C, S) with expected utility vector $(5, 1)$
- pure strategy profile (S, C) with expected utility vector $(1, 5)$
- mixed strategy profile $(\frac{1}{2}, \frac{1}{2})$ with expected utility vector $(2.5, 2.5)$

Example: Traffic game

Suppose we model the traffic game as follows, with actions S (for stop) and C (for continue).

	S	C
S	4 , 4	1 , 5
C	5 , 1	0 , 0

p is correlated equilibrium is following constraints satisfied:

- $4 \cdot p(SS) + 1 \cdot p(SC) \geq 5 \cdot p(SS) + 0 \cdot p(SC)$
- $5 \cdot p(CS) + 0 \cdot p(CC) \geq 4 \cdot p(CS) + 1 \cdot p(CC)$
- $4 \cdot p(SS) + 1 \cdot p(CS) \geq 5 \cdot p(SS) + 0 \cdot p(CS)$
- $5 \cdot p(SC) + 0 \cdot p(CC) \geq 4 \cdot p(SC) + 1 \cdot p(CC)$

Example: Traffic game

Suppose we model the traffic game as follows, with actions S (for stop) and C (for continue).

	S	C
S	4 , 4	1 , 5
C	5 , 1	0 , 0

Equivalently:

- $p(SC) \geq p(SS)$
- $p(CS) \geq p(CC)$
- $p(CS) \geq p(SS)$
- $p(SC) \geq p(CC)$

Possible solution:

$$p(CC) = 0.1, p(SS) = 0.2,$$
$$p(SC) = 0.3, p(CS) = 0.4$$

with expected utility $(3.1, 2.7)$

Example: Traffic game

Suppose we model the traffic game as follows, with actions S (for stop) and C (for continue).

	S	C
S	4 , 4	1 , 5
C	5 , 1	0 , 0

Equivalently:

- $p(SC) \geq p(SS)$
- $p(CS) \geq p(CC)$
- $p(CS) \geq p(SS)$
- $p(SC) \geq p(CC)$

Another possible solution:

$$p(SC) = p(SS) = p(CS) = \frac{1}{3}$$

with expected utility $(\frac{10}{3}, \frac{10}{3})$

Maxmin strategies

What if we don't know the other players are rational ?
Or what if they are purposely trying to hurt us ?
What is the "safest" way for us to play ?

The maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$.

maximum utility which player *i* can guarantee no matter what the other players do

A maxmin strategy of player i is any $s_i \in S_i$ such that $\min_{s_{-i}} u_i(s_i, s_{-i})$ is i's maxmin value.

pick a strategy which gives the maxmin value

Minmax strategies

Now we consider the opposite situation, in which we want to choose a strategy which will minimize the utility of our adversary.

Here we assume there are only two players, but the definition can be generalized to an arbitrary number.

The minmax value for player $-i$ is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

assume player -*i* will try to maximize his utility,
player *i* wants to make maximum as low as possible

A minmax strategy for player i against player $-i$ is a strategy $s_i \in S_i$ such that $\max_{s_{-i}} u_{-i}(s_i, s_{-i})$ is $-i$'s minmax value.

choose strategy for player *i* which makes
player *-i's* utility as low as possible

Minmax Theorem

The following result was proven by von Neumann in 1928:

Minmax Theorem. In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

This theorem tells us three things:

1. A player's minmax and maxmin values are always the same.

2. A player's minmax and maxmin strategies are the same.

3. In every Nash equilibrium, a player's utility is exactly his maximin = minmax value. So all Nash equilibria give the same utilities.

Extensive form games

We now consider games with sequential moves.

An extensive form game consists of:

- a finite set $P = \{1, \ldots, n\}$ of players
- a finite tree in which:
  - each edge is labelled with an action
  - each non-terminal node (called a decision node) is labelled by the player whose turn it is to choose an action
  - each terminal node is labelled with a tuple of utility values, one for each player
- a partition of the decision nodes into information sets

Extensive form games, cont.

Some constraints:

- a node cannot have two outgoing edges with the same label
- nodes in the same information sets must have the same player and same possible actions

Perfect-information extensive form games:
every information set contains a single node

Imperfect-information extensive form games:
some information sets contain multiple nodes

Example: Bach or Stravinsky

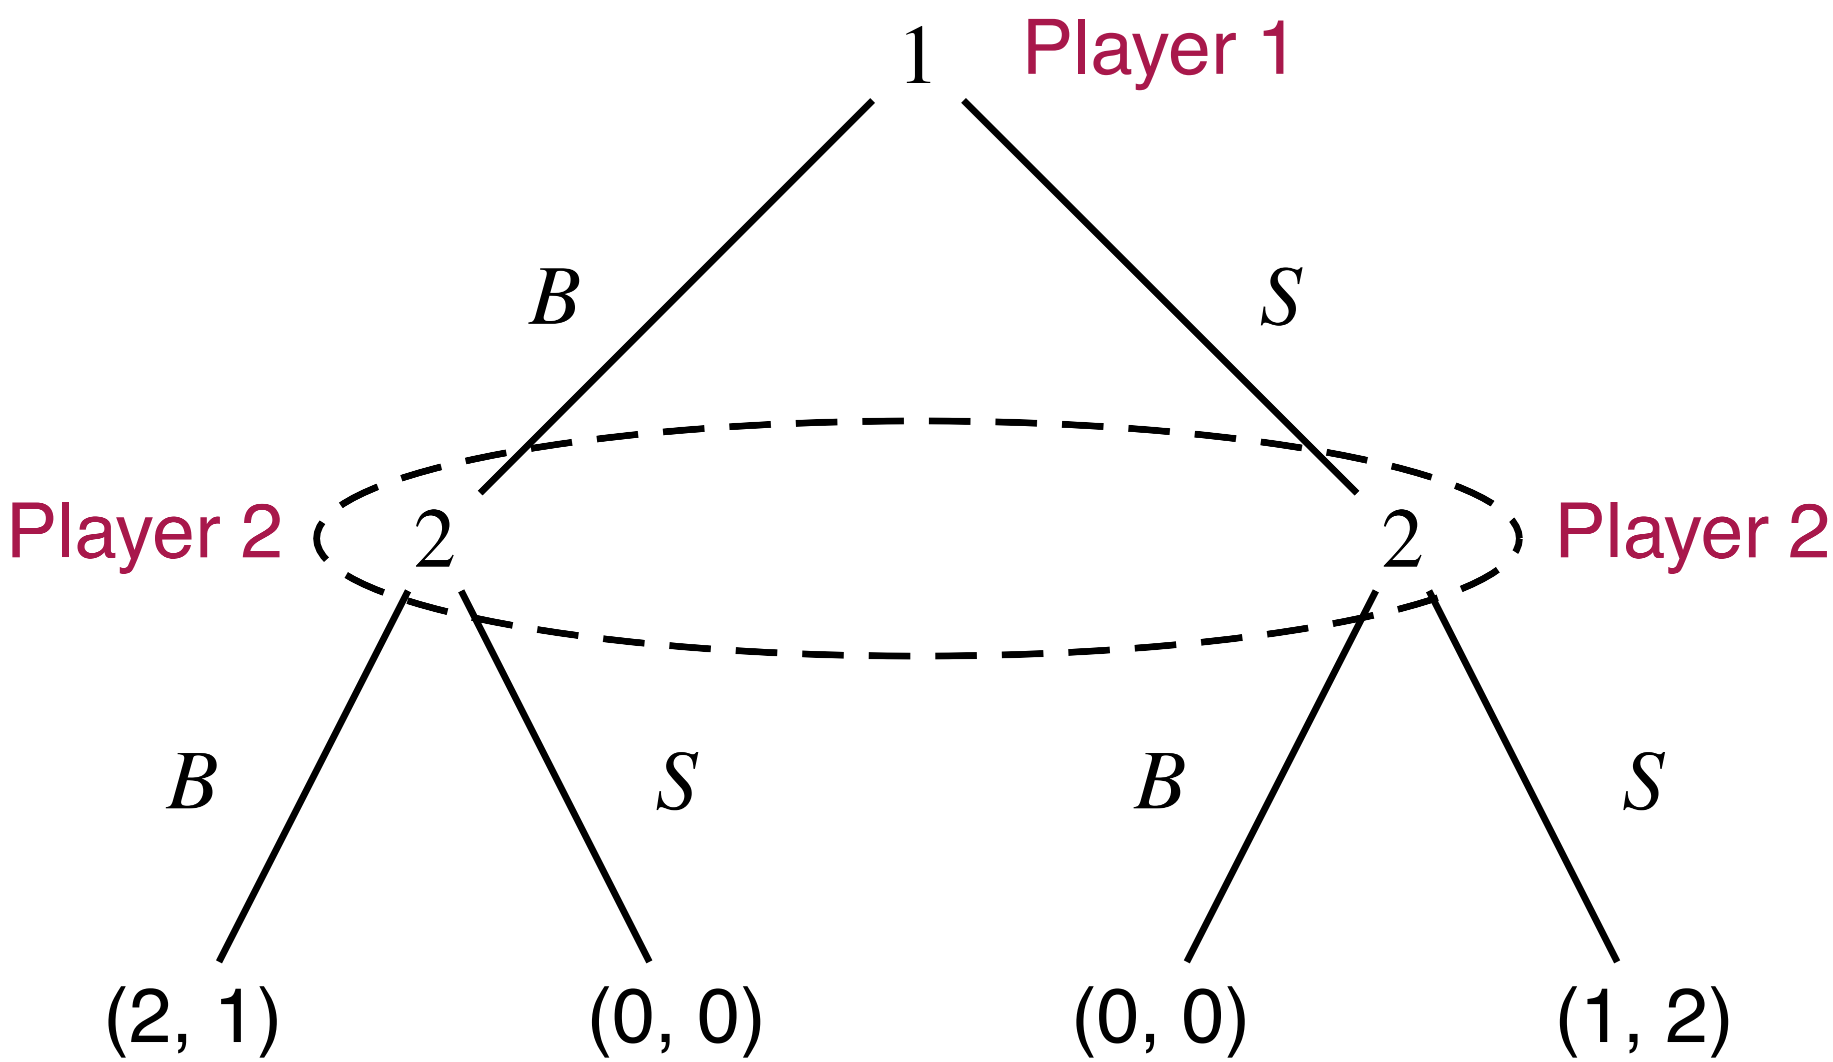

Information sets: $\{\lambda\}$, $\{B, S\}$

Since Player 2 ignores Player 1's choice, same game as before.

Example: Bach or Stravinsky

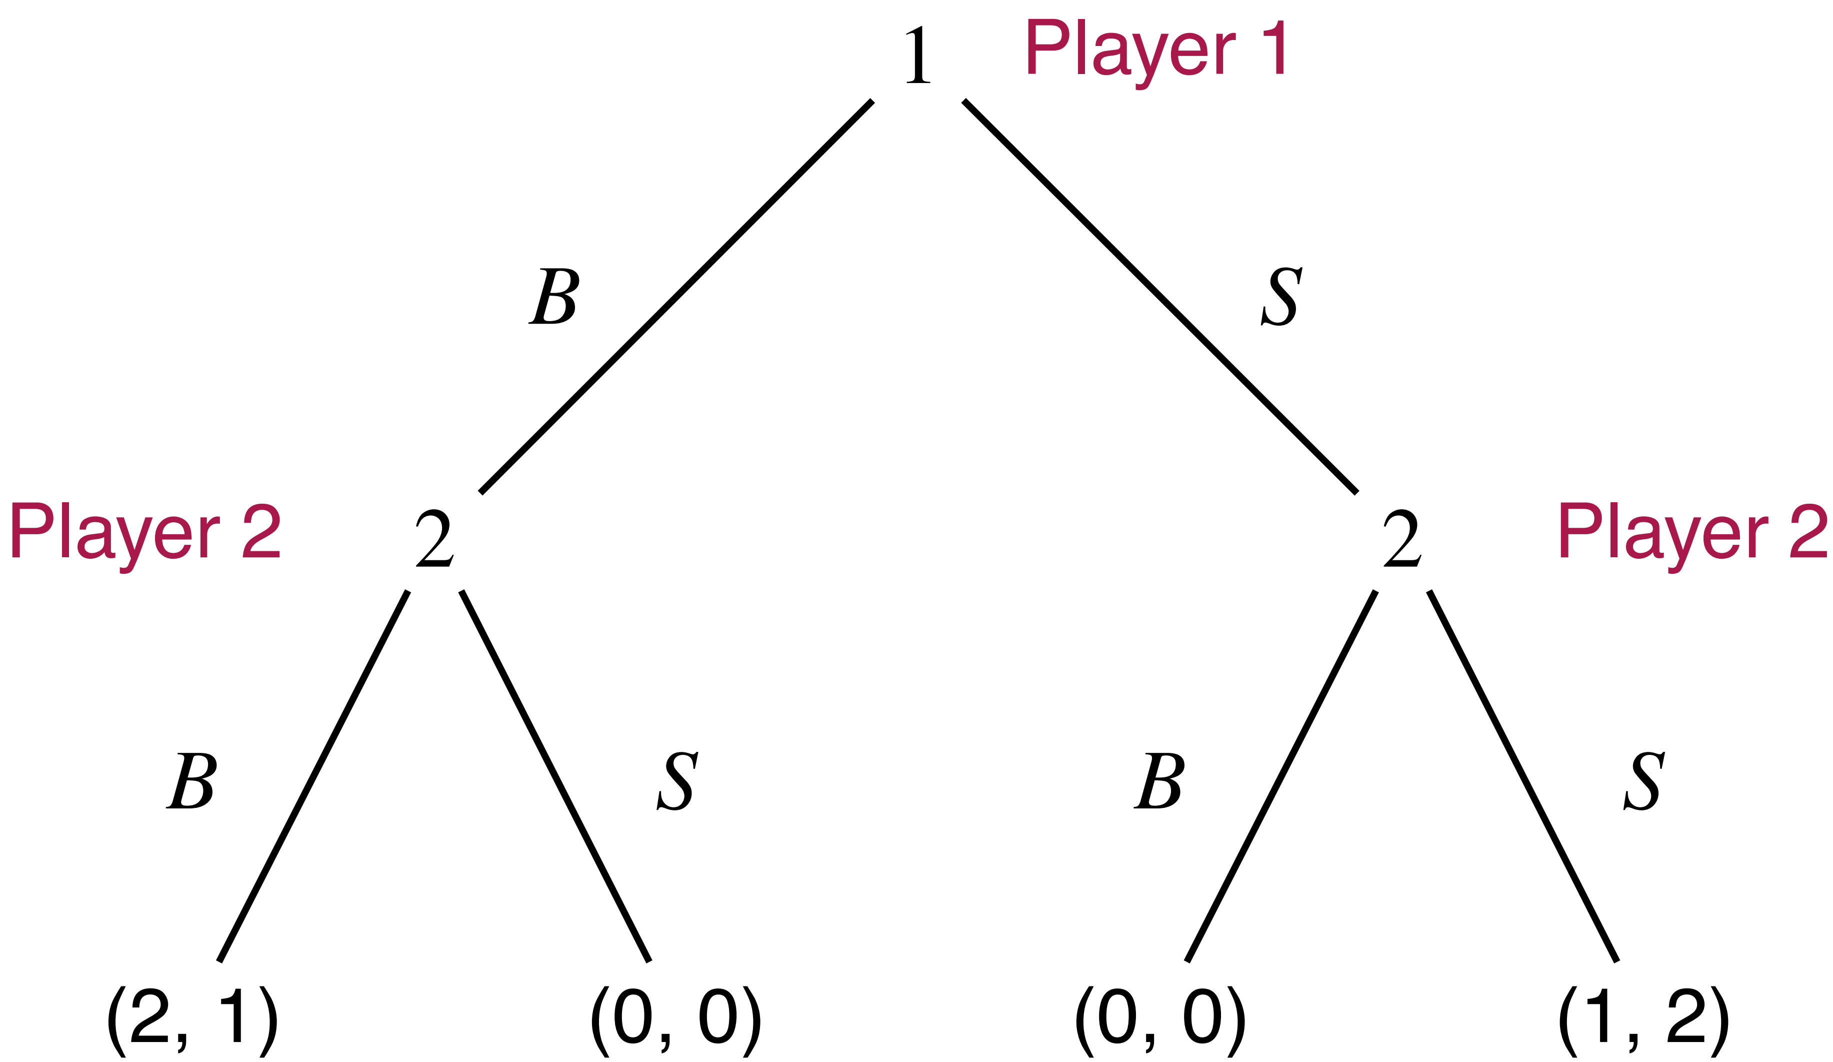

Information sets: $\{\lambda\}$, $\{B\}$, $\{S\}$

Now Player 2 knows what Player 1 has selected, and she can use this information to make her choice.

