0368.4170: Cryptography and Game Theory	Ran Canetti, Alon Rosen
Lecture 9	
16 December 2009	Fall 2009

Scribe: Opher Lieber

1 Overview

In this lecture we will cover the following topics:

- Subgame Perfect Equilibria (SPE).
- Games with imperfect information.

2 Subgame Perfect Equilibria

In previous lectures, we studied Nash Equilibria in normal form games. A strategy is in NE if no single player can gain by deviating from the strategy. In extensive form games the notion of NE remains the same, if no single player can gain by deviating in any way from the actions prescribed to him by the strategy, then the strategy is in NE. The problem in extensive form games is that players act in turns. This introduces what we call "empty threats": In some cases a player can deviate from his strategy because he knows the other player will see this and be forced to deviate as well, resulting in an overall profit for the first player. Hence the notion of a single player deviation from the overall strategy is too weak for extensive form games. SPE will eliminate these empty threats by requiring ideal play in any sub-game the players find themselves in, i.e. a strategy which is a "best response" at any history.

Example An empty threat.

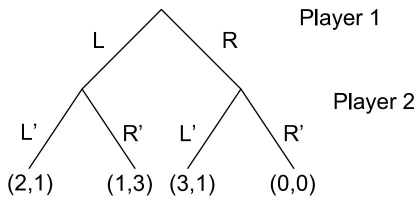

Here we have 2 Nash Equilibria (Recall that, in extensive form games a strategy has to be specified for every node in the game tree, even if the strategy doesn't reach it):

1. NE_1: $S_1(\epsilon) = L, S_2(h) = R' \quad \forall h \in \{L, R\}$
2. NE_2: $S_1(\epsilon) = R, S_2(h) = L' \quad \forall h \in \{L, R\}$

Let's look at NE_1 above. If player 1 deviates and goes right, the strategy dictates he will lose more, as it reaches a result of (0,0). However in extensive form games this is an empty threat, because player 2 will see this and go left, resulting in overall profit for player 1. So rationally, it is in fact profitable for player 1 to deviate from his strategy in this case.

In order to formally define SPE, we first need to introduce and define what exactly a sub-game is.

Definition 1. ***(Subgame)*** *The sub-game of an extensive form game* $G = (N, H, P, (A_i), (u_i))$ *at history* $h \in H$ *is the extensive form game:*

$$G|_h = (N, H|_h, P|_h, (A|_h), (u_i|_h))$$

Where:

- $h' \in H|_h \iff (h, h') \in H$
- $P|_h(h') = p(h, h')$
- $A|_h(h') = A(h, h')$
- $u_i|_h(h') = u_i(h, h')$

Given a strategy S_i and $h \in H$ denote by $S_i|_h$ the strategy that S_i induces in $G|_h$, that is:

$$S_i|_h(h') = S_i(h, h'), \forall h' \in H|_{h'}$$

Definition 2. ***(Subgame Perfect Equilibrium)*** $S = (S_1, \ldots, S_n)$ *is said to be in subgame perfect equilibrium (SPE) in G if:*

$$\forall i \in N, \forall h \in H \setminus Z \;\; s.t. \;\; p(h) = i :$$

$$u_i|_h(S_i|_h, S_{-i}|_h) > u_i|_h(S_i', S_{-i}|_h) \quad \forall S_i' \;\; in \;\; G|_h$$

In words: $S|_h$ is a NE in every subgame $G|_h$

Example Examine NE_1 in the game presented in the beginning of the lecture. This strategy is not an SPE, because if P_1 goes right, P_2 should go left and not right as required by the strategy. SPE requires a best-response by P_2 to any action taken by P_1, i.e. an optimal decision according to the sub-game it finds itself in. For example the following is an SPE for this game:

$$S_1(\epsilon) = R, S_2(h) = \begin{cases} L' & h = R \\ R' & h = L \end{cases}$$

This SPE strategy has P_2 behave according to which subgame (Left or Right) it finds itself in, and provides the best response in that subgame.

The Centipede Game Although SPE is advantageous over NE for extensive form games, it is not perfect itself. Sub-game perfection only considers maximizing a players utility in every sub-game, under the assumption that the next player to act will try to maximize their own utility under that same condition. There is no sense of "cooperation" between players to achieve more ideal utilities. Following is a strong example of a case where SPE gives a far from ideal outcome.

