AGTA Tutorial 3Solutions

Lecturer: Aris Filos-Ratsikas Tutor: Charalampos Kokkalis

March 7, 2025

Exercise 1. Consider the following game between Alice and Bob. First, Alice proposes to either go to the cinema or for a walk in the park and then Bob decides whether to accept the proposition or to reject it. If they agree to go to the cinema, each player gets a payoff of 2. If Alices proposes to go to the cinema and Bob rejects, Alice gets a payoff of 0 and Bob gets a payoff of 1. If Alice proposes to go for a walk in the park and Bob rejects, they each get a payoff of 1. Finally, if they agree to go for a walk in the park, Alice gets a payoff of 3 and Bob gets a payoff of 0.

A. Formulate the above scenario as a two-player complete information game in extensive form, explicitly presenting the game tree.

B. Convert the extensive form game to a game in normal form, explicitly writing down the utility matrix.

C. Find all the pure Nash equilibria of the game.

D. Explain how to use backwards induction to find a subgame perfect equilibrium of the game.

Solution 1. **A.** The game tree representing the game is the following:

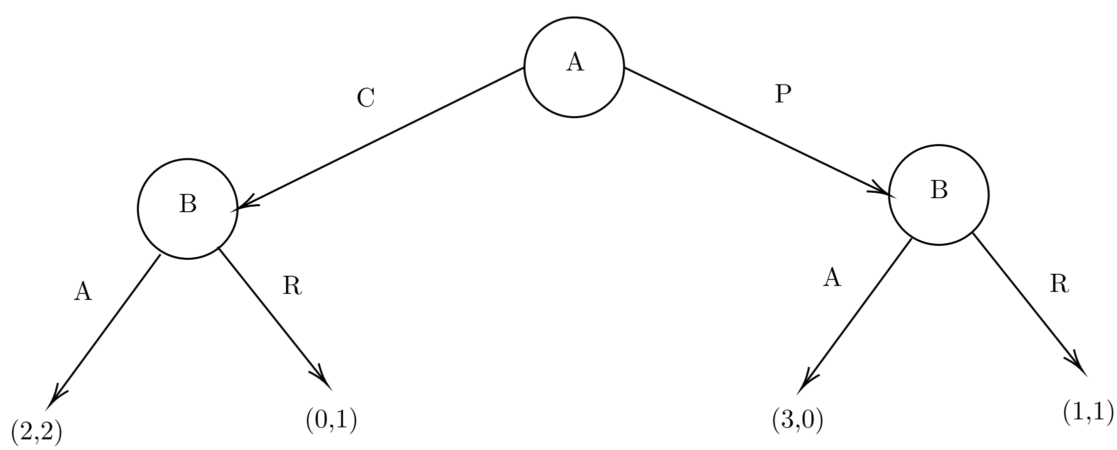

B. We first identify the strategy sets of the two players. For Alice, we have $S_A = \{C, P\}$. For Bob, we have $S_B = \{AA, AR, RA, RR\}$. The corresponding normal form game is the following:

Alice/Bob	AA	AR	RA	RR
C	(2,2)	(2,2)	(0,1)	(0,1)
P	(3,0)	(1,1)	(3,0)	(1,1)

C. We observe that stategy RA for Bob is strictly dominated by strategy AR. Indeed, the utility of RA against C and P is 1 and 0, respectively, whereas the utility of AR against C and P is 2 and 1 respectively. This means that we can eliminate the RA column from the matrix to obtain:

Alice/Bob	AA	AR	RR
C	(2,2)	(2,2)	(0,1)
P	(3,0)	(1,1)	(1,1)

By inspection, we can see that (C, AR) and (P, RR) are the only pure Nash equilibria of the game.

D. We have three subgames to consider. One of them is the whole game rooted at A, and the other two are the two subgrames rooted at B (left) and B (right). These are the *proper* subgames of the game. For the proper subgame on the left, where only Bob plays, playing A is the only PNE. For the proper subgame on the right, playing R is the only PNE. Now given these, in the subgame rooted at A (the original game), Alice's best response is to choose C (as she knows that if she chooses P, Bob will subsequently choose R. So the only subgame perfect equilibrium of this game (when expressed in normal form) is (C, AR).

Exercise 2. Consider the 2-player extensive form game (of imperfect information) described by the game tree below. At the leaves, the left payoff is for Player 1, and the right payoff is for Player 2. Describe all subgame-perfect Nash equilibria in this game, in terms of behavioural strategies.

