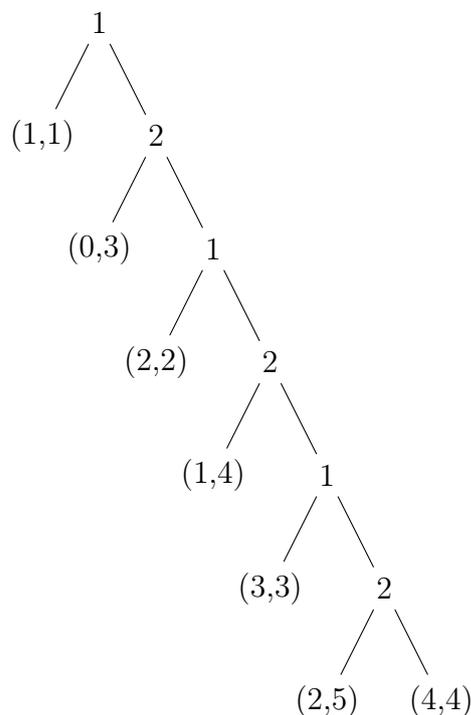


Tutorial sheet 4: solutions

1. (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 player “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 Equilibrium 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.

2. 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:

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

Or equivalently,

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

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 difficult 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 preferable 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.

3. (a) Suppose $x = (x_1, x_2, \dots, x_n)$ is a Nash equilibrium. Consider the product distribution $p(s_1, s_2, \dots, s_n) = \prod_{i=1}^n x_i(s_i)$. We will show that any such Nash equilibrium is also a *correlated equilibrium* (CE).

Recall that, by definition, for all players i and for all pure strategies $s_i, s'_i \in S_i$,

$$U_i^{s'_i}(p | s_i) = \sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) \cdot u_i(s_{-i}; s'_i)$$

is the conditional expected payoff of player i , for playing pure strategy s'_i , having received recommendation s_i from the distribution p , assuming other players play according to their recommendation s_j . Here, by definition,

$$p(s_{-i} | s_i) = \frac{p(s_1, \dots, s_n)}{\sum_{t_{-i} \in S_{-i}} p(t_{-i}; s_i)}$$

Recall that we define p to be a correlated equilibrium iff, for all players i and $s_i, s'_i \in S_i$, we have

$$U_i^{s_i}(p | s_i) \geq U_i^{s'_i}(p | s_i)$$

(In other words, the conditional expected payoff of using the strategy recommended by the recommender is at least as high as choosing any other pure strategy.)

Recall that from the Claim on page 6 of Lec.3, the fact that $x = (x_1, \dots, x_n)$ is a Nash equilibrium is equivalent to the following statement: for all pure strategies $s'_i \in S_i$, $U_i(x) \geq U_i(x_{-i}; s'_i)$. In other words:

$$\sum_{(s_1, \dots, s_n) \in S} \left(\prod_{i=1}^n x_i(s_i) \right) \cdot u_i(s_1, \dots, s_n) \geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} x_j(s_j) \right) \cdot u_i(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$$

Since $p(s_1, \dots, s_n) = \prod_{j=1}^n x_j(s_j)$, we see that

$$\begin{aligned} p(s_{-i} | s_i) &= \frac{\prod_{j=1}^n x_j(s_j)}{\sum_{t_{-i} \in S_{-i}} p(t_{-i}; s_i)} \\ &= \frac{\prod_{j=1}^n x_j(s_j)}{x_i(s_i) \sum_{t_{-i} \in S_{-i}} \prod_{j \neq i} x_j(t_j)} \\ &= \frac{\prod_{j=1}^n x_j(s_j)}{x_i(s_i)} \\ &= \prod_{j \neq i} x_j(s_j) \end{aligned}$$

Thus, by definition,

$$U_i^{s_i}(p | s_i) = \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} x_j(s_j) \right) \cdot u_i(s_{-i}; s_i).$$

Note that therefore $U_i^{s_i}(p | s_i) = U_i(x_{-i}; s_i)$, where $U_i(x_{-i}; s_i)$ is the expected payoff to player i if it were to unilaterally switch to pure strategy s_i , assuming everybody else plays according to the mixed profile x . Moreover, for any pure strategy s'_i for player i , we have

$$U_i^{s'_i}(p | s_i) = \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} x_j(s_j) \right) \cdot u_i(s_{-i}; s'_i)$$

In other words, $U_i^{s'_i}(p | s_i) = U_i(x_{-i}; s'_i)$ is the same as the expected payoff to player i if it were to unilaterally switch to pure strategy s'_i , assuming everyone else plays according to the mixed profile x .

Our aim is to show that, under the assumption that x is an NE, and that player i was recommended s_i under the joint distribution $p(s_1, \dots, s_n) = \prod_{i=1}^n x_i(s_i)$, we must have $U_i^{s_i}(p | s_i) \geq U_i^{s'_i}(p | s_i)$, or equivalently, that we must have $U_i(x_{-i}; s_i) \geq U_i(x_{-i}; s'_i)$.

But this follows immediately from the “useful corollary to Nash’s theorem”, because in order for s_i to be recommended to player i , s_i must be in the

support of x_i (meaning $x_i(s_i) > 0$), and since we are assuming x is a Nash Equilibrium, this means that the pure strategy s_i itself must be a best response to x_{-i} . Hence $U_i(x_{-i}; x_i) = U_i(x_{-i}; s_i) \geq U_i(x_{-i}; s'_i)$, and hence p is a correlated equilibrium.