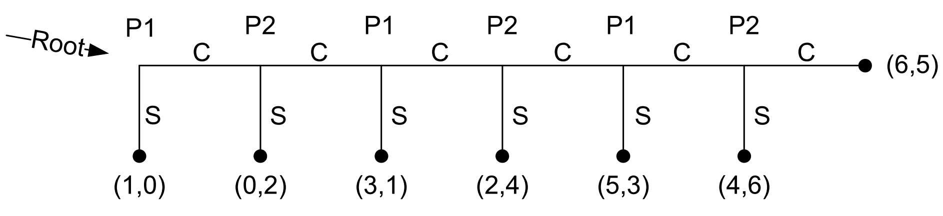

Here we have the Centipede Game for $T = 6$, which can be extended to any T. At each step the current player can choose to continue (go right), or stop (go down) and finish the game.

Define $C(t) = (c, \dots, c)$ to be the sequence in which "continue" is performed t times.
Define $S(t) = (c, \dots, c, s)$ to be the sequence in which "continue" is performed $t - 1$ times, after which "stop" is performed.

The current player function is:

$$p(c(t)) = \begin{cases} 1 & t\ is\ even \\ 2 & otherwise \end{cases}$$

It's easy to see that player $p(C(t))$ will always prefer $S(t+3)$ to $S(t+1)$ to $S(t+2)$. So the dilemma is if to "continue" in hopes the other player will continue as well in which case we increase our payout, on the other hand he risks decreasing his.

Is there an SPE here, and if so what is it? The way to look at it SPE-wise is to check what's the rational thing to do starting from the leaves and working upwards. In our above example let's look at the last decision point, $C(5)$ in this case, where P_2 must choose "continue" or "stop". Obviously the optimal payout strategy here for P_2 is "stop", as this gives him payout of 6 instead of 5, hence the only available strategy for this history which meets the SPE definition is "stop".

Now we work our way up to point $C(4)$ in which case it's P_1's turn. Since the only SPE strategy for P_2 if P_1 goes "continue" is "stop", P_1's expected payout in this case is 4, whereas if he chooses "stop" he gets 5. Hence the SPE strategy for him is also "stop".

We continue like this to the root and reach the conclusion that there is only one unique SPE for this game: Choose "stop" at any history.

So the payout SPE achieves here is (1,0), where in fact an optimal strategy could have given much larger payouts (Especially for a version of the game with a very large T)

In the following lemma we show an effective method to find and prove the existence of an SPE. The SPE construction method we show will employ a similar "Bottom to Top" method to that which we saw in the previous example.

Lemma 3. ***(One Deviation Property)*** *Strategy S is an SPE in an extensive form game G iff:*

$$\forall i \in n \forall h \in H\ s.t.\ p(h) = i :$$

$$u_i|_h(S_i|_h, S_{-i}|_h) > u_i|_h(S_i', S_{-i}|_h)$$

$\forall S_i'$ *in* $G|_h$ *that differs from* $S_i|_h$ *only in the action prescribed immediately after* h.

In other words: it's enough to check only that any deviation from the initial action at the root of the sub-game does not give better utility than the action prescribed by the strategy, in order to show that the strategy is in SPE. The idea is that since SPE strategies need to be optimal in *every sub-game*, then for a given sub-game you don't have to check deviations beyond the initial action, since those sub-trees are already checked in their respective sub-games.

Proof. The first direction is trivial: Let S be an SPE in G, hence in particular it satisfies the One Deviation Property (If S is profitable over all strategies in a sub-game, then in particular is is profitable over strategies which are only different in the initial prescribed action).

Suppose S is not an SPE. We need to show there exists a player i, history h, and strategy S_i' for player i in $G(h)$ which is profitable over $S_i|_h$ and differs from $S_i|_h$ only in the action prescribed immediately after h (i.e. $S_i|_h(h') = S_i'(h')\ \forall\ h' \in H|_h\ s.t.\ h' \neq \epsilon$). Since S is not an SPE, there exists an i, h and at least one strategy S_i' for G(h) s.t.:

$$u_i|_h(S_i|_h, S_{-i}|_h) < u_i|_h(S_i', S_{-i}|_h) \tag{1}$$

For any two strategies S_i, S_i' for a player i define:

$$\mathrm{diff}(\mathrm{S_i}, \mathrm{S_i'}) = |\{\mathrm{h}|\mathrm{p(h)} = \mathrm{i\ and\ S_i(h)} \neq \mathrm{S_i'(h)}\}|$$