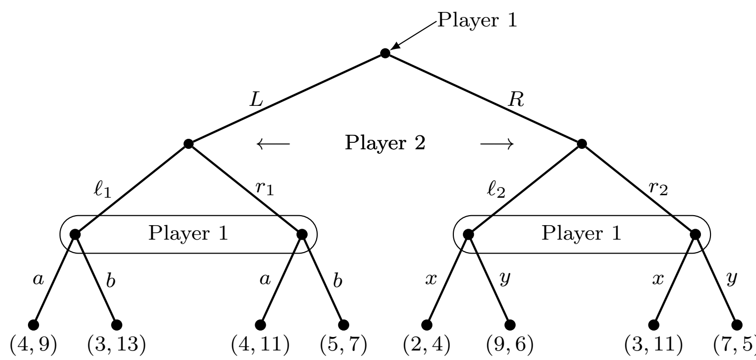

Solution 2. To compute a subgame perfect equilibrium, we need first to define our subgames. We have one subgame which is the whole game rooted at the Player 1 node at the top, and two proper subgames rooted at Player's 2 choices on the left on the right, after Player 1 has chosen L and R respectively. Notice that for a subtree to induce a subgame, it must contain all of the nodes that are in the same information set. Therefore the subtrees rooted at ℓ_1 and r_1 for Player 1 are not subgames, because there are other nodes in the same information set which are not part of the subtree.

To compute subgame perfect Nash equilibria (SPNE), we first compute the Nash equilibria of the left proper subgame. To do that, we have transform this into an equivalent normal form game with the following utility matrix:

Player 1/Player 2	ℓ_1	r_1
a	(4,9)	(4,11)
b	(3,13)	(5,7)

If we solve this game (e.g., using the "Educated Guess" technique), we find that the game has a unique mixed Nash equilibrium (MNE), in which Player 1 plays $(3/4, 1/4)$ for a utility of 4, and Player 2 plays $(1/2, 1/2)$ for a utility of 10.

We next consider the right proper subgame. Again, we can transform the subgame into an equivalent normal form game with the following utility matrix:

Player 1/Player 2	ℓ_2	r_2
x	(2,4)	(3,13)
y	(9,6)	(7,5)

The game has a unique pure Nash equilibrium, namely (y, ℓ_2) which gives Player 1 a utility of 9 and Player 2 a utility of 6. The game does not have any other MNE.

Now, knowing that Player 1's utility from the left subgame is 4 and its payoff from the right subgame is 9, at the root of the tree corresponding to the whole game, Player 1 is going to choose R with probability 1.

So, in the end, at the unique SPNE of the game, the behavioural strategy of Player 1 is to

- play R with probability 1 at the root, and L with probability 0,
- play a with probability 3/4 at the left subtree and b with probability 1/4, and
- play x with probability 0 at the right subtree and y with probability 1.

The behavioural strategy of Player 2 is to

- play ℓ_1 with probability 1/2 at the left subtree, and r_1 with probability 1/2, and
- play ℓ_2 with probability 1 at the left subtree, and r_2 with probability 0.

Exercise 3. Consider the following finite extensive form game of perfect information. There are two players. Each player receives 1 British pound at the beginning of the game. The two players then alternate moves, starting with player 1. In each move, the player whose turn it is to move either chooses `stop` or `give`. If a player chooses `give` then the referee takes 1 pound from that player and gives 2 pounds to the other player. If it chooses `stop`, then the game stops immediately, and both players keep the money they have already accumulated. In any case, the game stops immediately if we reach a state where both players have accumulated exactly 4 pounds.

A. Draw the finite game tree for this game, indicating the payoffs to the two players at the leaves.

B. Compute a subgame perfect Nash Equilibrium (SPNE) for this game.

C. Is there more than one SPNE? Explain.

D. Are there any other NEs? Explain.

E. How would you play this game if you were, say, player 1? Is there a plausible game-theoretic explanation for how you would play it?

Solution 3.