Pure strategies in extensive form games

A pure strategy tells us what choice a player will make in any possible state of the game.

More formally, a pure strategy for player i is a function mapping each information set of i to one of the actions which is possible from the information set.

For perfect-information games, this means the function selects an action at each of player i's nodes.

Important: the strategy definition requires a decision at every information set, even if it is not possible to reach that information set given earlier moves.

Pure strategies in extensive form games

In the perfect-information version of Bach or Stravinsky:

2 pure strategies for player 1: B and S

4 pure strategies for player 2:

- play B at choice points, written as $B(B), B(S)$
- play S at choice points, written as $S(B), S(S)$
- play B following B and S following S, written as $B(B), S(S)$
- play S after B and B after S, written as $S(B), B(S)$

Pure strategies in extensive form games

In the imperfect-information version of Bach or Stravinsky:

Same 2 pure strategies for player 1: B and S

Only 2 pure strategies for player 2:

- play B at both nodes, written as $B(\{B, S\})$
- play S at both nodes, written as $S(\{B, S\})$

From extensive form to strategic form

Every extensive form game can be translated into a strategic game as follows:

- the set of players is the same in both games
- the action set A_i of player i is the set of pure strategies of player i in the extensive form game
- the utility of an action profile for player i is simply the utility value of i at the terminal node which occurs when all players follow their pure strategies in the action profile

Note: strategic game may be exponentially larger

From extensive form to strategic form

For the perfect-information version of Bach or Stravinsky, we get:

	B(B),B(S)	S(B),S(S)	B(B),S(S)	S(B),B(S)
B				
S				

From extensive form to strategic form

For the perfect-information version of Bach or Stravinsky, we get:

	B(B),B(S)	S(B),S(S)	B(B),S(S)	S(B),B(S)
B	2 , 1	0 , 0	2 , 1	0 , 0
S	0 , 0	1 , 2	1 , 2	0 , 0

From strategic form to extensive form

Every strategic game can be translated into an imperfect-information extensive form game as follows:

- the set of players is the same in both games, $P = \{1, \dots, n\}$
- at root node, player 1's turn, create one outgoing edge for each action in A_1
- at second level, all nodes belong to player 2 and are part of a single information set, outgoing edges correspond to the actions in A_2
- repeat procedure for players $3, \dots, n$
- at terminal node, if a is the action profile obtained by taking the actions along the path from the root node, then we assign the utility vector $(u_1(a), \dots, u_n(a))$

From strategic form to extensive form

Bach or Stravinsky game yields:

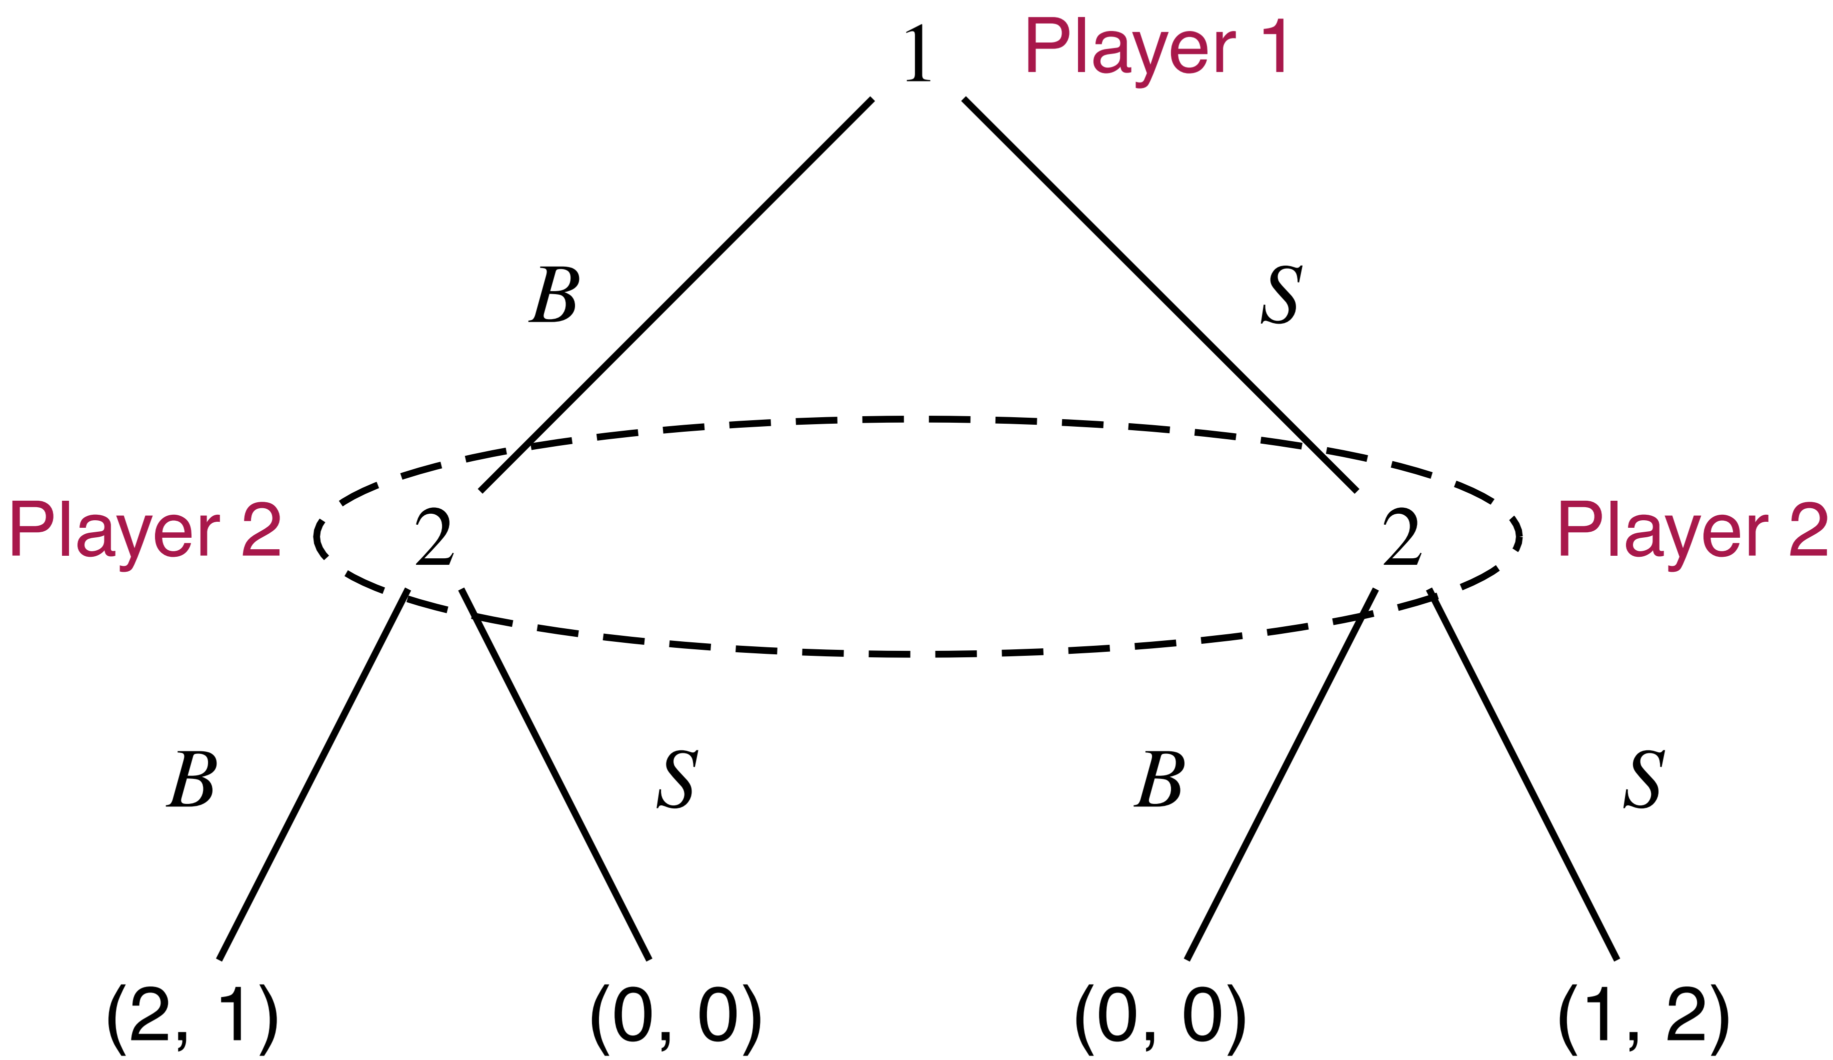

From strategic form to extensive form

which translates back to...

	B({B,S})	S({B,S})
B	2 , 1	0 , 0
S	0 , 0	1 , 2

Same as original game, just different names for actions.

Nash equilibrium in extensive form games

Definition of Nash equilibrium also can be used for extensive form games.

For imperfect-information extensive form games, a pure strategy Nash equilibrium need not exist. **Why** ?

For perfect-information extensive form games, there is always at least one pure strategy Nash equilibrium.

Two methods for finding a PNE:

1. Translate into equivalent strategic form game
2. Backward induction

First method: use strategic games

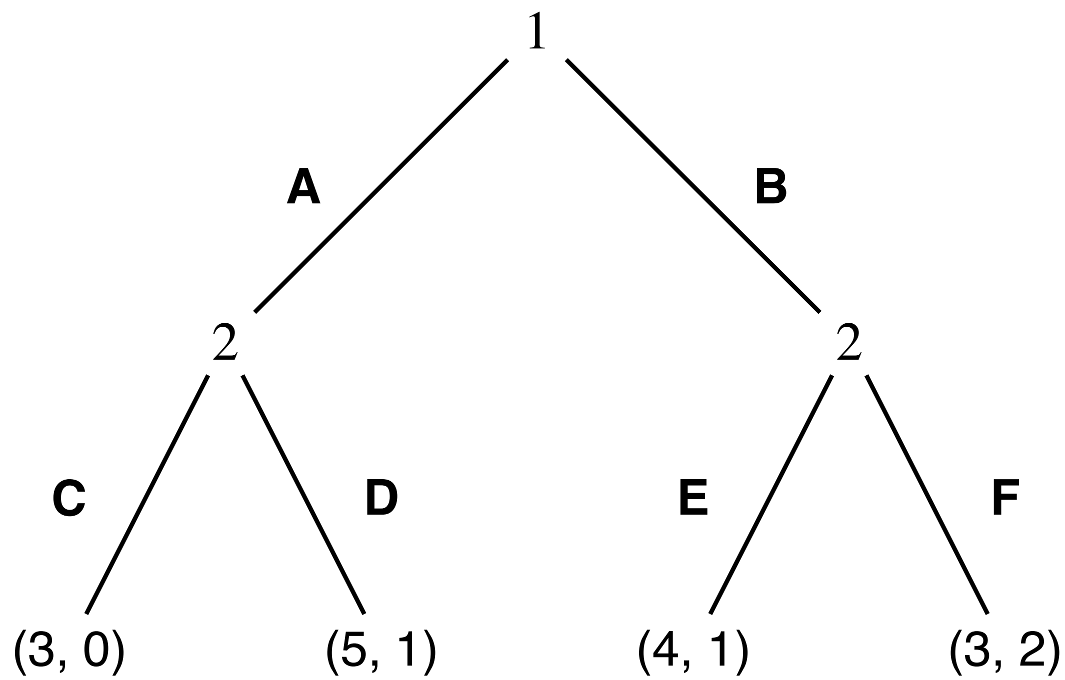

First method: use strategic games

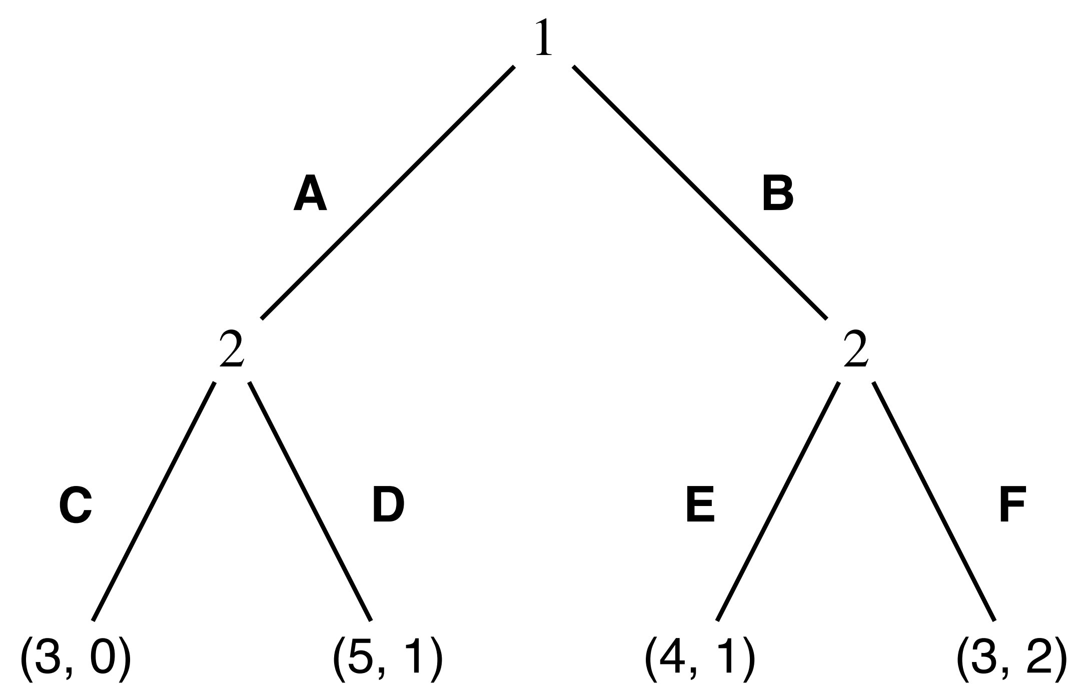

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

First method: use strategic games

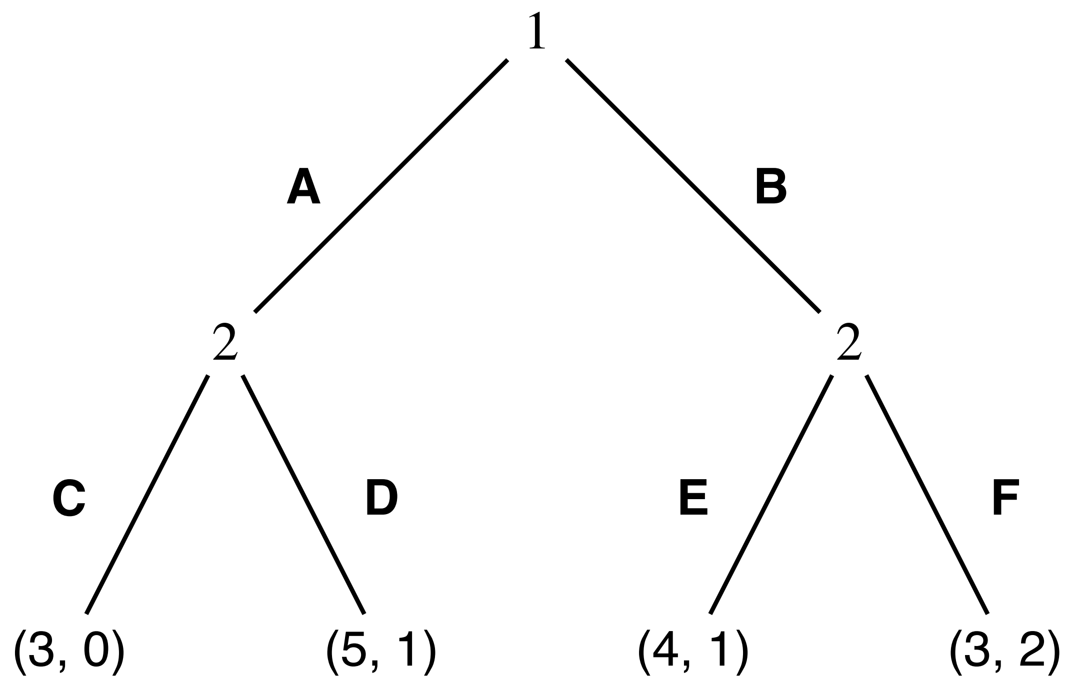

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

Second method: backward induction

Basic idea: start at the terminal nodes, and select always the actions leading to best payoff