In other words the number of histories for which the prescribed action for player i differs between the two strategies. From all the strategies which satisfy (1), choose the strategy S_i' with the minimal amount of different prescribed actions (Such a strategy exists since we assume the game is finite, hence there are a finite number of different prescribed actions for any two strategies). Let h' be the longest history in $G|_h$ for which player i is the active player and the action prescribed by $S_i|_h$ differs from that prescribed by S_i'. Since it's the longest such history, it follows that strategies $S_i|_{h,h'}$ and $S_i'|_{h'}$ differ only in the action prescribed immediately after (h, h'). It remains only to show that $S_i'|_{h'}$ is a profitable deviation in $G(h, h')$, i.e.:

$$u_i|_{h,h'}(S_i|_{h,h'}, S_{-i}|_{h,h'}) < u_i|_{h,h'}(S_i'|_{h'}, S_{-i}|_{h,h'})$$

Suppose by contradiction it is not a profitable deviation in $G(h, h')$. Hence at history (h, h'), player i cannot lose and only gain by acting according to strategy $S_i|_h$ and not S_i'. We can therefore define a strategy S_i'' in $G(h)$ which is equivalent to S_i', except we change the action prescribed for player i at history h' to be the one prescribed by $S_i|_h$, and not the one prescribed by S_i'. Since we can only gain by using S_i'' over S_i' (If the strategy happens to reach h'), S_i'' is also a profitable deviation from $S_i|_h$ in $G(h)$ (Since S_i' was as well). But by construction we have:

$$\mathrm{diff}(\mathrm{S_i|h}, \mathrm{S_i''}) = \mathrm{diff}(\mathrm{S_i|h}, \mathrm{S_i'}) - 1$$

Since S_i'' is just S_i' with the action at h' matched to $S_i|h$. But this is in contradiction to S_i' being the profitable strategy in G(h) with the minimal number of different prescribed actions for player i, hence $S_i'|_{h'}$ is a profitable deviation from $S_i|_{h,h'}$ in $G(h, h')$

□

Proposition 4. *Every* finite *extensive form game with perfect information has an SPE, and* moreover*:*

1. *The SPE consists of pure strategies (No mixing).*
2. *If all utilities for a player are different, then the SPE is* unique.
3. *The proof is constructive, hence it gives you the SPE too.*

The idea for proving the proposition is to start at the leaves. The player whose turn to play is last uses a pure strategy, so the utility at the last decision node is fixed (Just choose the one which gives him a higher utility, or if some are equal pick an arbitrary one amongst them). Then replace this subtree with this "fixed" utility (which is determined by the action we chose in the last level), and continue recursively:

- Determine optimal action for last player.
- Determine optimal action for previous player, given that last player plays optimal.

This method is called "backwards" induction.

Proof. We build an SPE by induction on the length of the subgame $G(h)$ (The length of a subgame being the longest history in the subgame, with the length of a leaf/terminal history being 0). In addition we associate a function $R(h)$ for every history h, which returns the terminal history reached by following the constructed strategy in the subgame $G(h)$. In particular we will end up with a strategy for $G(\epsilon) = G$ which is an SPE, with the exact construction of the strategy being derived from the proof, and the outcome of following the strategy being $R(\epsilon)$.

If the length of $G(h)$ is 0 (i.e. it is a terminal history), define $R(h) = h$, and there is no required action to define for the strategy. Assume by induction we defined strategy S and function R for all histories whose length are at most k. Now let h be a history of length $k+1$ with $p(h) = i$. Since G(h) is of length $k+1$, for any action $a \in A(h)$ it holds that the length of $G(h, a)$ is at most k, hence for any subgame under history h, S and R are already defined. Define $S_i(h)$ to be the action from $A(h)$ which maximizes the utility player i receives at the terminal history defined by R for that child node, i.e.:

$$S_i(h) = max_{a \in A(h)}\{u_i(R(h, a))\}$$

Now, define $R(h)$ to be the terminal history reached by following this action, i.e. $R(h) = R(h, S_i(h))$. Inductively we end up with a strategy S for G.