A. Going left in the tree indicates stopping, and going right indicates giving.

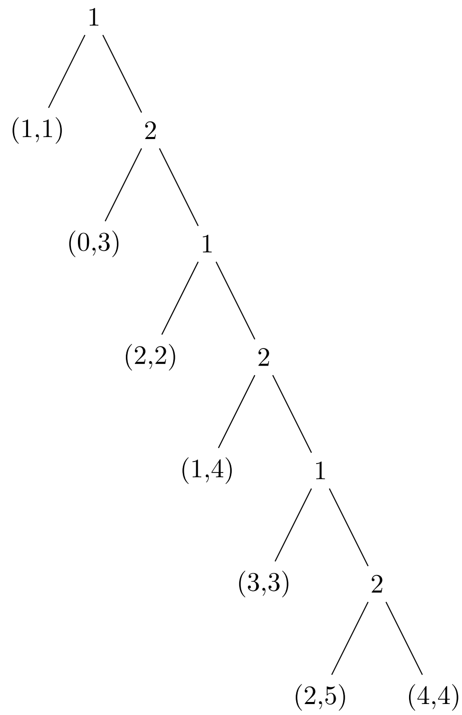

B. Recall that a pure strategy for player i is a function that maps each node controlled by player i to an action available at that node. In this way, the strategy tells Player i what to do at each node controlled by it. (More generally, in the case of imperfect information games, a pure strategy for player i is a function that maps each information set controlled by player i to an action available at that information set.) In this game, each player controlled 3 nodes. We can hence describe a pure strategy for each player as just a tuple, e.g., (G, G, S) is the strategy where the play "gives" in the first node it controls, and "stops" in both of the other two nodes (lower down the tree) that it controls.

We can compute an SPNE for this game by using backwards induction algorithm discussed in class, in the context of Kuhn's theorem. In the lowest proper subgame, rooted at the last node controlled by player 2, choosing "Stop" (S) yields a payoff of 5 to player 2, which is strictly higher than the payoff of 4 that player 2 would obtain by choosing "Give" (G). This, in the (unique) pure SPNE of that subgame, player 2 chooses action S, yielding payoff $(2, 5)$ to the two players.

Knowing this, in the step before, player 1 gets strictly higher payoff of 3 by choosing S, than choosing G and getting payoff 2 (in the SPNE of the subgame below). Hence, in the (unique) SPNE of the subgame rooted at the lowest node for player 1, the action taken by player 1 is S. And so forth, we can work our way back up the game tree, until we reach the root. Thus the SPNE is given in short hand notation by $((S, S, S), (S, S, S))$. In other words, both players choose action "Stop" at every node that they control.

C. Working backwards in the above argument, we see that at each stage the choice S made by the player is because it gets a *strictly* higher payoff by making that choice than by making the other choice G. There is never the case where either player would get exactly the same payoff by choosing either S or G (assuming the already computed SPNE for the lower subgame). This allows us to establish by induction that each subgame, starting from the lower most subgame and working our way up toward the root, has a unique SPNE. Therefore the entire game has a unique SPNE.

D. Consider any pure strategy pair $((S, *, *), (S, *, *))$ for the two players, where each player's first move is S, but thereafter their move can be either G or S (it doesn't matter). We claim that ANY such combination of pure strategies for the two players is a Nash Equilbrium in this game.

To see this, note that indeed, since player 1 starts with S, player 2 cannot possibly improve its own payoff by unilaterally deviating from its own strategy, because against such a strategy for player 1 player 2 can't even change its own payoff no matter what strategy it changes to.

On the other hand, since player 2 plays S at the first node it controls, we know that player 1 cannot improve its own payoff by unilaterally changing its own pure strategy, because against such a pure strategy for player 2, if player 1 chooses G instead of S at the root of the tree then its payoff will decrease from 1 to 0. Moreover, if player 1 only changes its actions elsewhere lower in the tree, it will have no effect on its own payoff (because its own first action makes the game stop immediately).

Thus any pair of strategies of the form ((S,*,*),(S,*,*)) is a pure NE for the game. Likewise, in terms of mixed/behavior strategy NEs, note that any behavior strategy profile $((S, -, -), (S, -, -))$ where the first action chosen by both players is action S with probability 1, and where the subsequent choices at the two lower nodes controlled by each player is ANY probability distribution on the two actions S and G, forms a Nash Equilibrium.

E. This game is indeed very odd. In particular, it doesn't feel that the SPNE or NEs of the game are a good reflection of how humans might actually behave when playing this game.

Consider the same kind of game, but rather than having just 3 nodes belonging to each player, imagine the game was extended to 100 rounds, so to 50 nodes for each player.