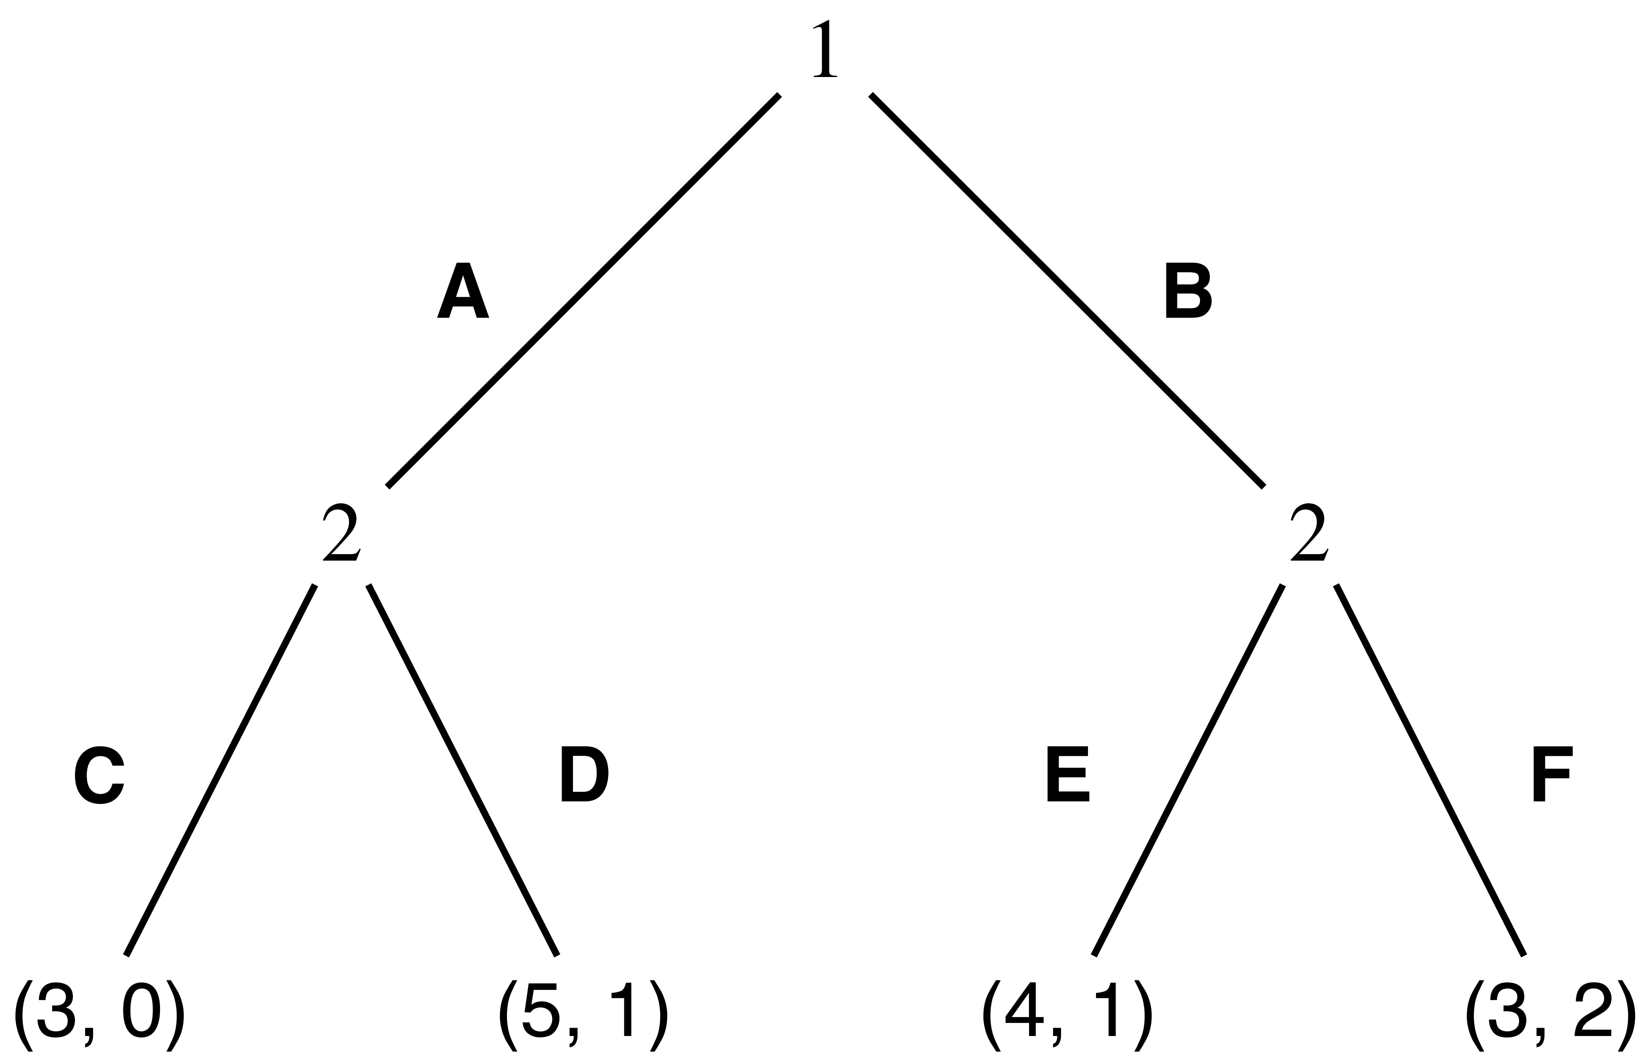

Second method: backward induction

Basic idea: start at the terminal nodes, and select always the actions leading to best payoff

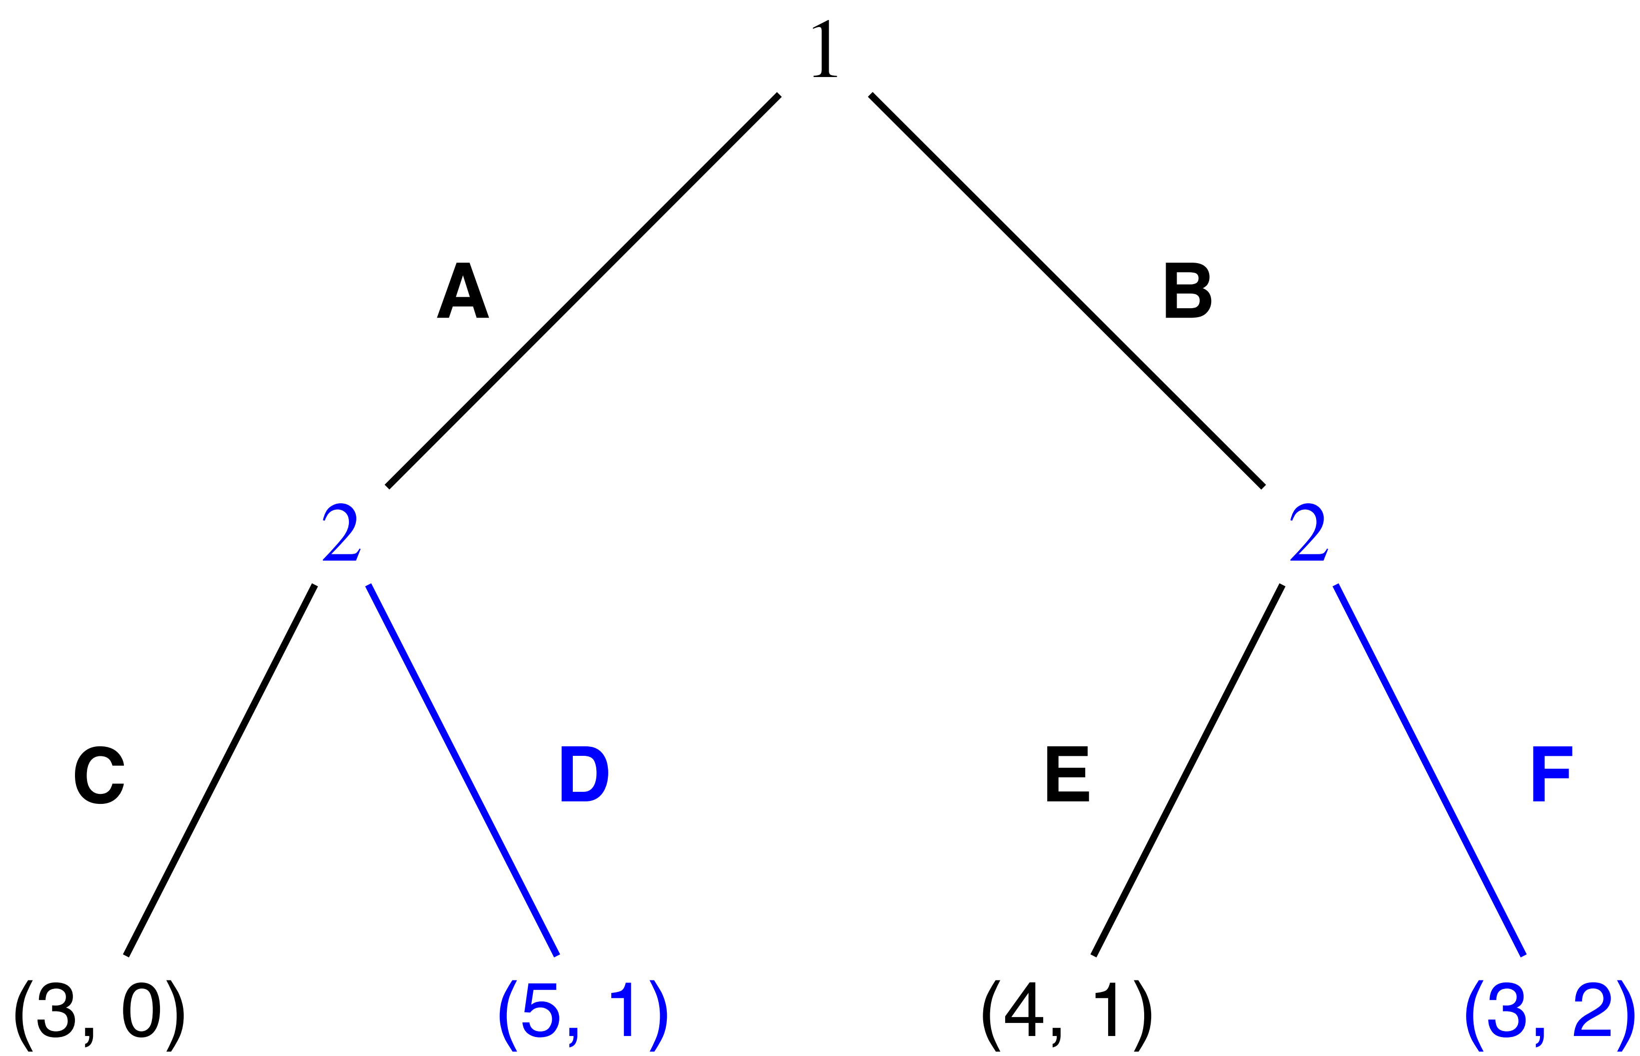

Second method: backward induction

Basic idea: start at the terminal nodes, and select always the actions leading to best payoff

Player 1: A

Player 2: D(A), F(B)

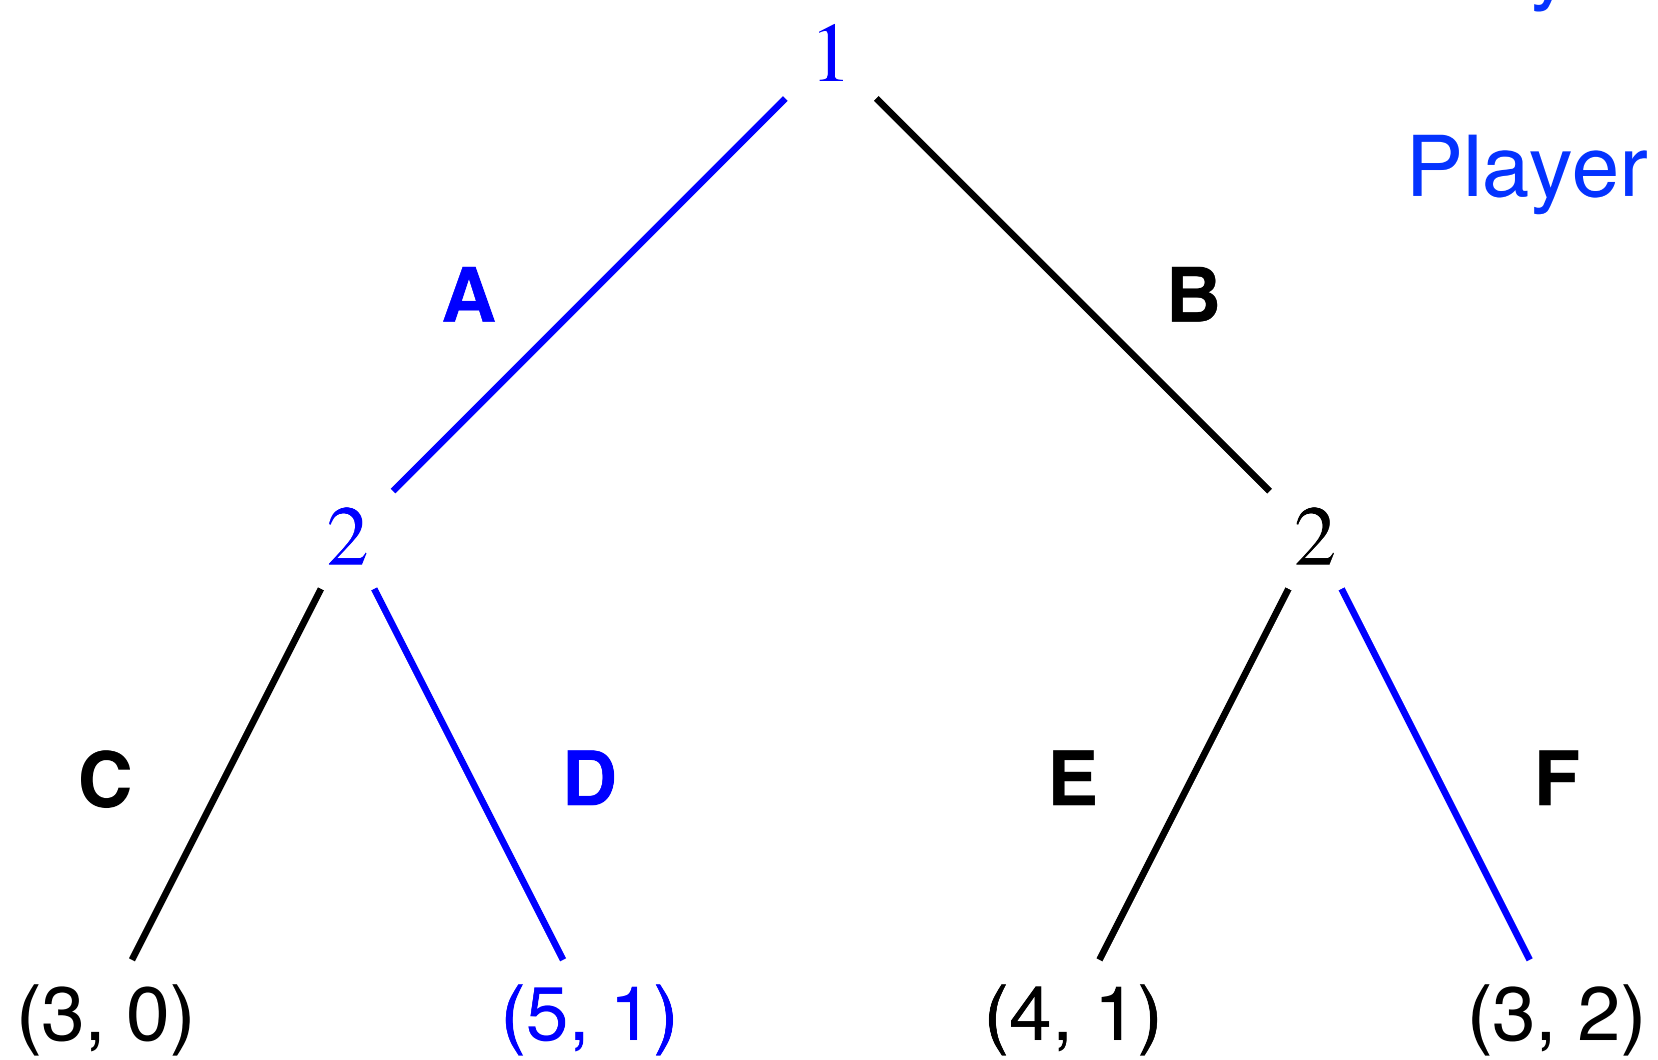

Subgame perfect equilibrium

Nash equilibria well-defined for extensive form games, but sometimes gives unnatural solutions.