Using the One Deviation Property (Lemma 3), we now show this strategy is an SPE for G. Let i be any player, and h be any history such that $p(h) = i$. Let S'_i be a strategy for subgame $G(h)$ which differs from $S_i|_h$ only in the action it prescribes immediately after h. By construction, the action $S_i(h)$ maximizes utility for player i over all other possible actions available at history h, and in particular it maximizes with respect to action $S'_i(h)$ since the rest of the strategy is the same (i.e. the rest of the strategy of S'_i is the one we constructed in the proof), hence by the lemma S is an SPE.

□

SPE Construction Example 3-Player Game

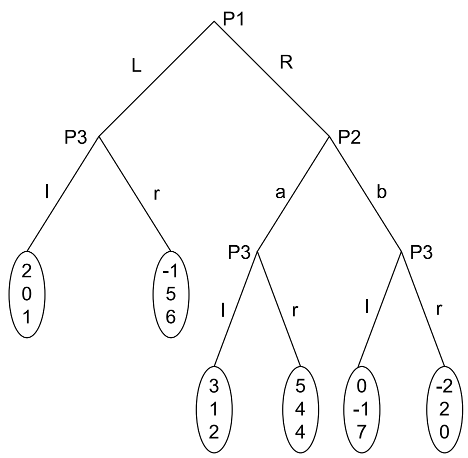

Let's look at this 3-player game. The utilities for players 1,2 and 3 are respectively denoted at the leaves. From the previous theorem there exists an SPE for this game, and moreover, since the utilities for each player are unique across all the leaves, the SPE is unique.

To go about finding the SPE, we first go to all the root decision nodes, i.e. decisions which end the game whichever choice is made. In the above example we have 3 such nodes, all for player 3. In the first case his strategy is obviously "r" as this gives him utility 6 instead of 1. In the middle case it's "r" as well, while in the last it is "l" (7 instead of 0).

Next, we replace each of these subtrees with a leaf node, whose utility is that of the action chosen in the previous step. We now have a new game, which looks like this:

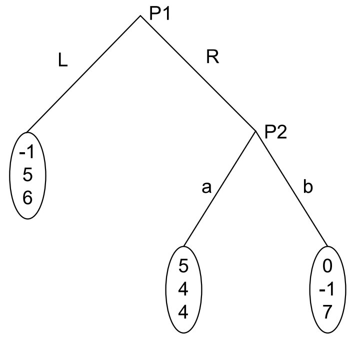

Next, we go to replace the node for player 2. With these utilities, his choice would be "a", giving him utility 4 instead of -1. Replacing this subtree, we have yet another new game:

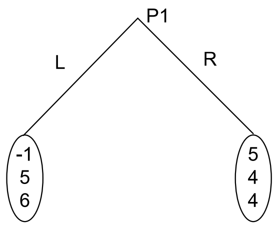

We are left with a single choice game, in which player 1 will prefer "R", giving us a final outcome of (5,4,4). Notice that there was no "choice" here, and only one possible outcome. Because all the utilities for a certain player are different, we will always have a fixed choice: The decision which gives the higher utility. Only in cases where we reach a decision point which gives the same utility for both decisions do we have an arbitrary choice between the two, which results in more than one possible SPE strategy.

3 Extensive games with imperfect information

In perfect information games, players are fully aware of everything that has happened in the game, i.e. what decisions other players have made at previous nodes. More precisely, they know exactly where they are in the tree when it's their turn to act, hence they know all the actions performed to that point. But in reality this is not always the case. We want to model games where this information is not always available.

Intuition Example Matching pennies game in "extensive form".

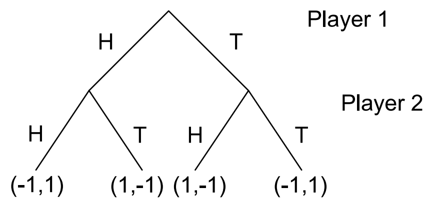

The idea is the same as in the simultaneous version: If the pennies match, player 2 wins, otherwise player 1 wins. In this extensive form version player 2 always wins, as he sees what player 1 did.

We want a way to model simultaneous moves. One way to achieve this is to unite a group of nodes into an "information set":

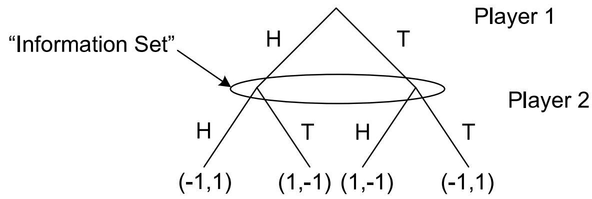

Player 2 doesn't know what action was chosen out of all the nodes in the information set, i.e. doesn't know where he is in the tree.