I think that if I was confronted with such a game in the "real world", for the first rounds of play I would "take a risk" and Give to the other player, to see if the other player is willing to return the favor and "cooperate with me for a while" so we can both make some money.

It is much harder to argue why, at the very last step of the game, the player whose turn it is to move would do anything other than pick the unique choice (Stop) which maximizes its own payoff. After all, we assume a "rational" player always make choices that maximize its own (expected) payoff.

But that's the troubling aspect: if the other player "knows" that Stop will be chosen at the very last step, then it is also incentivized to choose "Stop" in the prior step, and so on. But this kind of backward reasoning (which is very much related to "iterated illimination of strictly dominated strategies"), would yield both players to choose Stop from the beginning of the game.

If a player could somehow "commit" to the other player that it will play G, for example by yelling out "I promise that I will play (G,G,G)", and if the other player was convinced by this, then the other player's best response to (G, G, G) would give both players a better payoff than just playing the SPNE.

However, there is no mechanism within such a 2-player non-cooperative game for "making firm commitments" about how you will play in the future, since we assume the players choose their moves independently, and we assume that each player is "rational", meaning that its only objective is to maximize its own (expected) payoff.

Exercise 4. Consider the following game between Alice and Bob. Out of a regular deck of 52 cards, one is selected uniformly at random. Alice goes first and decides whether to play or quit. If she quits, the game ends an no player gains or loses any money. If she decides to play, Bob guesses whether the selected card is a King or not. If he guesses correctly, Alice pays him 1000 British pounds, otherwise no exchange of money takes place.

A. State the game above as an imperfect information extensive form game by explicitly writing down the game tree. Make sure to depict the information sets. *Hint: Assume that the card is chosen by a third player that is called* nature *and does not participate in the game, but is placed at the root of the tree.*

B. Convert the game into a normal form game by explicitly writing down the payoff matrix.

C. Find all the mixed Nash equilibria of the game. Would you say that any of these equilibria are counter-intuitive?

Solution 4.

A. The game can me modelled as an imperfect information extensive form game with the following game tree:

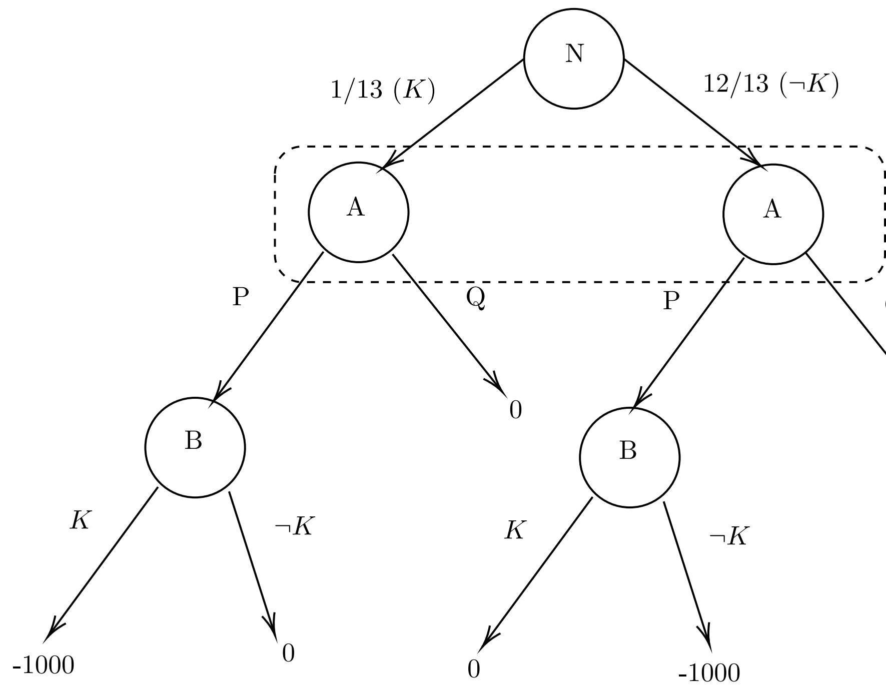

Notice that Alice is aware of whether she has drawn a King or not, but Bob does not have this information. Therefore, the two nodes of Alice resulting from the two choices of nature are in the same information set. Additionally, we do not need to add a node for every possible card, because the only relevant information here is whether the card is a King or not.