Subgame perfect equilibria is a refinement of Nash equilibria which addresses this problem.

A subgame of a perfect-information extensive form game is a restriction of a game to a node and its descendants.

A subgame perfect equilibrium (SPE) of a game G is a strategy profile s such that for any subgame G' of G, the restriction of s to G' is a Nash equilibrium.

Note: backward induction always gives SPE.

Subgames

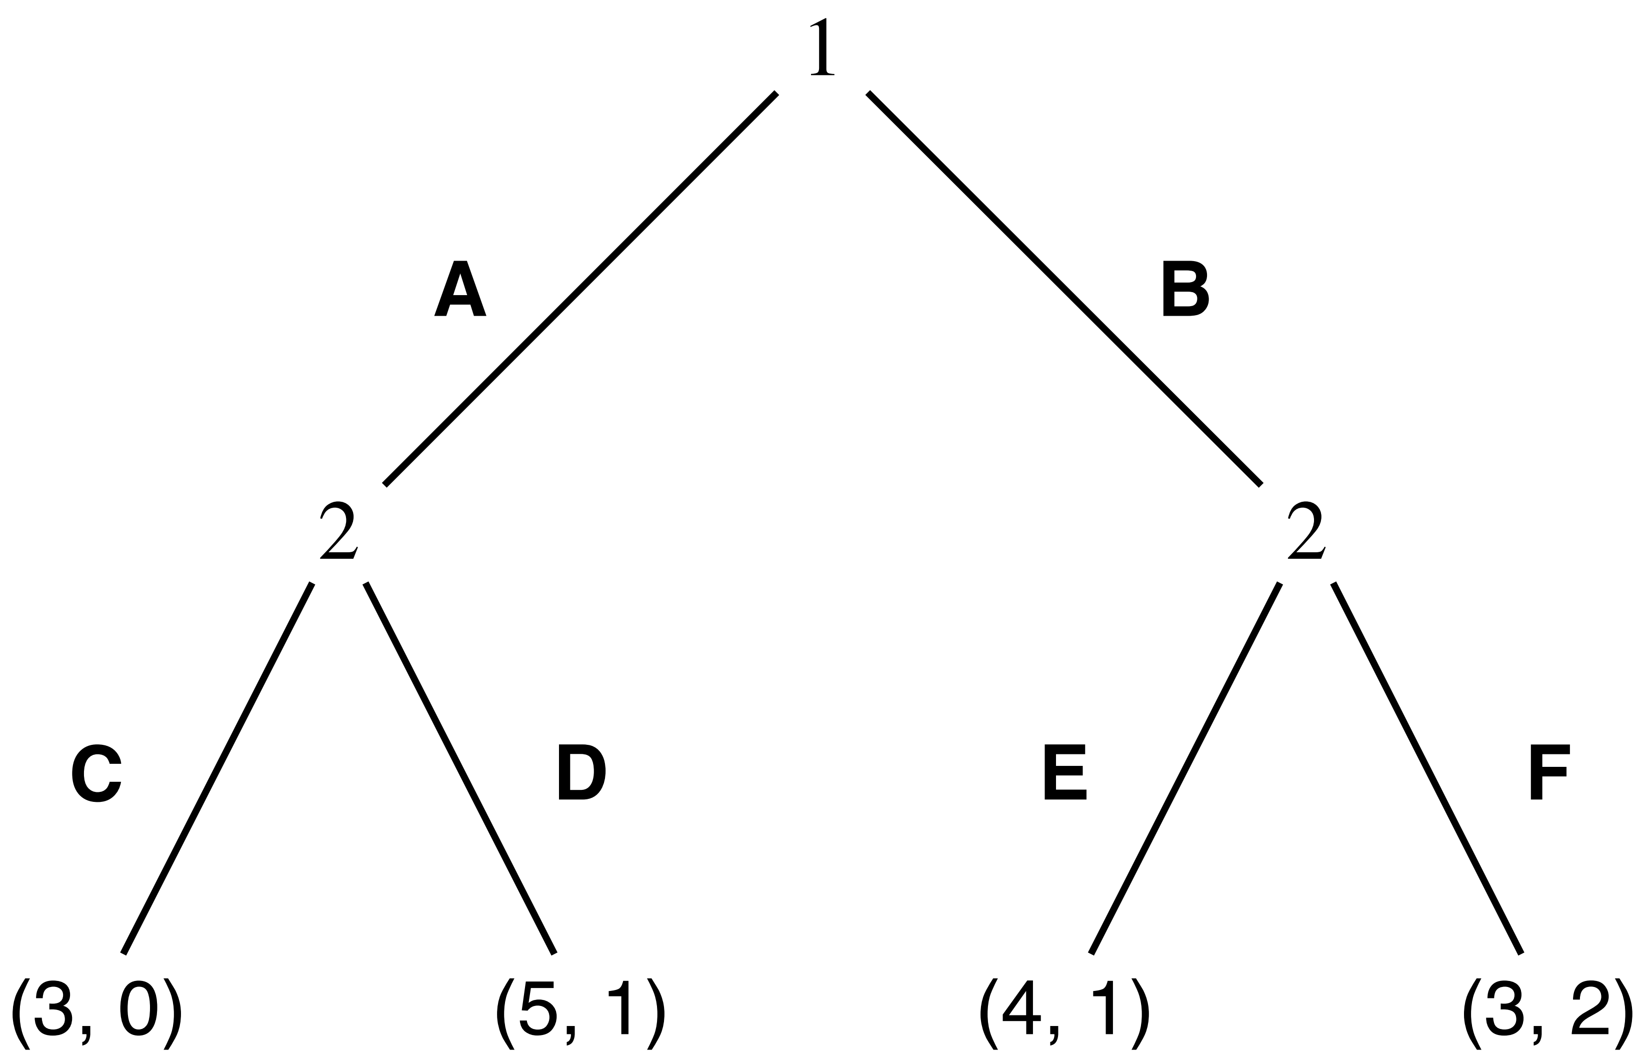

3 subgames: whole game, subtree of A, subtree of B

Subgames

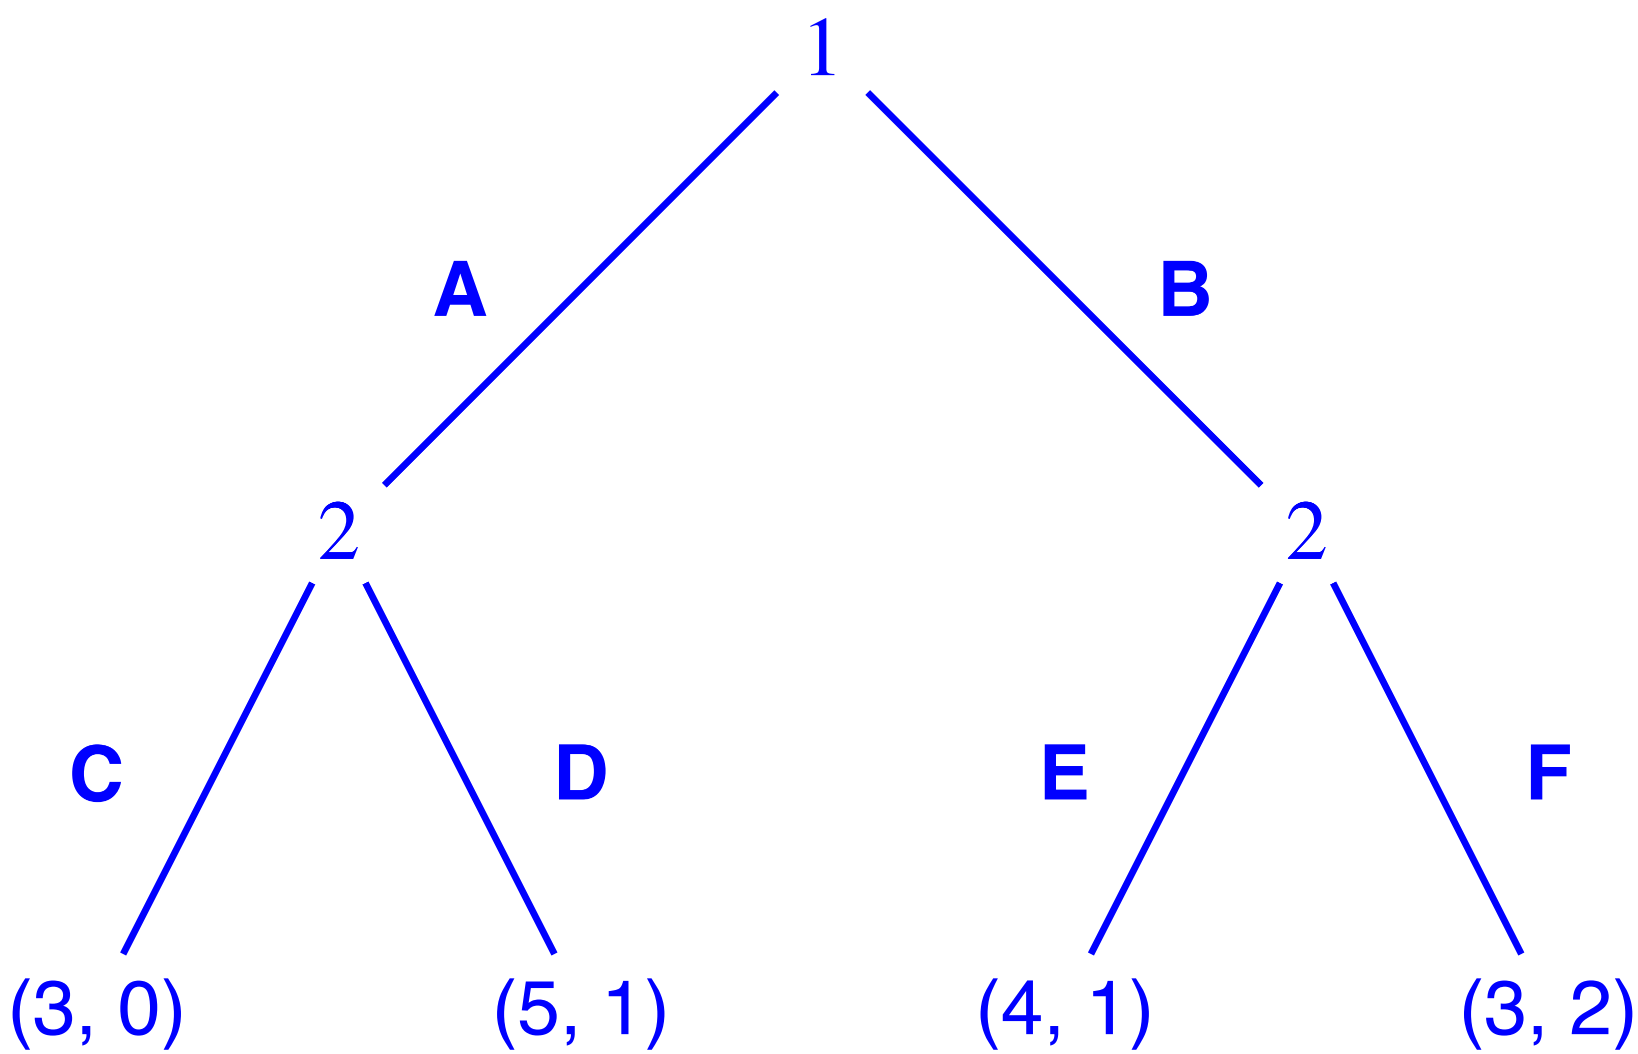

3 subgames: whole game, subtree of A, subtree of B

Subgames

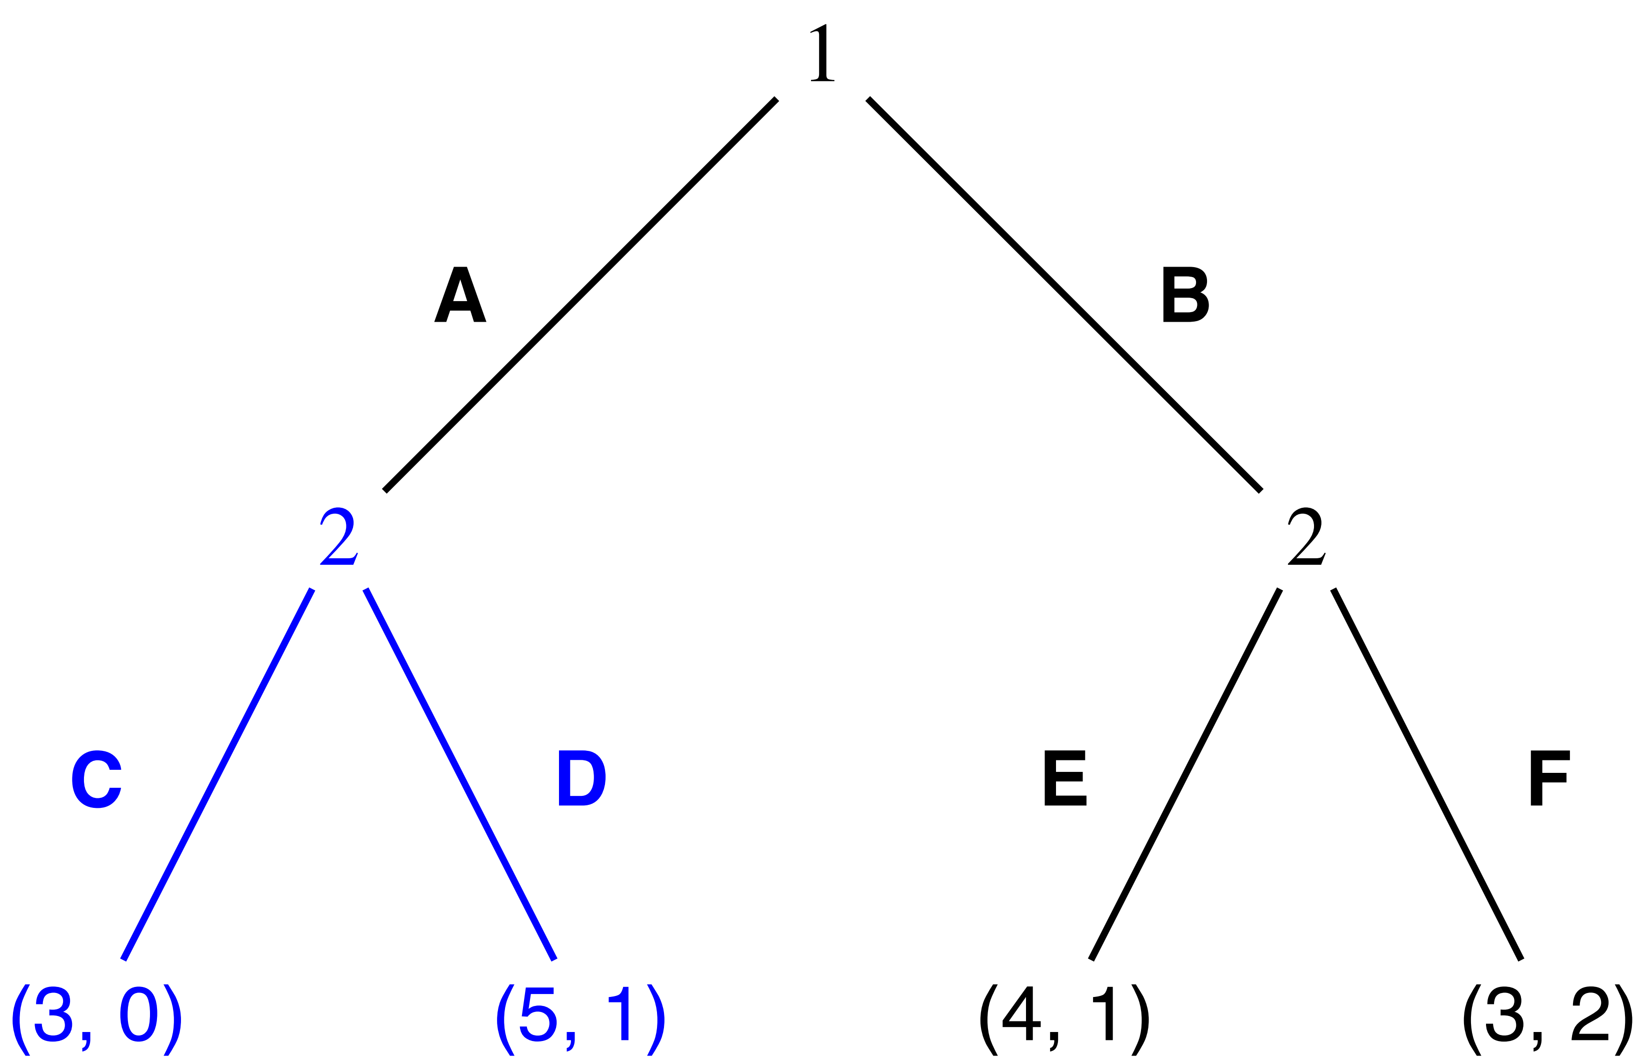

