AGTA Tutorial 3Solutions

Lecturer: Aris Filos-Ratsikas Tutor: Charalampos Kokkalis

February 25, 2025

Exercise 1. Assume that you have developed 10 different algorithms for finding mixed Nash equilibria in games. You would like to test these algorithms on 100 instances. In particular, the running time of an algorithm on the set of 100 instances is its maximum running time over all instances in the set. You may assume that for an (algorithm, instance) pair, the running time is always the same.

A randomised algorithm is a probability distribution over your 10 algorithms, i.e., a probability vector $(p_1, \ldots, p_{10})$, where p_i for $i \in \{1, \ldots, 10\}$ is the probability of selecting algorithm i out of your 10 algorithms. Every randomised algorithm naturally has an expected running time on an instance, and its running time is the worst-case expected running time over all instances in the set.

Describe a polynomial-time algorithm for finding the best randomised algorithm for the 100 instances, i.e., the one with the smallest expected running time.

Solution 1. We will model the problem of finding the best randomised algorithm as a 2-player zero-sum game. The maximiser will be the algorithm designer, and its set of pure strategies will be the 10 algorithms for finding mixed Nash equilibria. The minimiser will be the *adversary*, and its set of pure strategies will be the 100 instances on which these algorithms will be tested. The utility of the maximiser given an algorithm A and an instance I will be $u(A, I) = -t(A, I)$, where $t(A, I)$ is the running time of algorithm A on instance I.

We know from the lectures that an optimal mixed strategy of the maximiser should maximise the minimum utility of the maximiser against any pure strategy of the minimiser. A mixed strategy of the maximiser is a randomised algorithm, as defined in the exercise statement, and it's expected utility is the expected running time of the randomised algorithm against the instances used by the minimiser. An optimal strategy maximises the minimum possibly such utility, i.e., it minimises the worst-case running time over the 100 instances. Finding an optimal strategy of the maximiser can be done in polynomial time, via linear programming, as we saw in the lectures.

Exercise 2. Consider the following linear program:

$$
\begin{array}{lll}
\text{maximise} & v & \\
\text{subject to} & v - 2x_1 - 7x_2 & \leq 0 \\
& v - 9x_1 & \leq 0 \\
& v - 4x_1 - 3x_2 & \leq 0 \\
& x_1 + x_2 & = 1 \\
& x_1, x_2 & \geq 0
\end{array}
$$

A. Write down the dual to this linear program.

B. Describe a 2-player zero-sum game (by writing down its utility matrix) that is consistent with this linear program computing an optimal strategy of the maximiser and the dual computing an optimal strategy of the minimiser.

Solution 2.

A. The dual to the given program is the following:

$$\begin{array}{lll} \text{minimise} & w & \\ \text{subject to} & w - 2y_1 - 9y_2 - 4y_2 & \geq 0 \\ & w - 7y_1 - 3y_3 & \geq 0 \\ & y_1 + y_2 & = 1 \\ & y_1, y_2, y_3 & \geq 0 \end{array}$$

B. A 2-player zero-sum game that is consistent with the two linear programs is the following:

$$\begin{bmatrix} 2 & 9 & 4 \\ 7 & 0 & 3 \end{bmatrix}$$

Exercise 3. Solve the following linear program using the simplex method.

$$\begin{array}{ll} \text{maximise} & 6x_1 + 6x_2 + 5x_3 + 9x_4 \\ & \\ \text{subject to} & 2x_1 + x_2 + x_3 + 3x_4 \leq 5 \\ & x_1 + 3x_2 + x_3 + 2x_4 \leq 3 \\ & x_1, x_2, x_3, x_4 \geq 0 \end{array}$$

Show each dictionary and each basic feasible solution produced during the execution of the algorithm. Explain which variable is the entering variable and which one is the leaving variable and why.