B. The payoff matrix of the corresponding normal form game is the following:

Alice/Bob	Guess K (K)	Guess $\neg K$ ($\neg K$)
P with K (PK)	$-1000 \cdot (1/13)$	0
P with $\neg K$ ($P\neg K$)	0	$-1000 \cdot (12/13)$
Q with K (QK)	0	0
Q with $\neg K$ ($Q\neg K$)	0	0

C. We first observe that QK and $Q\neg K$ are both pure Nash equilibria (PNE) of the game. Indeed, Alice receives the maximum possible utility of 0, so she does not want to deviate. Bob also has disutility 0, which is the same as his disutility if he deviates to the other strategy. These equilibria seem quite natural and expected.

Now let's consider any mixed strategy (x_1, x_2, x_3, x_4) of Alice against the mixed strategy (y_1, y_2) of Bob. Assuming that Bob's strategy has both pure strategies in the support (i.e., $y_1, y_2 > 0$), we can write the following system of linear equations to find Alice's strategy:

$$\frac{-1000x_1}{13} = \frac{-12000x_2}{13}$$
$$x_1 + x_2 + x_3 + x_4 = 1.$$

From this, we have that $x_1 = 12x_2$. Note that the system is undertermined, so there are multiple choices of x_1 and x_2 that satisfy the system of equations (depending on the value of x_3 and x_4). Let's assume that $x_3 = x_4 = 0$, meaning that only strategies PK and $P\neg K$ are in the support of Alice's strategy. In that case, we obtain that $x_1 = 12/13$ and $x_2 = 1/13$.

We can also set up a similar system of equations for Bob's strategy:

$$\frac{-1000y_1}{13} = \frac{-12000y_2}{13},$$
$$y_1 + y_2 = 1$$

and compute $y_1 = 12/13$ and $y_2 = 1/13$. Plugging in the equations for Alice and Bob, we get the value $(-1000 \cdot 12)/13^2 \approx -71$, which is the value of the game. Notice that it cannot be the case that Alice plays QK or $Q\neg K$ in any optimal strategy, because then the equations for Bob would be

$$\frac{-1000y_1}{13} = \frac{-12000y_2}{13} = 0,$$
$$y_1 + y_2 = 1,$$

which is infeasble.

This latter MNE is counter-intuitive, because Alice is receiving a negative expected payoff, when, obviously, by always quitting, she can receive a zero payoff. More intuitively, Alice can engage in a game where she can only lose or quite - it seems that rationally it would make sense to quit! But still, there is an equilibrium in which she still plays with some probability.

What makes this equilibrium counter-intuitive, in game-theoretic terms, is that in it, Alice's strategy is weakly dominated by any pure strategy in which she quits the game. There are certain cases for which these weakly dominated equilibria might be ruled off as unnatural (an example is overbidding equilibria in the second-price auction, which we will see later).

Exercise 5.

Consider the following extensive form game of imperfect information with chance nodes. Compute a SPNE of the game. Are there any other NEs for this game?