3 subgames: whole game, subtree of A, subtree of B

Subgames

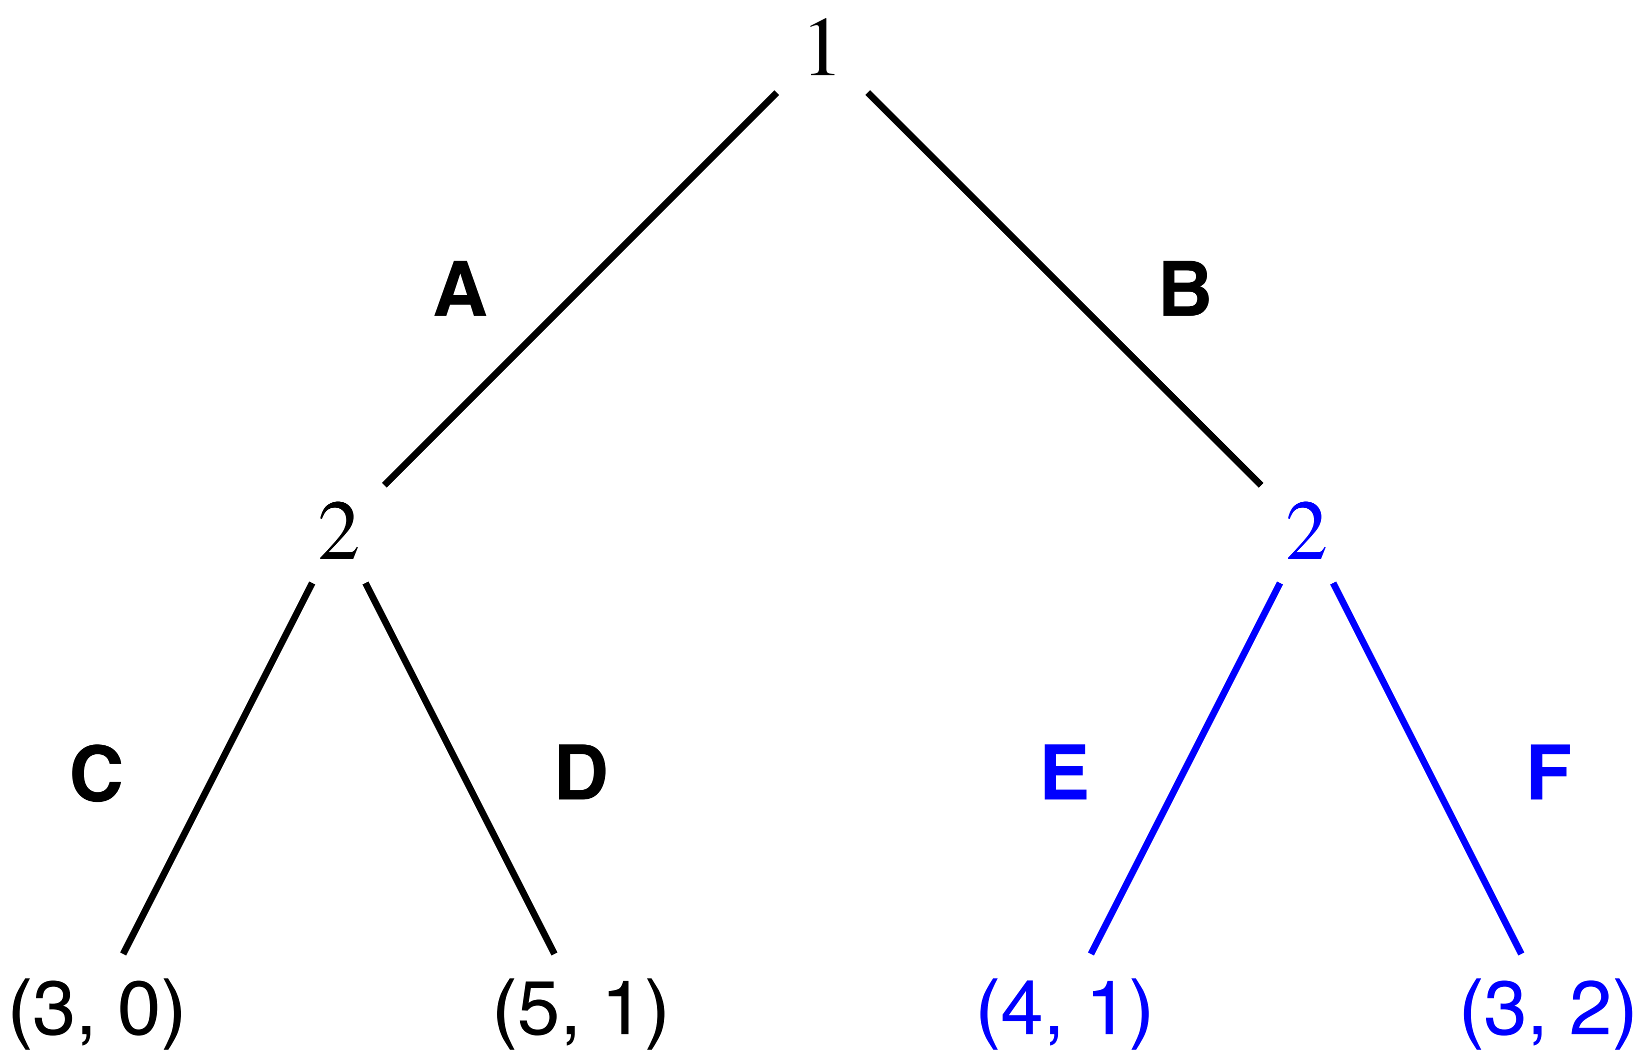

3 subgames: whole game, subtree of A, subtree of B

Subgame perfect equilibrium

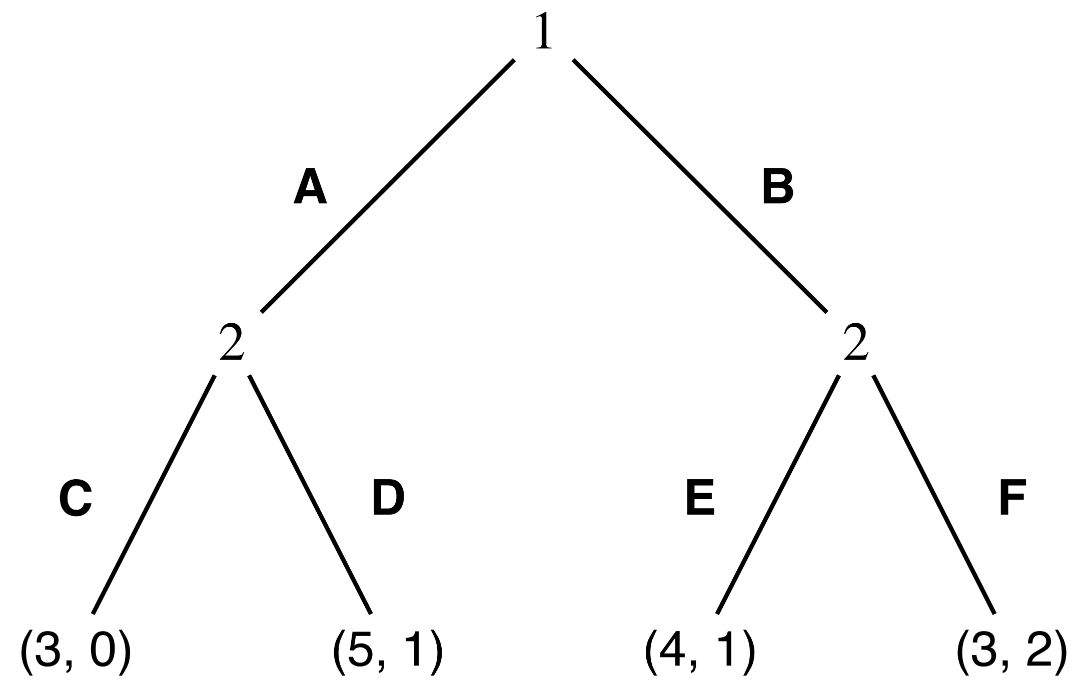

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

3 PNE. How many are subgame perfect equilibria?

Subgame perfect equilibrium

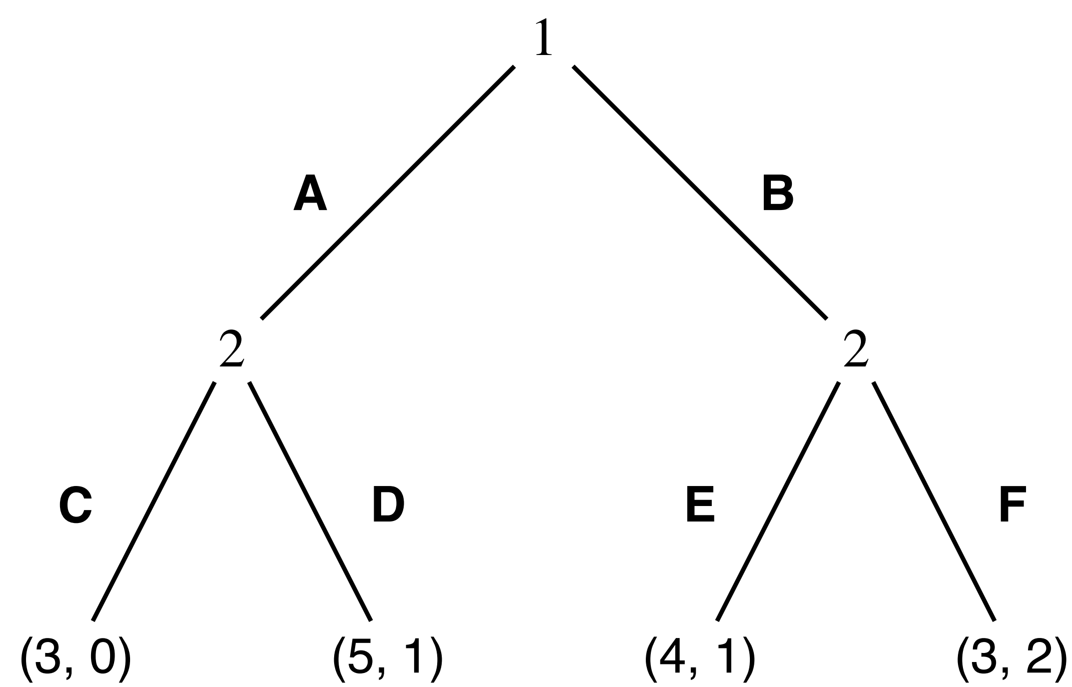

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

Not SPE since C not NE in subgame rooted at A.

Subgame perfect equilibrium

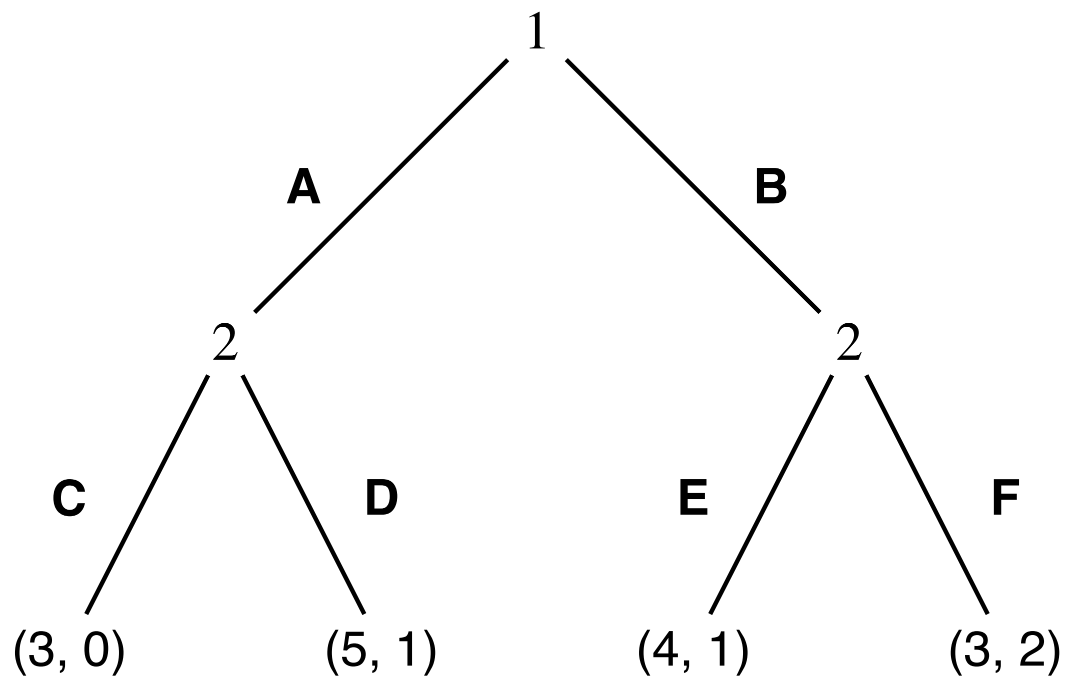

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

Not SPE since E not NE in subgame rooted at B.