Note that at every node in an information set a player needs to have the same set of actions. This is because if he were to be told what he can do, he can deduce where he is, and if he isn't told, he doesn't know what he can and can't do.

Definition 5. ***(Game with imperfect information)*** *An extensive form game with imperfect information is a tuple:*

$$G = (N, H, P, (A_i)_{i \in N}, (\mathcal{I}_i)_{i \in N}, (U_i)_{i \in N})$$

Where:

- $N = 1 \dots n$ *is a set of players*
- H *is a (possibly infinite) set of histories, as defined for a regular extensive form game.*
- $P : (H \setminus Z) \to N$ *is a next player function*
- $\forall i \in N$, $\mathcal{I}_i$ *is a partition for all histories* $\{h \in H | p(h) = i\}$ *s.t.*
  - $A(h) = A(h')$ *whenever* h *and* h' *belong to the same* $I_i \in \mathcal{I}_i$
  - *The set* $A(h)$ *for* $I_i \in \mathcal{I}_i$ *is denoted* $A(I_i)$
  - $p(I_i)$ *denotes* $p(h)$ $\forall h \in I_i$
- $u_i : Z \to \mathbb{R}$ *a payoff function, as defined for a regular extensive form game.*

$\mathcal{I}_i$ *is called the "information partition" of* $i \in N$, *and* $I_i \in \mathcal{I}_i$ *is called an "information set" of* $i \in N$.

The idea is that all the histories in $I_i \in \mathcal{I}_i$ are indistinguishable to player i.

Note that perfect information games are a special case of imperfect information games (Information sets are singletons)

Strategies for imperfect information games are pretty much the same as for perfect information games, except we assign an action to each information set, rather than to each history.

Definition 6. *A pure strategy of a player* $i \in N$ *in* G *is a function that assigns an action from* $A(I_i)$ *to each* $I_i \in \mathcal{I}_i$

Example

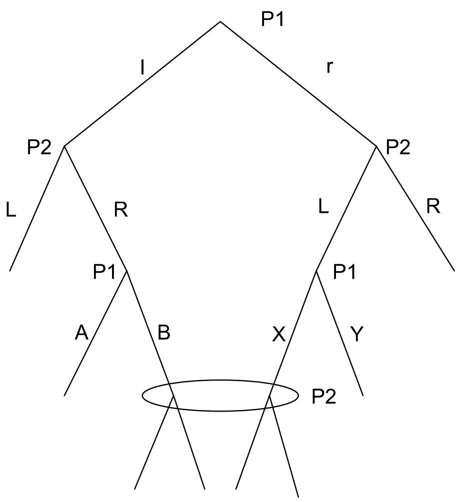

In the above example, we have one non-singleton information set for player 2, denote it X. Suppose now player 2 finds himself at X. By the definition, he can't distinguish between the 2 possible histories in X. But player 2 made a move earlier in the game, at level 2 of the tree, following player 1's first move, 'l' or 'r'. At that point player 2 can distinguish between these 2 histories 'l' and 'r', so he can remember which one it was, and when it finds itself in X, it can know which of the 2 histories it is in (i.e. which side of the tree it is in). This model of a game is called "imperfect recall". In this model players forget the choices and locations made beforehand (i.e., in the above case, whether player 1 initially chose 'l' or 'r').

Definition 7. ***(Perfect Recall)*** *Let $X_i(h)$ be the sequence consisting of all information sets that player $i \in N$ encounters in history $h \in H$ and actions he takes (in the order that they have occurred), then we say the game has "perfect recall" if $\forall i \in N X_i(h) = X_i(h')$ whenever h and h' are in the same information set.*

In words: Consider every pair of nodes in the same information set. Consider the two sequences of nodes on the paths from these nodes to the root. Remove any nodes where player i was not the one making the decision. Then the two sequences must have the same length, and corresponding nodes in the two sequences must belong to the same information sets. A player has imperfect recall if this is not the case, i.e., if there exists a pair of nodes for which the above fails (As in the previous example).