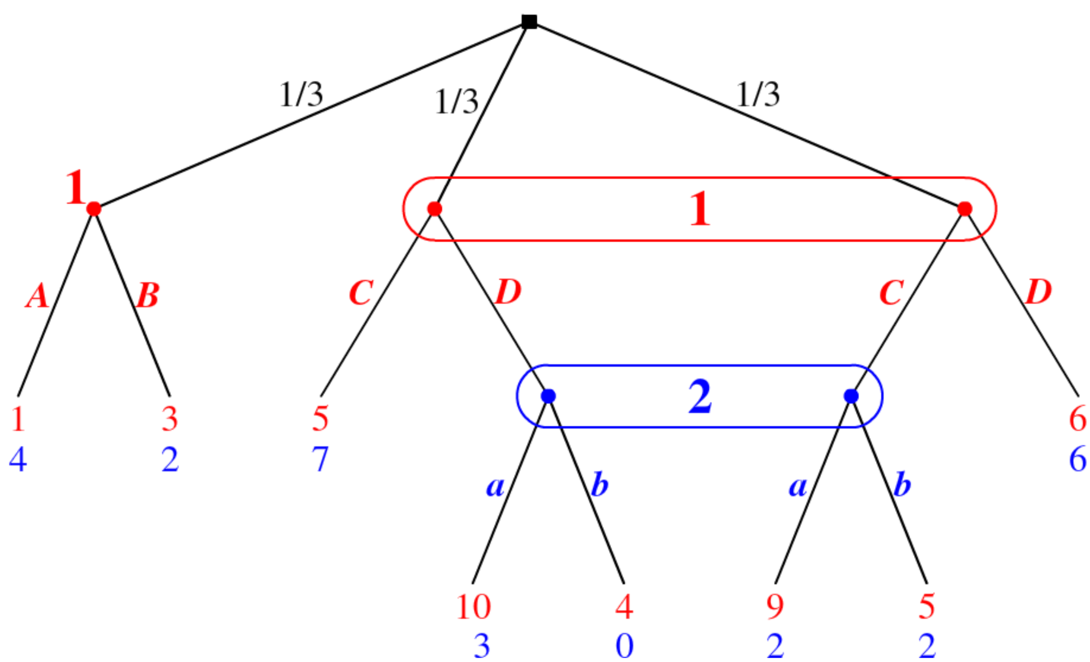

Solution 5. First of all, it is clear that Player 1 will always choose B whenever facing the choice at the leftmost node. Thus, we can and will from now on assume that player 1 will always play B in that leftmost subgame. Thus with 1/3 probability, the payoff to player 1 will be 3, and the payoff to player 2 will be 2. This is in fact the only proper subgame of the game, as a subgame must consist of a subtree with self-contained information sets, and say starting from player 2s information set doesn't form a subtree (it is a forest). Now let us consider the expected payoff overall, to both players. In effect, let us construct the normal form game corresponding to this extensive form game, after the action B at the leftmost node for player 1 has been fixed.

It is not difficult to calculate the expected payoffs to both players under the remaining combinations of pure strategies (actions) for both players.

Specifically, we get the following payoff table:

	a	b
BC	$((3+5+9)/3,(2+7+2)/3)$	$((3+5+5)/3,(2+7+2)/3)$
BD	$((3+10+6)/3,(2+3+6)/3)$	$((3+4+6)/3,(2+0+6)/3)$

Or equivalently,

	a	b
BC	$(17/3,11/3)$	$(13/3,11/3)$
BD	$(19/3,11/3)$	$(13/3,8/3)$

To see the above, note that, for example, if Player 1 plays B and C and player 2 plays "a" then the expected utility (payoff) for Player 1 is $(3+5+9)/3 = 17/3$. We can likewise calculate all of the entries of the above table. (Note that in all these entries, it is always assumed that in the leftmost subtree player 1 plays B, because that is the unique optimal action in that subgame. So, without loss of generality, we can assume player 1 has two possible pure strategies: BC and BD, and of course it can also mix (randomize) between these two strategies.)

Now that we have the above normal form, we can easily calculate the Nash equilibria in this game, all of which will be "subgame perfect", because they already incorporate the fact that player 1 plays B in the leftmost subgame.

Note, in particular, that $((BD),(a))$ is a SPNE for the game, by inspection of the above payoff table: neither player can improve its payoff by switching strategies. Likewise $((BC),(b))$ is also an SPNE for the game, since both players can not *strictly* improve their payoff by unilaterally switching their strategy.

It is also not diffcult to check that there are no other, mixed NEs in this 2×2 normal form game. This is because as soon as player 1 puts positive probability on BD, it is preferrable for player 2 to switch its strategy to put probability 1 on pure strategy "a". Likewise, as soon as player 2 puts any positive probability on strategy "a", it is preferable for player 1 to put probability 1 on pure strategy BD.

The above two (pure) Nash Equilibria are both subgame perfect. So, there are exactly two SPNEs, both of which are pure.

Moreover, there are no other Nash Equilibria of any kind in the game. The reason is that, firstly, the only proper "subgame" of this game is the one in the leftmost subtree, rooted at the node controlled by player 1. But since there is a 1/3 probability that the game will end up in that subgame, player 1 MUST play B with probability 1 in that subgame. Otherwise, if it puts positive probability on the action A, then it can always increase its own expected payoff (no matter what the other player does), by playing action B with probability 1 in that subgame. Hence, in all Nash equilibria (not just in all subgame perfect Nash equilibria), player 1 plays the action B with probability 1 in the leftmost subgame. Hence, there are no other NEs, other than the two pure NEs we have mentioned above.