Subgame perfect equilibrium

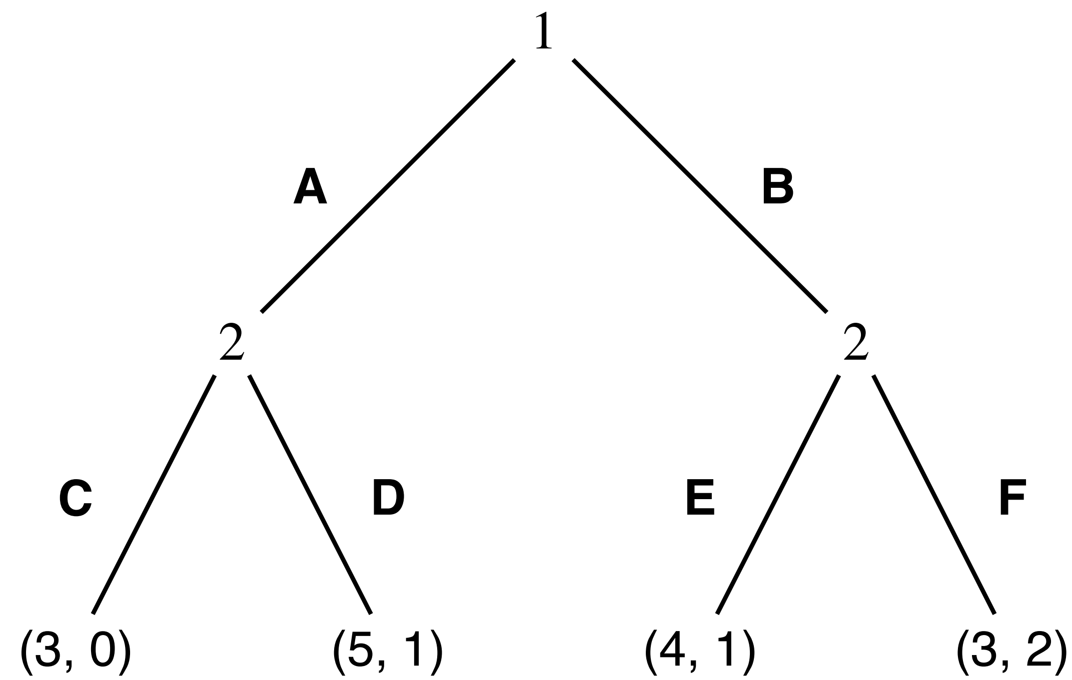

	C(A), E(B)	C(A), F(B)	D(A), E(B)	D(A), F(B)
A	3 , 0	3 , 0	5 , 1	5 , 1
B	4 , 1	3 , 2	4 , 1	3 , 2

Only 1 SPE, corresponding to intuitive solution.

Computing subgame perfect equilibrium

We have shown that backward induction returns a subgame perfect equilibrium.

<u>Good news</u>: backward induction takes only polynomial time in the size of the extensive game.

<u>Bad news</u>: for many realistic games, the extensive form is simply too large to be represented.

<u>Example</u>: chess can be represented as a perfect-information extensive form game, but requires around 10^{150} nodes !!

In general, no way to avoid looking at all nodes in the tree, but for zero-sum games, some improvements exist...

Minmax algorithm and backward induction

For two-player zero-sum games, only need to keep track of the utility for player 1, since player 2's utility is just the opposite.

Player 1 wants to maximize this value, and player 2 wants to minimize it.

This means that backward induction corresponds to taking the maximum at 1 nodes, and the minimum at 2 nodes.

Let's see this on an example...