- (b) We can express the fact that p is a Correlated Equilibrium via a system of LP constraints as follows. Firstly, we express that p must define a distribution on the set S of combinations of pure strategies, as follows:

$$(1) \quad p(s_1, \dots, s_n) \geq 0, \forall (s_1, \dots, s_n) \in S$$

$$(2) \quad \sum_{(s_1, s_2, \dots, s_n) \in S} p(s_1, \dots, s_n) = 1$$

Next, we need to express all constraints of the form: $U_i^{s_i}(p | s_i) \geq U_i^{s'_i}(p | s_i)$, for every pair of pure strategies s_i, s'_i such that s_i can possibly be recommended to player i , in other words, such that $\sum_{t_{-i} \in S_{-i}} p(t_{-i}; s_i) > 0$.

If we write these out explicitly the constraint $U_i^{s_i}(p | s_i) \geq U_i^{s'_i}(p | s_i)$, it looks as follows:

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) \cdot u_i(s_{-i}; s_i) \geq \sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) \cdot u_i(s_{-i}; s'_i)$$

Which is equivalent to:

$$(3) \quad \sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) \cdot (u_i(s_{-i}; s_i) - u_i(s_{-i}; s'_i)) \geq 0$$

But since $p(s_{-i} | s_i) = \frac{p(s_{-i}; s_i)}{\sum_{t_{-i} \in S_{-i}} p(t_{-i}; s_i)}$, and since we know the denominator is positive, by multiplying both sides of the inequality by this denominator, we can rewrite the inequality (3) as:

$$(4) \quad \sum_{s_{-i} \in S_{-i}} p(s_{-i}; s_i) \cdot (u_i(s_{-i}; s_i) - u_i(s_{-i}; s'_i)) \geq 0$$

Note that this is an LP constraint.

Thus the set of Correlated Equilibria of a game can be defined as the set of feasible solutions to a system of linear inequalities. The fact that the set of feasible solutions to a system of linear inequalities is convex is of course well-known, and was discussed in the lectures on LP. (You can prove this yourself easily if you wish, from the definition of convexity.)

(c) Given the game

	A	B
a	(5, 2)	(0, 0)
b	(0, 0)	(2, 5)

Using the methods we have already learned, we see that the only three NE in this game are:

- $[(\frac{5}{7}, \frac{2}{7}), (\frac{2}{7}, \frac{5}{7})]$ with expected payoff both for pl.1 and for player 2 : $\frac{10}{7}$.
- $[(0, 1), (0, 1)]$ with exp. payoff for player 1 = 2 and exp. payoff for player 2 = 5.
- $[(1, 0), (1, 0)]$ with exp. payoff for player 1 = 5 and exp. payoff for player 2 = 2.

Now consider the correlated distribution, p , with probabilities: $p(a, A) = 1/2, p(a, B) = 0, p(b, A) = 0, p(b, B) = 1/2$.

We claim that p is a correlated equilibrium, which is not a NE.

To see that p is a CE, we need to show that the constraints

$$(5) \quad U_i^{s_i}(p | s_i) \geq U_i^{s'_i}(p | s_i)$$

are satisfied for all pure strategies s_i, s'_i for each player i .

But in the context of this specific game, for example, if we let $i = 1$ and we let $s_i = a$ and $s'_i = b$, we have that inequality (5) is equivalent to:

$$U_1^a(p | a) = p(A|a)u_1(a, A) = 5 \geq 0 = p(A | a)u_1(b, A) = U_1^b(p | a)$$

In the same way, it can be checked that all the inequalities of form (5) hold true. Thus, p is a CE. It is clearly not a NE.

Moreover, the expected payoff to each player under this CE is basically the weighted average of their payoff, weighted by the probabilities of the two possible recommendations $p(a, A)$ and $p(b, B)$. Thus, the expected payoff to each player is $5 + 2/2 = 3.5$, yielding a sum total expected utility of 7 for both players, which is as high as in any Nash equilibrium (and note that this CE is also not “biased” toward either player, unlike the two NEs where the sum total expected payoff is 7).

(d) Some possible thoughts are:

- In the game in part (c), the CE given in the solution has social welfare (sum total of expected utilities for all players) that is as high as any NE (but not higher), and is more “fair” than the pure two NEs, which are the only NEs that give the same social welfare. The CE is “fairer” in the sense that both players get the same expected payoff in the CE, by contrast to the pure NEs.

So, as long as the two players have a “randomized recommender” device available to them (which in the case of deciding on a meeting location can be as simple as flipping a fair coin), they we may prefer the CE.

- We might sometimes prefer a CE because it actually gives *strictly* higher social welfare than any Nash equilibrium. For example, consider the following alternative game (see [Aumann,1974], [Ashlagi-Monderer-Tennenholtz,2008]):

$$\begin{array}{cc}
 & A & B \\
 a & (5, 1) & (0, 0) \\
 b & (4, 4) & (1, 5)
 \end{array}$$

Here the two pure NEs have social welfare 6, and there is one other (mixed) NE, in which both players play both pure strategies with probability $1/2$. Under that NE each player gets expected payoff $\frac{5}{2}$, so the social welfare is 5.

By contrast, the following correlated distribution, p , defined by letting $p(a, A) = p(b, A) = p(b, B) = 1/3$ is a CE (it's not hard to check this; feel free to do so), and gives a strictly higher social welfare (namely $\frac{20}{3} = \frac{1}{3}(6 + 8 + 6)$). So, CEs could provide strictly better social welfare than any NE.

- Computing a CE is also less computationally “expensive” (only involving LP) than computing an NE, so in some situations we may find it easier just to get a hold of a CE.