Solution 3. First, we reformulate the linear program by defining slack variables w_1 and w_2 as below and constraining their values to be non-negative:

$$\begin{aligned} w_1 &= 5 - 2x_1 - x_2 - x_3 - 3x_4 \\ w_2 &= 3 - x_1 - 3x_2 - x_3 - 2x_4, \end{aligned}$$

giving us the following dictionary:

$$\begin{array}{rcccccc} \zeta = & & +6x_1 & +6x_2 & +5x_3 & +9x_4 \\ \hline w_1 = & 5 & -2x_1 & -x_2 & -x_3 & -3x_4 \\ w_2 = & 3 & -x_1 & -3x_2 & -x_3 & -2x_4 \end{array}$$

$$x_1,\ x_2,\ x_3,\ x_4,\ w_1,\ w_2 \geq 0$$

We select the origin as our initial feasible solution, giving us

$$x_1 = 0,\ \ x_2 = 0,\ \ x_3 = 0,\ \ x_4 = 0,\ \ w_1 = 5,\ \ w_2 = 3,\ \ \zeta = 0.$$

Our first observation is that, since the coefficients of ζ are all positive, we can increase our objective function by increasing any of x_1, x_2, x_3, x_4. Suppose we choose to increase x_4; i.e., we make x_4 our *entering variable.* By how much should we increase x_4? As much as we can, without violating any of our constraints. From our dictionary, we see that only constraints which might apply are w_1 and w_2. In particular, increasing x_4 by $5/3$ would make constraint w_1 go tight; but even before we hit this constraint, we would first hit constraint w_2 at $x_4 = 3/2$. So we set $x_4 = 3/2$ and define w_2 as the *leaving variable* of this iteration. Our new, improved solution is thus

$$x_1 = 0,\ \ x_2 = 0,\ \ x_3 = 0,\ \ x_4 = \frac{3}{2},\ \ w_1 = \frac{1}{2},\ \ w_2 = 0,\ \ \zeta = \frac{27}{2}.$$

Now we wish to reformulate our LP so that ζ and the constraints are written in terms of x_1, x_2, x_3 and w_2, rather than x_1, x_2, x_3, x_4. For this, we note that the equation for w_2 above implies that

$$x_4 = \frac{1}{2}(3 - w_2 - x_1 - 3x_2 - x_3).$$

Making this substitution into our original LP dictionary above gives us the following updated dictionary:

$$\begin{array}{rrrrrr}
\zeta = & \frac{27}{2} & -\frac{9}{2}w_2 & +\frac{3}{2}x_1 & -\frac{15}{2}x_2 & +\frac{1}{2}x_3 \\
\hline
w_1 = & \frac{1}{2} & +\frac{3}{2}w_2 & -\frac{1}{2}x_1 & +\frac{7}{2}x_2 & +\frac{1}{2}x_3 \\
x_4 = & \frac{3}{2} & -\frac{1}{2}w_2 & -\frac{1}{2}x_1 & -\frac{3}{2}x_2 & -\frac{1}{2}x_3 \\
\end{array}$$
$$x_1,\, x_2,\, x_3,\, x_4,\, w_1,\, w_2 \geq 0$$

We continue on in a similar fashion. From our revised LP, we see that we can increase the value of our objective function ζ by increasing x_1 or x_3. Suppose we choose x_1 as our entering variable. Looking at constraints w_1 and x_4, we see that increasing x_1 will hit constraint w_1 (our new leaving variable) first at $x_1 = 1$, since $\min\{1, 3\} = 1$. This gives the following improved solution:

$$x_1 = 1, \;\; x_2 = 0, \;\; x_3 = 0, \;\; x_4 = 1, \;\; w_1 = 0, \;\; w_2 = 0, \;\; \zeta = 15.$$

Rewriting ζ and the constraints in terms of w_1, w_2, x_2, x_3 using the fact that $x_1 = 1 - 2w_1 + 3w_2 + 7x_2 + x_3$ (from w_1 in the previous dictionary) gives the following updated dictionary:

$$\begin{array}{rrrrrr}
\zeta = & 15 & -3w_1 & & +3x_2 & +2x_3 \\
\hline
x_1 = & 1 & -2w_1 & +3w_2 & +7x_2 & +x_3 \\
x_4 = & 1 & +w_1 & -2w_2 & -5x_2 & -x_3 \\
\end{array}$$
$$x_1,\, x_2,\, x_3,\, x_4,\, w_1,\, w_2 \geq 0$$