Naturally, following the introduction of SPE for perfect information games, we would like to try to define SPE for imperfect information games. We will show some examples and see why it is not so simple and a bit messy.

Example

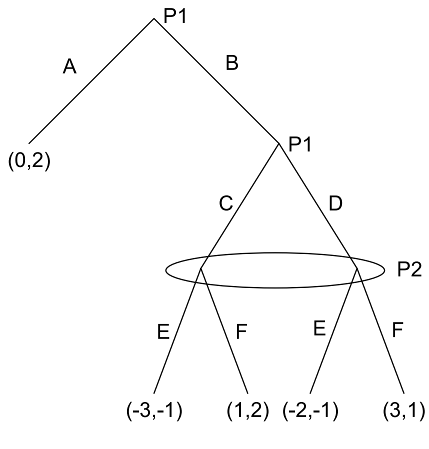

Looking at this example, we have 3 pure NE strategies:

1. $S_1(\epsilon) = A, S_1(B) = D, S_2(B) = E$
2. $S_1(\epsilon) = A, S_1(B) = C, S_2(B) = E$
3. $S_1(\epsilon) = B, S_1(B) = D, S_2(B) = F$

What about an SPE? We need to look at the sub-games, but how can we root a sub-game at an information set? If we look for example at these sub-trees:

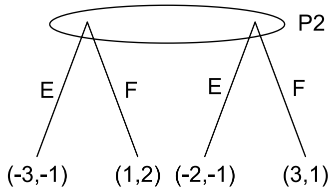

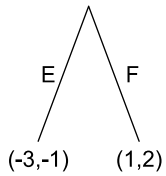

The first is not a valid sub-game as it has no root.

The second is not valid sub-game, as it is not well defined, player 2 can't be "only" in this sub-game, as he can only be in the whole information set.

So essentially, we have only 2 sub-games here, the root, and this one:

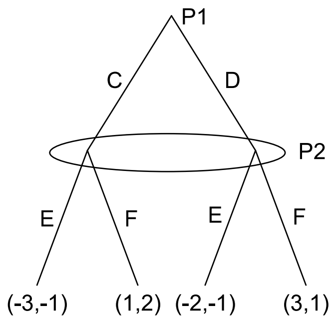

In this case, the only NE for this sub-game is (D, F), in which case the optimal play for player 1 in the root is "B" (Utility 3 rather than 0). Hence the SPE for this game is NE 3 denoted above.

Example

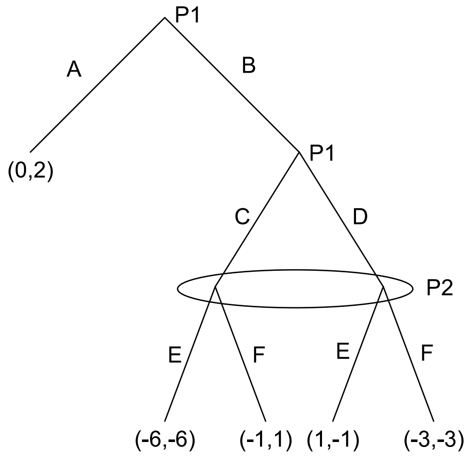

Again, we have only 2 sub-games. Analyzing the sub-game defined by the right-hand side of the tree, we get a simultaneous game equivalent to the following:

		P2: E	P2: F
P1	C	(-6.-6)	(-1,1)
	D	(1,-1)	(-3,-3)

We have 2 pure NE here:

1. (D,E)
2. (C,F)

Hence using SPE we reach the following 2 possible sub-games at the root:

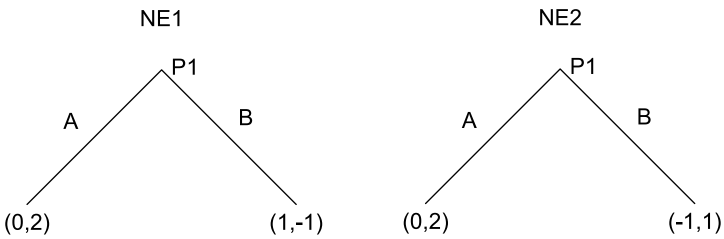

In one case player 1 should go "A", in the other "B", what does SPE dictate in this case? It's not clear, and this example emphasizes one of the issues when trying to define SPE for imperfect information games.

References

[1] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, MIT Press, 1994