Minmax algorithm

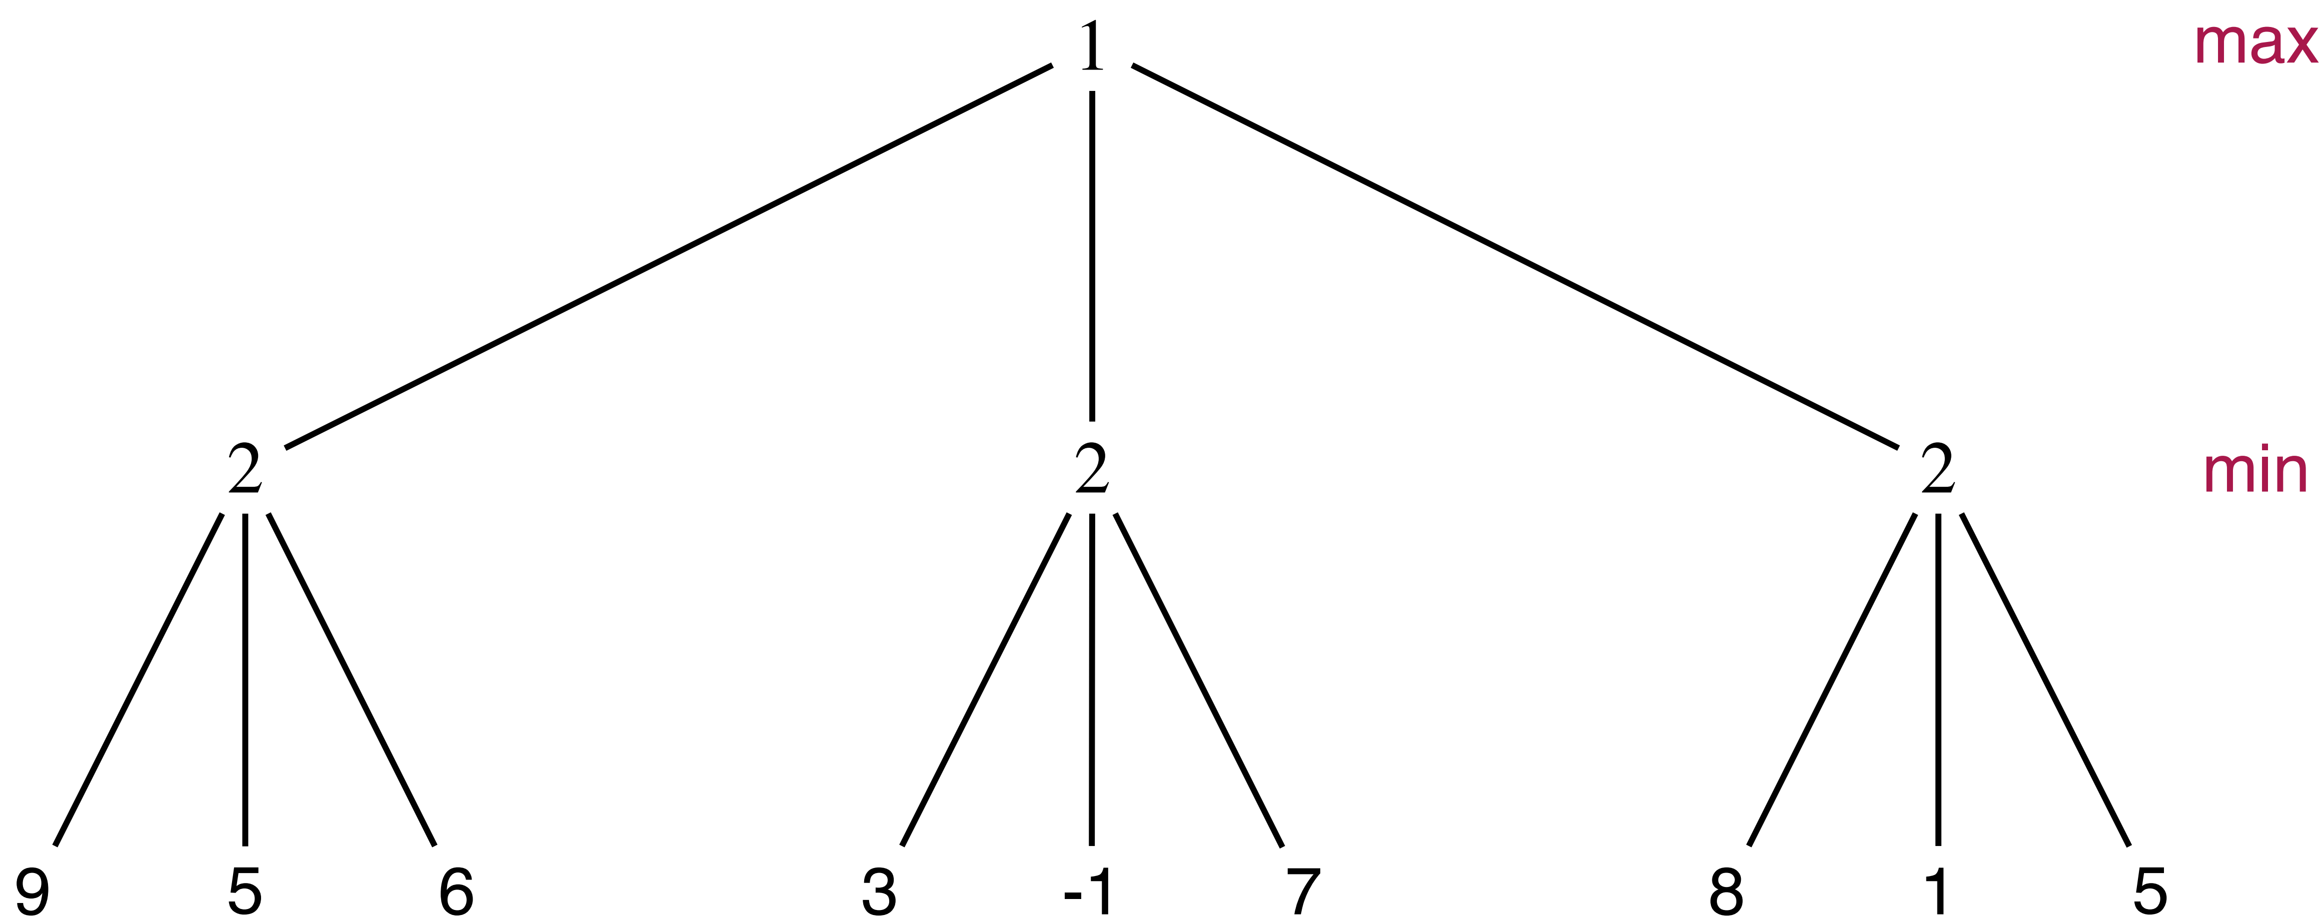

Minmax algorithm

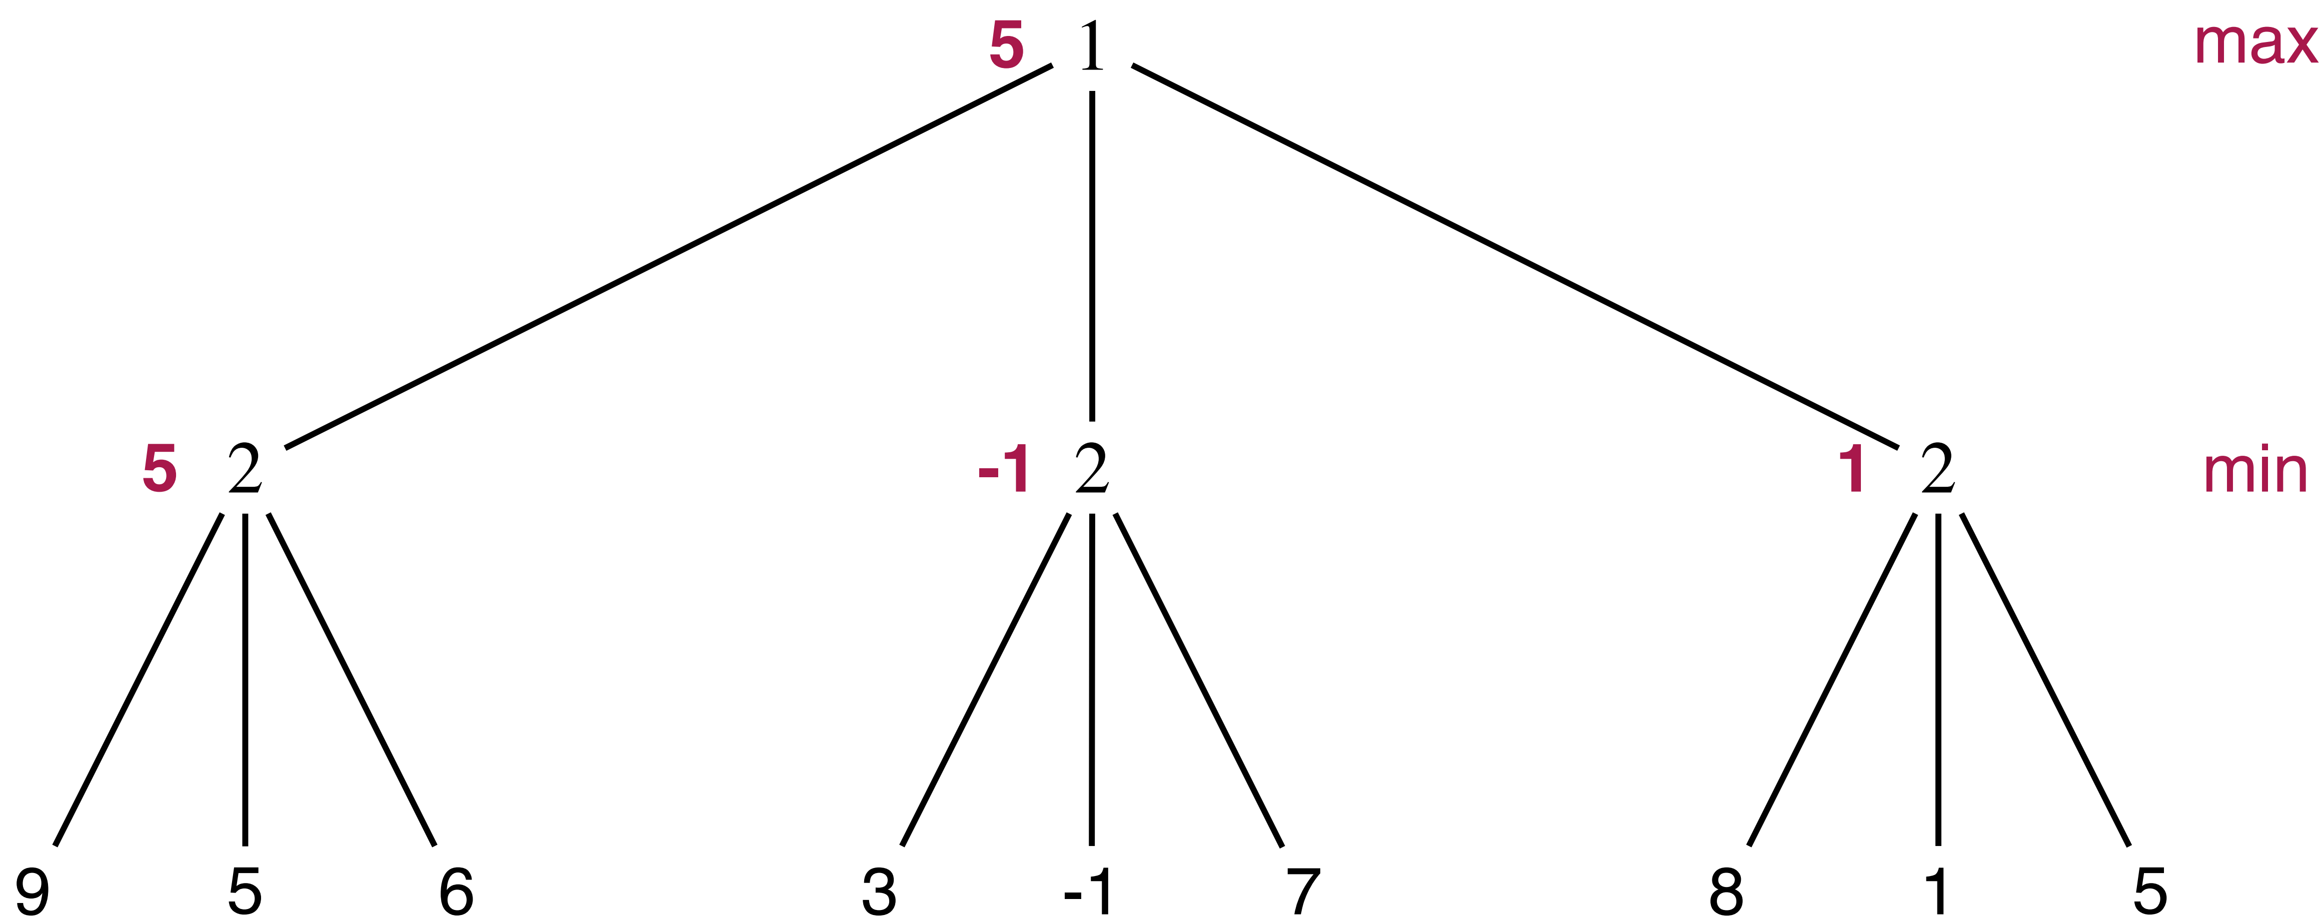

Alpha-beta pruning

Idea: use depth-first traversal, skip nodes whenever possible

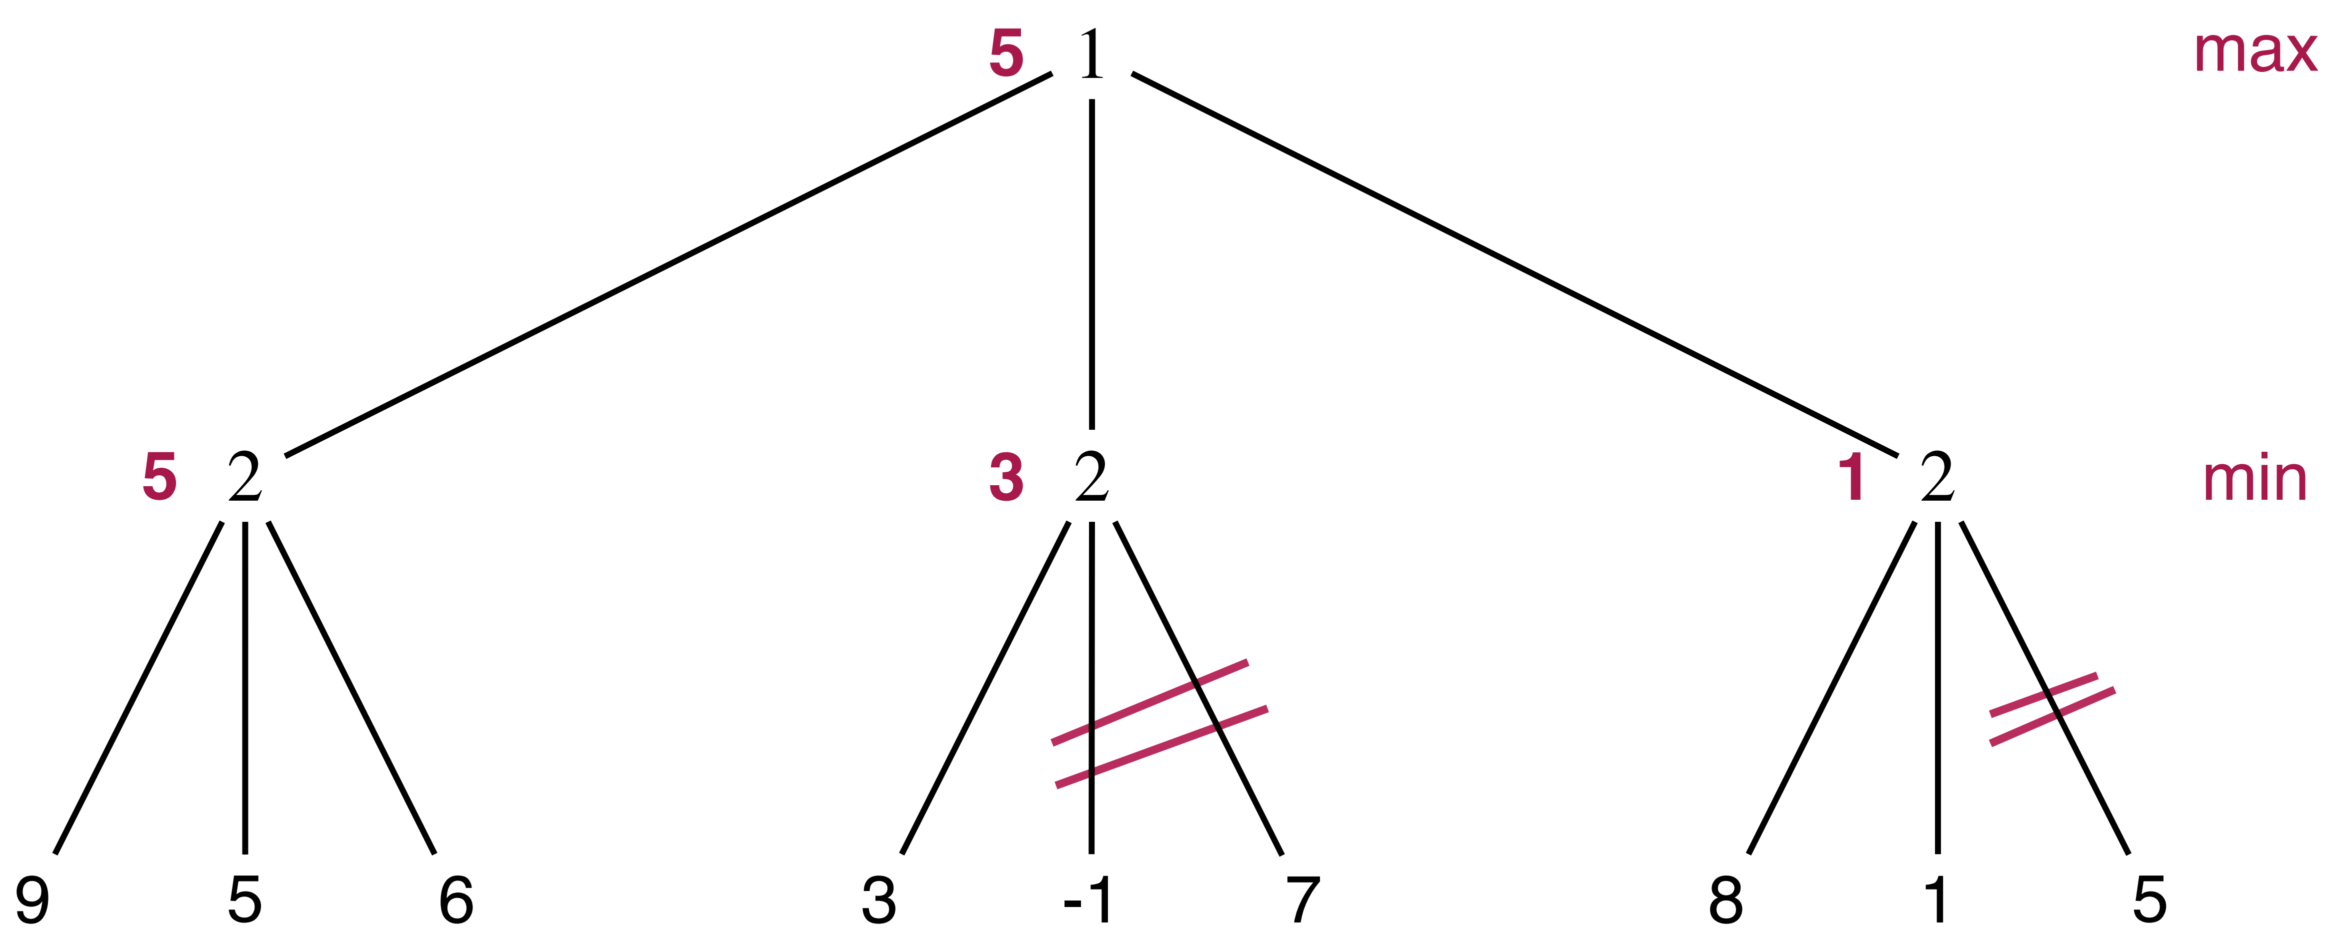

Once 3 is visited, know that the value of the node is at most 3, which is lower than 5. So player 1 will never choose this branch.

Once 1 is seen, know that the value of the node is at most 1, which is lower than 5. So can safely skip the final node.

Imperfect-information games

In imperfect-information extensive form games, there might be no pure strategy Nash equilibria.

In order to ensure the existence of a NE, we can introduce mixed strategies for extensive form games.

A mixed strategy for player i in an extensive form game is a probability distribution over the set of pure strategies of player i.

Can also define another type of probabilistic strategy where we associate a probability distribution over actions to each of the player’s information sets. Known as behavioral strategies.

Example

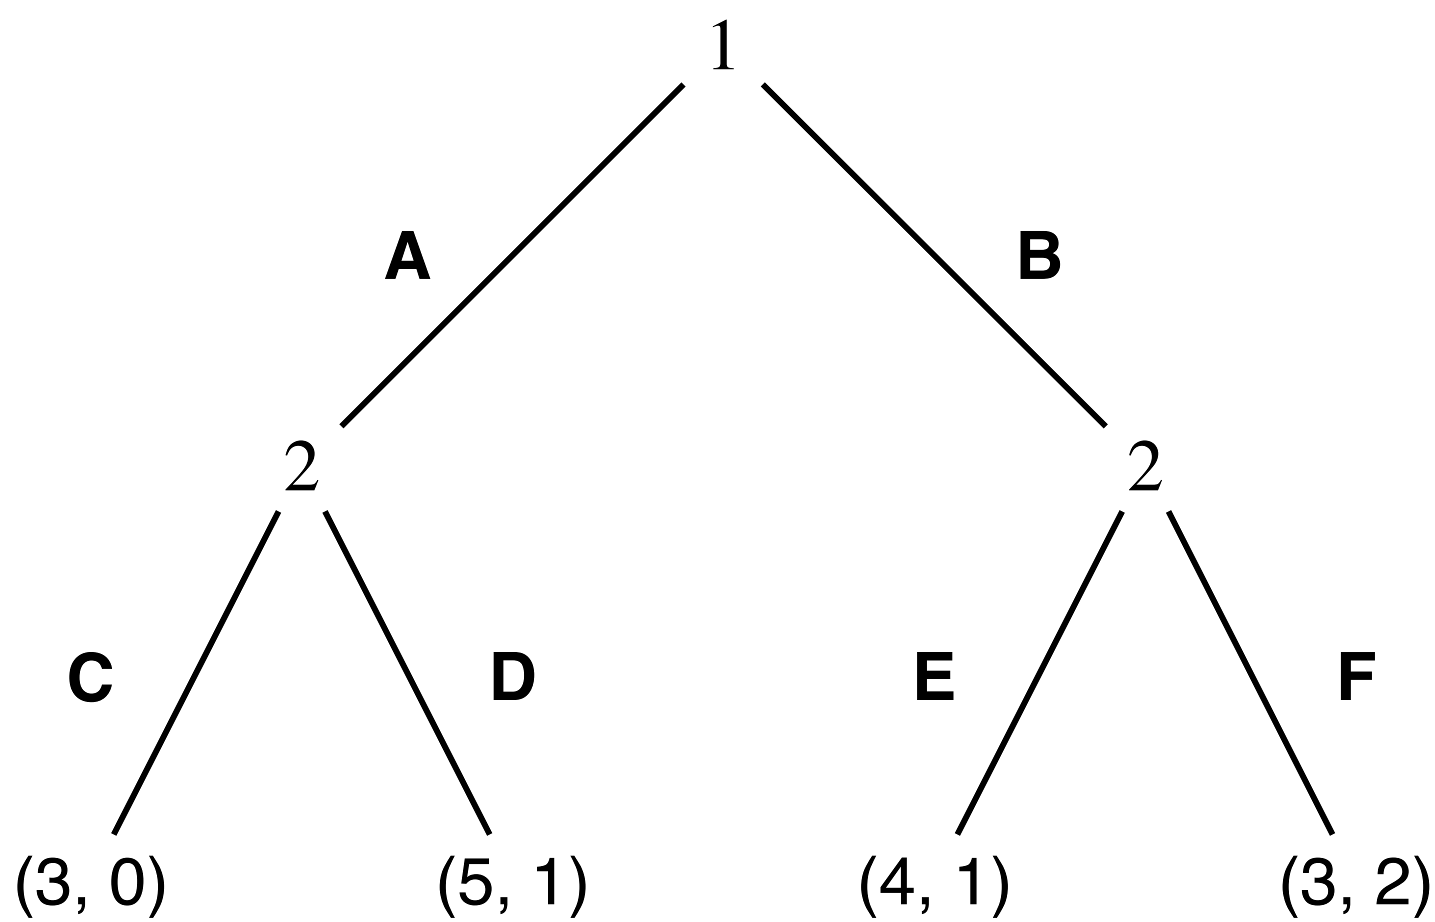

Mixed strategy for player 2: $\frac{1}{2}[C(A), E(B)]$ and $\frac{1}{2}[C(A), F(B)]$

Behavioral strategy for player 2: play $[\frac{1}{3}C, \frac{2}{3}D]$ at A and $[\frac{3}{5}E, \frac{2}{5}F]$ at B

Question: do we get the same equilibrium notions ?