From here, we see that we have two possible entering variables: x_2 and x_3. Suppose we choose x_3 and increase it until we hit constraint x_1 or x_4. Since we see from the equation for x_1 that this constraint is infeasible (i.e., there is no positive value of x_3 which makes $x_1 = 0$), we know that the leaving vairable must be x_4, with the constraint going tight at $x_3 = 1$. Using the fact that $x_3 = 1 + w_1 - 2w_2 - 5x_2 - x_4$, we rewrite the LP dictionary as follows:

$$\begin{array}{rrrrrr}
\zeta = & 17 & -w_1 & -4w_2 & -7x_2 & -2x_4 \\
\hline
x_1 = & 2 & -w_1 & +w_2 & +2x_2 & -x_4 \\
x_3 = & 1 & +w_1 & -2w_2 & -5x_2 & -x_4 \\
\end{array}$$
$$x_1,\, x_2,\, x_3,\, x_4,\, w_1,\, w_2 \geq 0$$

Since our objective function ζ has all-negative coefficients, increasing any of w_1, w_2, x_2, x_4 can only decrease its value. Thus, we have maximized our solution at $(x_1, x_2, x_3, x_4) = (2, 0, 1, 0)$ and $(w_1, w_2) = (0, 0)$, with objective value $\zeta = 17$.

Exercise 4. Consider the following linear program.

$$\begin{array}{ll}
\text{maximise} & 2x_1 + x_2 \\
& \\
\text{subject to} & 2x_1 + x_2 \leq 4 \\
& 2x_1 + 3x_2 \leq 3 \\
& 4x_1 + x_2 \leq 5 \\
& x_1 + 5x_2 \leq 1 \\
& x_1, x_2 \geq 0
\end{array}$$

A. Solve the LP above using the simplex method. Show each dictionary and each basic feasible solution produced during the execution of the algorithm. Explain which variable is the entering variable and which one is the leaving variable and why.

B. Solve the LP above by drawing the feasible region in two dimensions and checking the objective function value on each of its corners.

Solution 4.

A. Original dictionary:

$$\begin{array}{rcrr}
\zeta = & & +2x_1 & +x_2 \\
\hline
w_1 = & 4 & -2x_1 & -x_2 \\
w_2 = & 3 & -2x_1 & -3x_2 \\
w_3 = & 5 & -4x_1 & -x_2 \\
w_4 = & 1 & -x_1 & -5x_2 \\
\multicolumn{4}{c}{x_1,\, x_2,\, w_1,\, w_2,\, w_3,\, w_4, \geq 0}
\end{array}$$

Initial solution:

$$x_1 = 0,\ \ x_2 = 0,\ \ w_1 = 4,\ \ w_2 = 3,\ \ w_3 = 5,\ \ w_4 = 1,\ \ \zeta = 0.$$

Possible entering variables: x_1, x_2. Suppose we select x_1. Then the corresponding leaving variable would be w_4, since $\min\{2, 3/2, 5/4, 1\} = 1$. Updated solution:

$$x_1 = 1,\ \ x_2 = 0,\ \ w_1 = 4,\ \ w_2 = 3,\ \ w_3 = 5,\ \ w_4 = 0,\ \ \zeta = 2.$$

Rewriting LP dictionary using the fact that $x_1 = 1 - w_4 - 5x_2$:

$$\begin{array}{rcrr}
\zeta = & 2 & -2w_4 & -9x_2 \\
\hline
w_1 = & 2 & +2w_4 & +9x_2 \\
w_2 = & 1 & +2w_4 & +7x_2 \\
w_3 = & 1 & +4w_4 & +19x_2 \\
x_1 = & 1 & -w_4 & -5x_2 \\
\multicolumn{4}{c}{x_1,\, x_2,\, w_1,\, w_2,\, w_3,\, w_4, \geq 0}
\end{array}$$

Since both coefficients of ζ are negative, we know that our solution, $(x_1, x_2) = (1, 0)$ and $\zeta = 2$, must be maximal.

B. Since this optimization problem is a linear program, we know that its solution lies at one of the vertices of the feasible region (assuming the feasible region is bounded). Graphing each of the constraints, we find that the feasible region lies at the intersection of the constraints $x_1 \geq 0$, $x_2 \geq 0$, and $x_1 + 5x_2 \leq 1$, see Figure 1.

The feasible region is the dark purple triangle in the bottom left. The vertices of this feasible region are $(0, 0)$, $(0, 1/5)$ and $(1, 0)$; and value of our objective function ζ at each of these vertices is as follows:

$$\zeta(0, 0) = 0,\ \ \zeta(0, 1/5) = 1/5,\ \ \zeta(1, 0) = 2.$$

Thus, we confirm what we showed in part A: that the solution $x_1 = 1$, $x_2 = 0$ maximizes ζ over the feasible region, and that $\zeta(1, 0) = 2$.

