ADFOCS Exercises for *Equilibrium Computation in Games*

Bernhard von Stengel, Part 1, 28 August 2024

Exercise 1. In the following introductory 2×2 zero-sum game,

Max \ min	l	r
T	1	0
B	−2	2

draw the max-min payoff and min-max cost against the mixed opponent strategy, and find the max-min and min-max strategies and corresponding values.

Exercise 2. Consider the following Penalty game. It gives the probability of scoring a goal when the row player (the striker) adopts one of the strategies L (shoot left), R (shoot right), and the column player (the goalkeeper) uses the strategies l (jump left), w (wait then jump), r (jump right). The row player is interested in maximizing and the column player in minimizing the probability of scoring a goal.

Max \ min	l	w	r
L	0.6	0.7	1.0
R	1.0	0.8	0.7

(a) Find all equilibria of this game in mixed (including pure) strategies, and their equilibrium payoffs. Why is the payoff unique?

(b) Now suppose that player I has an additional strategy M (shoot down the middle), so that the payoff matrix is

Max \ min	l	w	r
L	0.6	0.7	1.0
R	1.0	0.8	0.7
M	1.0	0.0	1.0

Find an equilibrium of this game, and the equilibrium payoff. [*Hint*: *The result from* (a) *will be useful.*]

Exercise 3. Prove that if $(\overline{x}, \overline{y})$ is a Nash equilibrium in the zero-sum game $(A, -A)$, then $\overline{x}$ is a max-min strategy and $\overline{y}$ is a min-max strategy.

Exercise 4. Consider the following simultaneous game between two players: Each player chooses an integer between 1 and 5 (e.g., some number of fingers shown with your hand), and the higher number wins, except when it is just one larger than the lower number, in which case the lower number wins. So 4 beats 2 and 5 but it loses against 3. Equal numbers are a draw.

(a) Write down the matrix of payoffs to the row player of this zero-sum game. A winning player gets payoff 1, a draw gives payoff 0.
(b) Show all pairs of strategies of the row player where one strategy weakly or strictly dominates the other, and indicate the type of domination.
(c) Argue carefully, without calculations, why the value of this game has to be zero.
(d) Find an equilibrium of this game in mixed (including pure) strategies, and explain why it is an equilibrium.

Exercise 5. *(Challenging exercise.)* Consider the following *infinite* zero-sum game between an *inspectee* and an *inspector*. The players choose their actions from the unit time interval. The inspectee violates at some time s in $[0, 1)$, and the inspector inspects at some time $t \in [0, 1)$, in addition to an automatic inspection that always takes place at time 1. The *time to detection* is maximized by the inspectee and minimized by the inspector, where detection happens at time t if $s < t$ and at time 1 if $s \geq t$. The inspection at time t is *not observed*, so s and t are chosen simultaneously. The payoff $V(s, t)$ to the inspectee is therefore

$$V(s,t) = \begin{cases} t - s & \text{if } s < t, \\ 1 - s & \text{if } s \geq t. \end{cases}$$

The optimal solution will involve randomization. Which times will be chosen by the players? Find a continuous optimal distribution for the inspector (it will be restricted to a certain interval), and if you can also for the inspectee.

Exercise 6. Consider the symmetric 3×3 game (A, B) with

$$A = \begin{pmatrix} 0 & 3 & 0 \\ 2 & 2 & 2 \\ 3 & 0 & 0 \end{pmatrix}, \qquad B = A^{\top} = \begin{pmatrix} 0 & 2 & 3 \\ 3 & 2 & 0 \\ 0 & 2 & 0 \end{pmatrix}.$$

Find all its symmetric and non-symmetric equilibria. [*Hint*: *Change the best-response diagram for a symmetric game shown in the lecture to the pair of diagrams for the bimatrix game.*]

ADFOCS Exercises for *Equilibrium Computation in Games*

Bernhard von Stengel, Part 2, 30 August 2024

Exercise 7. Consider the following game in extensive form between players 1 and 2, and the chance player 0. It is a zero-sum game with payoffs to player 1.