Single-player imperfect-information game

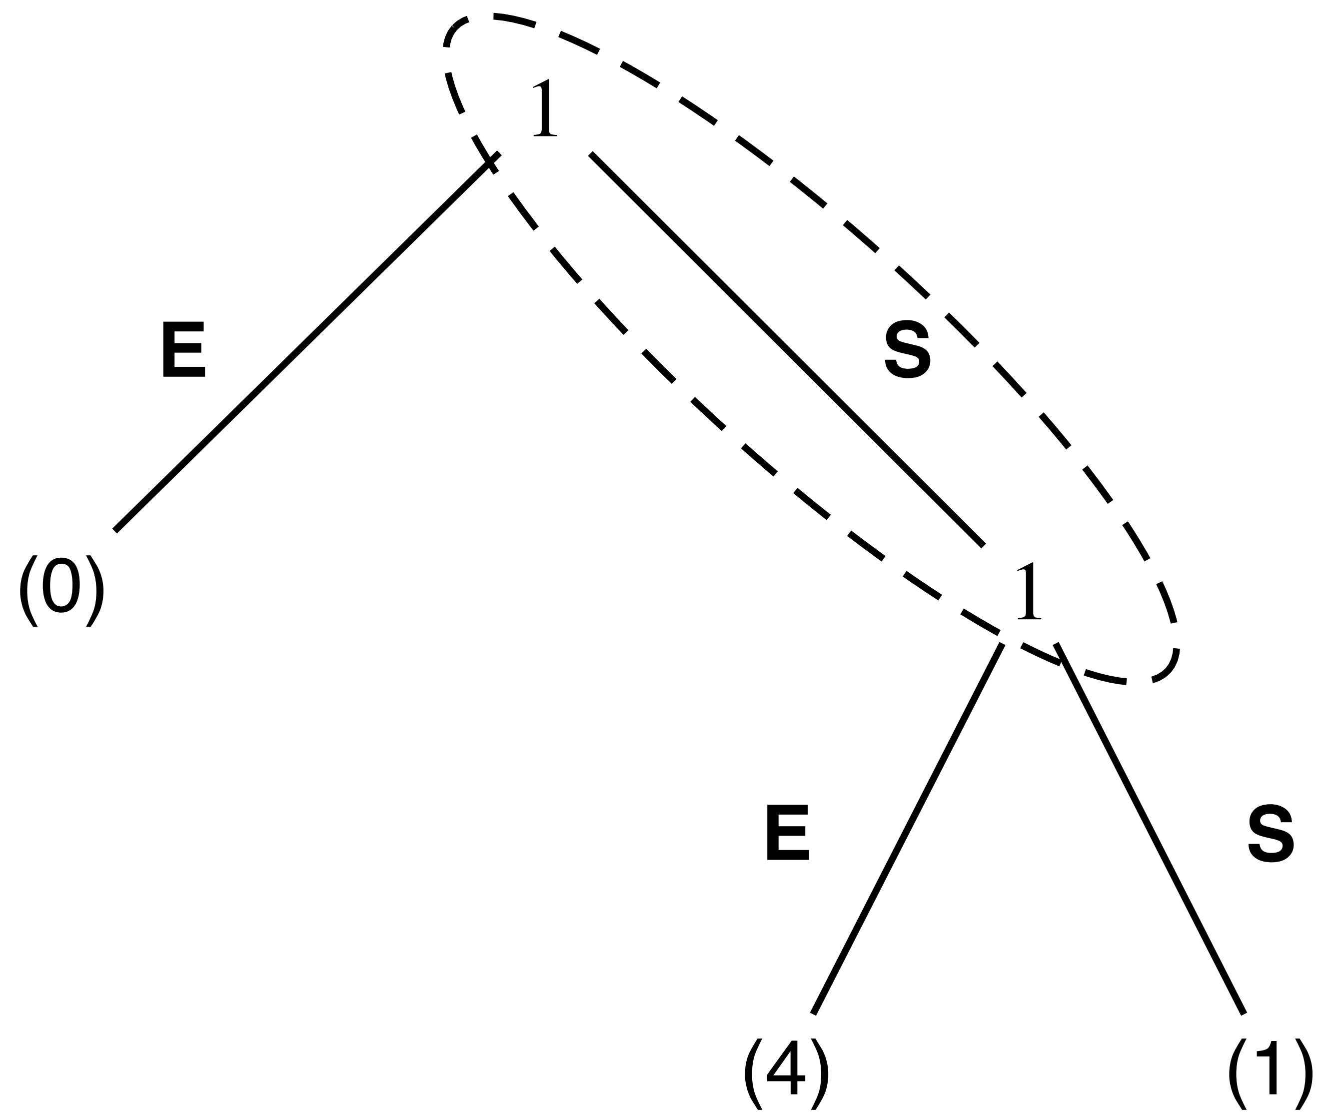

Single-player imperfect-information game

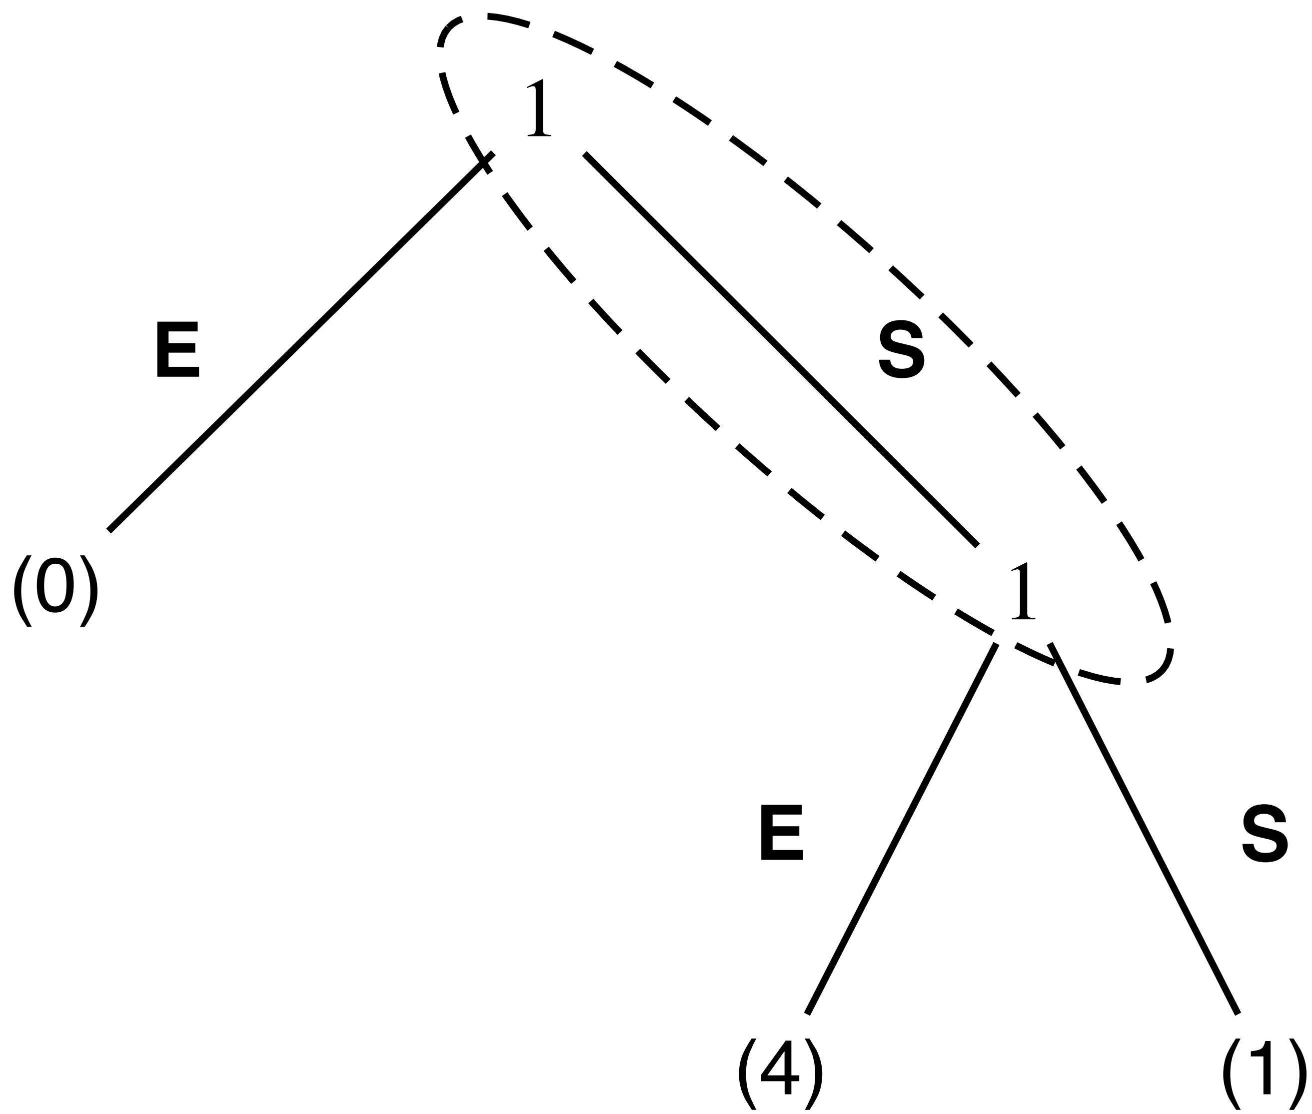

Pure strategies: E and S

Mixed strategies: play pure strategy E with probability p, S with probability $1 - p$

Behavioral strategies: at each 1-node, play E with probability p, S with probability $1 - p$

Single-player imperfect-information game

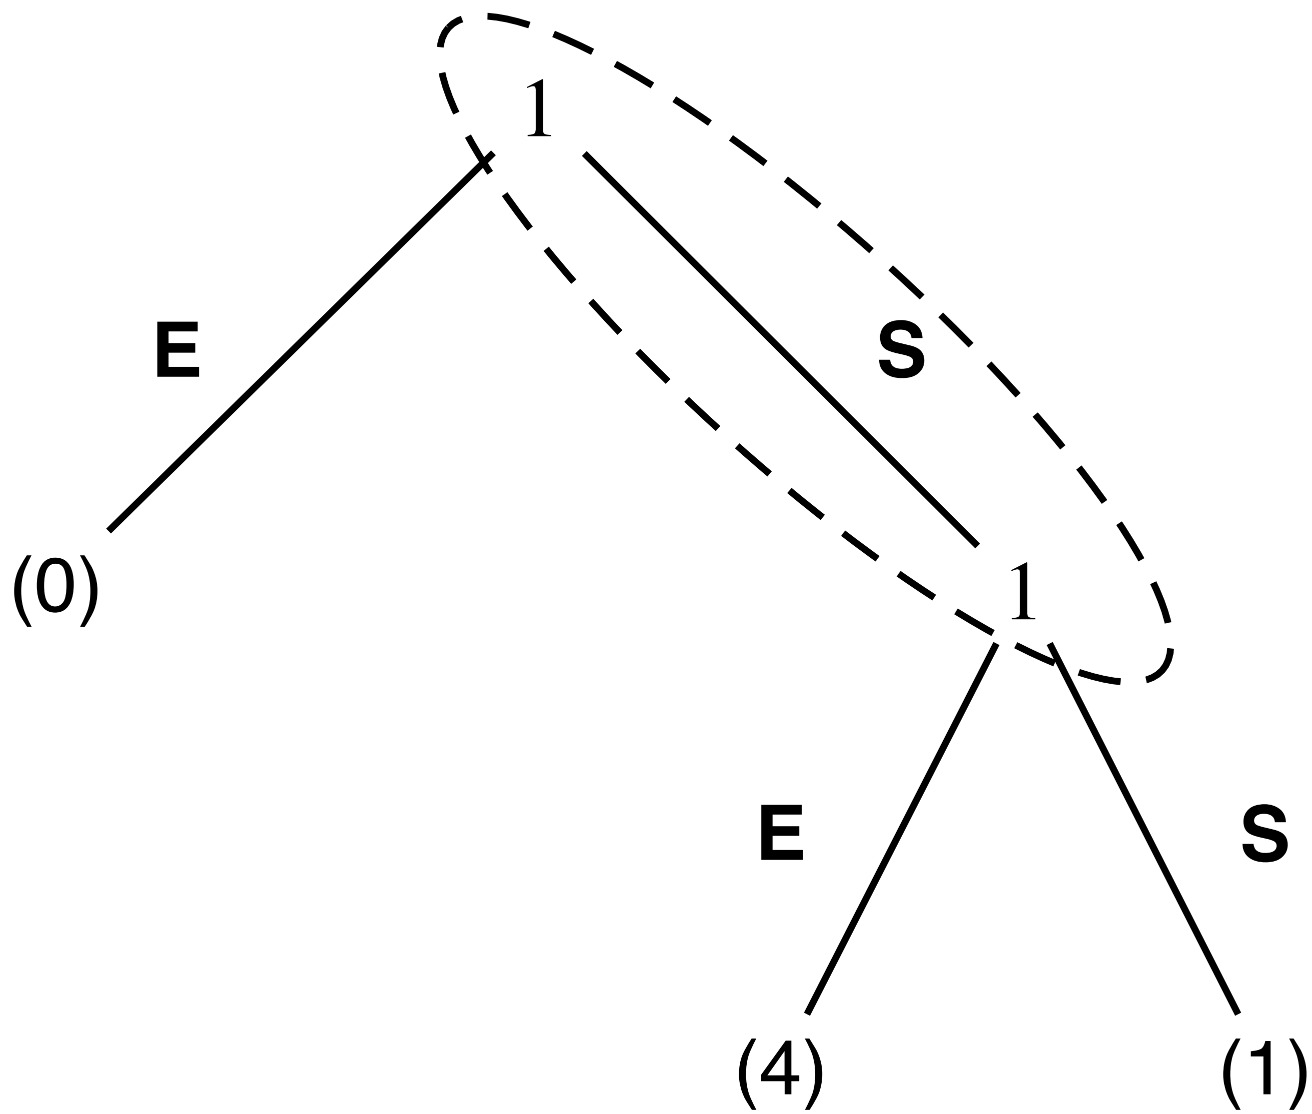

Value of best mixed strategy: 1 (when $p = 0$)

Value of best behavioral strategy: $\frac{4}{3}$ (when $p = \frac{1}{3}$)

Perfect recall games

The previous example involved a player who forgot which moves he had played. What happens if we require players to have perfect memories ?

An extensive form game is called a perfect recall game if whenever x and y belong to the same information set of a player i, then the paths to x and y, restricted to i's information sets and actions, are the same.

Two strategies for a player are called equivalent if they give rise to the same probability distributions over terminal nodes, for every choice of strategies for the other players.

Kuhn's Theorem: for perfect recall games, every mixed strategy is equivalent to some behavioral strategy, and vice-versa.

Useful since sometimes easier to work with behavioral strategies.

Equivalent strategies

The following strategies for player 2 are equivalent:

- Mixed strategy: equal probability to 4 pure strategies
- Mixed strategy: $\frac{1}{2}$ [C(A),F(B)], $\frac{1}{2}$ [D(A), E(B)]
- Behavioral strategy: $\frac{1}{2}$ C at node A, $\frac{1}{2}$ E at node B

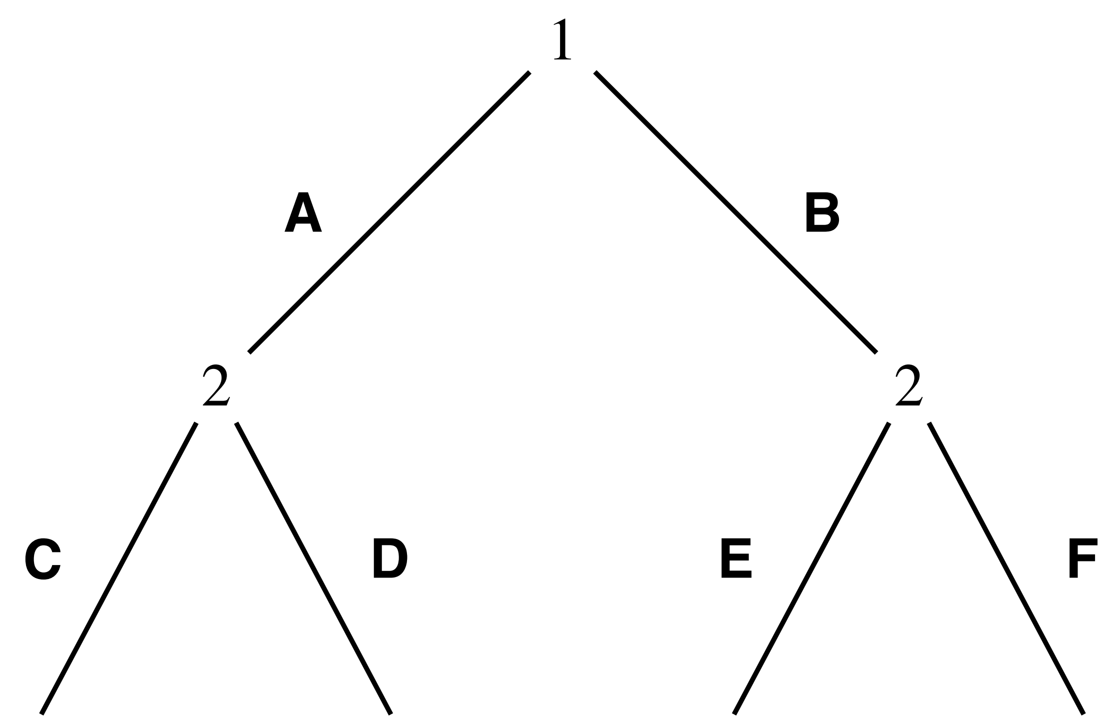

SPE in imperfect-information games

A subgame (of an imperfect-information game) is a subtree of a game tree such that if a node belongs to the subtree then so does all the other nodes in its information set.

In other words, a subgame can contain either all or none of the nodes of an information set.

Given this new definition of subgame, we can extend the notion of subgame perfect equilibria to imperfect-information games.

Backward induction can still be used, but now some of the subgames will involve several players, and so we cannot easily determine the Nash equilibria in the subgames.

Even more refined notions of equilibrium exist for imperfect-information games.