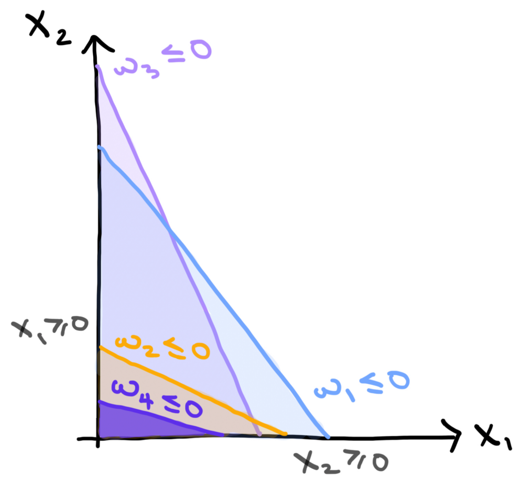

Figure 1: The feasible region of the given LP. *Credit for the figure goes to Kat Molinet.*

Exercise 5 (Game of Chicken)**.** Consider the following game of *Chicken*, where two cars are headed towards each other in opposite directions at full speed. Each driver can choose to either swerve, (or "chicken") or to continue going straight, (or "dear"). If they both dear, they crush and their cars get destroyed. If they both chicken, they both leave unharmed, but they don't get the satisfaction of besting their opponent. If one chickens and the other dares, then the daring driver is the winner. The utility bimatrix of the game is given below.

Driver 1 / Driver 2	**Chicken**	**Dare**
Chicken	6,6	2,7
Dare	7,2	0,0

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

B. Find a mixed Nash equilibrium of the game which is not a pure Nash equilibrium.

C. Assume that some trusted party (e.g. a traffic warden or a traffic light) is giving the two players some *advice* on how to play the game. In particular, the party first announces to the players that the possible strategy profiles of the game are either (C,D), (D,C) or (C,C) and each one happens with equal probability 1/3. Then, the party chooses one of the three profiles at random and lets the players know of their strategies in that profile, without letting them know of the strategy of the other player. We will consider some different kind of equilibrium, where the the players do not want to deviate from their prescribed advice, assuming their opponent is following the advice.

- Assume that Player 1 witnesses the advice "Dare". Prove that assuming that Player 2 sticks to his prescribed strategy (whatever that is), Player 1 does not want to deviate from playing "Dare".
- Assume that Player 1 witnesses the advice "Chicken". Prove that assuming that Player 2 sticks to his prescribed strategy (whatever that is), Player 1 does not want to deviate from playing "Chicken". Here, not wanting to deviate from playing "Chicken" means that his expected payoff from playing "Chicken" (over the randomness of the profile distribution of the trusted party) is higher than that of playing "Dare".

D. Compare the total utility (sum of players' utilities) of all the equilibria that you have computed. What do you observe?

Solution 5.

A. It is easy to see by inspection that (C,D) and (D,C) are the only pure Nash equilibria (PNE) of the game. Indeed, in (D,C) for example, Player 1 would receive a utility of 6 by deviating to C, which is smaller than 7, the utility which the player currently receives. If Player 2 deviated to D, the player would receive a utility of 0, which is smaller than 2. The argument for (C,D) being a PNE is symmetric.

B. Since we have found the PNE, we can now find a mixed Nash equilibrium (MNE) with full support, i.e., one in which both strategies are played with positive probability for every agent. Let (x_1, x_2) be the mixed strategy of Player 1 and (y_1, y_2) be the mixed strategy of Player 2. By a proposition that we saw in the lectures, we know that each pure strategy in the support of a player's equilibrium strategy must yield the same utility against the mixed equilibrium strategy of the other player. This gives rise to the following system of inequalities.

For Player 1 we have:

$$\begin{aligned} 6y_1 + 2y_2 &= 7y_1 \\ x_1 + x_2 &= 1 \end{aligned}$$

For Player 2, we have:

$$\begin{aligned} 6x_1 + 2x_2 &= 7_1 \\ y_1 + y_2 &= 1 \end{aligned}$$

Solving this system we get the following MNE: $(2/3, 1/3), (2/3, 1/3)$.

C. The solution concept that we will explore in this question is known as a *correlated equilibrium*, and was proposed by Robert Aumann in 1974 [1]. An intuitive interpretation is that there is a probability distributionσ over possible pure strategy profiles. A trusted third party (often referred to as "the mediator") is going to draw a pure strategy profile s from the distribution σ, and will prescribe the corresponding strategy to each player as advice. In a correlated equilibrium, each player maximises their expected utility when they follow the advice of the mediator.

Formally, let σ be joint distribution $\sigma \in \Delta(S_1 \times S_2 \times \ldots \times S_n)$ over strategy profiles, which we will refer to a *correlated strategy* in the game. Such a correlated strategy profile σ is a *correlated equilibrium (CE)* if, for any player $i \in N$ and any *swap function* $\delta_i : S_i \to S_i$, it holds that

$$\mathbb{E}_{s\sim\sigma}[u_i(s)] \geq \mathbb{E}_{s\sim\sigma}[u_i(\delta_i(s_i), s_{-i})]$$

An alternative, equivalent definition is the following:

$$\mathbb{E}_{s\sim\sigma_{-i}|s_i}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{s\sim\sigma_{-i}|s_i}[u_i(s_i', s_{-i})]$$

for any $s_i, s_i' \in S_i$, where $\sigma_{-i}|s_i$ is the conditional distribution of profiles s_{-i}, induced when conditioning on the strategy of Player i being s_i.

Notice that when σ is a product distribution, then each player i receives a recommendation x_i about the probability of playing each of its own pure strategies $s_j \in S_i$, independent of the strategies of the others. This is precisely a mixed strategy for the player, and $(x_1, \ldots, x_n j)$ is a mixed Nash equilibrium. Hence, every MNE is a CE (but the converse is not true, as we will see in the exercise below).

We now move on to answering the questions of the exercise.

- If Player 1 witnesses the advice "Dare", then they know that the profile that was drawn from σ was either (D,C) or (D,D) (because (C,C) or (C,D) are not consistent with their observation of "Dare"). The conditional distribution $\sigma_{-1}|s_1$ (conditioned on s_1 = D) is C with probability 1 and D with probability 0 (as the pair (D,D) appears with probability 0 in σ, and hence in $\sigma_{-1}|s_1$). In other words, since Player 1 received "Dare", they know that the strategy of Player 2 is "Chicken" with probability 1. Their expected utility from following the advice is 7; this is the highest possible utility they can receive in the game, so obviously they do not want to deviate.

- If Player 1 witnesses the advice "Chicken", then they know that profile that was drawn from σ was either (C,C) or (C,D) (because (D,C) or (D,D)) are not consistent with their observation of "Chicken"). The conditional distribution $\sigma_{-1}|s_1$ (conditioned on $s_1 = \text{C}$) is C with probability 1/2 and D with probabity 1/2 (since the probability of (C,C), as well as the probability of (C,D) in σ was 1/3). Their expected utility from following the advice of "Chicken" is therefore $6/2+2/2 = 4$. If they were to choose "Dare" instead, their expected utility would be $7/2+0/2 = 3.5$, hence smaller than the one of their prescribed advice.

We conclude that the distribution where (C,D), (D,C) and (C,C) happens with equal probability 1/3 is a correlated equilibrium.

From the PNE equilibria computed in the first subquestion, the total utility (also known as the *social welfare*) is $7+2 = 9$. The MNE found in the second subquestion has total (expected) utility $\frac{4}{9} \cdot 12 + 2 \cdot \frac{2}{9} \cdot 9 + \frac{1}{9} \cdot 0 = \frac{28}{3} \approx 9.3$. Finally the CE found in the third subquestion has total (expected) utility $2 \cdot \frac{1}{3} \cdot 9 + \frac{1}{3} \cdot 6 = 10$. We observe that the total utility of the CE is higher than that of any MNE that is not also a CE.

This means that if we "relax" the solution concept to allow more profiles as equilibria, we can sometimes find some that have higher total utility.

Exercise 6 (Bonus Coding Exercise)**.** In your programming language of choice, code a solver for 2-player zero-sum games. Your solver should take as input a utility matrix in appropriate format (e.g., in Python this could be a list of lists) and output the optimal strategies for the maximiser and the minimiser, as well as the value of the game.

References

[1] Robert J. Aumann. Subjectivity and correlation in randomized strategies. In *Journal of Mathematical Economics* 1(1): 67-96. 1974.