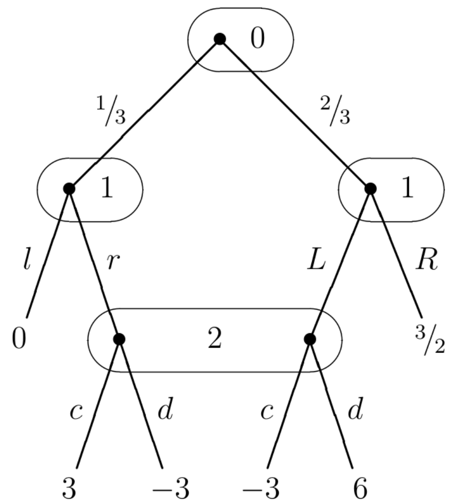

(a) Write down the strategic form of this game, and find optimal strategies and the value of the game.
(b) Write down the sequence form of the game with the sequences for the players, their constraint systems with $Ex = e$ and $Fy = f$ for the realization plans x and y, and the payoff matrix A. Find optimal realization plans x and y for the corresponding LP.

Exercise 8. Consider the following zeros-sum game in extensive form between players 1 and 2, with payoffs to player 1.

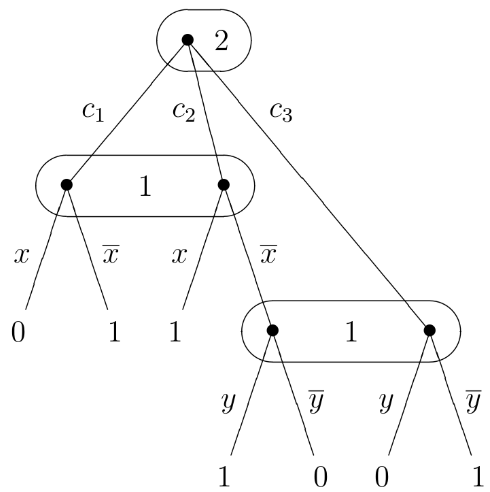

(a) Explain why this game does not have perfect recall.
(b) Write down the strategic form of this game, and find optimal strategies and the value of the game.

(c) This game has been constructed from the Boolean formula in conjunctive normal form with three clauses (in parentheses) $(\overline{x}) \wedge (x \vee y) \wedge (\overline{y})$ with Boolean variables x and y. Explain why this formula is not satisfiable, and why this fact can be told from the value of the game. Explain how to construct in this way from a general Boolean formula in conjunctive normal form an extensive game with imperfect recall where the value of the game tells whether the formula is satisfiable or not.

(d) *(A bit more challenging.)* Suppose that the players are only allowed to use behavior strategies, and find a max-min strategy for the maximizing player 1 in the above game in behavior strategies, and the corresponding value. Compare this with the max-min value when player 1 is allowed to play a mixed max-min strategy.

(e) In light of (c) and (d), what does this say about the possibility of solving a zero-sum game with imperfect recall in polynomial time?

Exercise 9. *("Pivoting changes signs".)* Let $y, z \in \mathbb{R}^m$ be adjacent vertices of a simple polytope P with facet normal vectors (in $\mathbb{R}^m$) $c, a_2, \ldots, a_m$ for y and $d, a_2, \ldots, a_m$ for z. (These assumptions mean the following: There are reals $\beta_0, \beta_1, \ldots, \beta_m$ such that $c^\top x \le \beta_0$, $d^\top x \le \beta_1$, and $a_i^\top x \le \beta_i$ for $i = 2, \ldots, m$ are some of the inequalities that define that $x \in P$, and turning any such inequality into an equality defines a facet of P. The polytope is simple means that every vertex is on exactly m, but not more, facets. The corresponding m vectors for these facets are then linearly independent, such as $c, a_2, \ldots, a_m$ that define the facets that y lies on, and similarly $d, a_2, \ldots, a_m$ that define the facets of z.) Show that the determinants $|c\ a_2 \cdots a_m|$ and $|d\ a_2 \cdots a_m|$ have opposite sign. [*Hint*: *Use the linear dependence of the* $m+1$ *vectors* $c, d, a_2, \ldots, a_m$.]

Exercise 10. A nonnegative $m \times m$ matrix Y is a *Markov chain* if its rows define probability distributions on m "states", that is, if $Y\mathbf{1} = \mathbf{1}$. A probability distribution x on m states is a *stationary* distribution for Y if $x^\top Y = x^\top$. Show that if Y is a Markov chain on two states, then the off-diagonal elements of Y, in their columns, are proportional to the probabilities of a stationary distribution x of Y. That is, if

$$Y = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

then $(c\ \ b) = (c + b)x^\top$. Exactly when is x unique? Argue carefully.

Solution to Exercise 1

The left shows the diagram for the max-min payoff, which is given by 0.4 when prob(B) = 0.2 as the maximum of the lower envelope of the possible responses l and r of the minimizer. The right shows the diagram for the min-max cost, which is given also by 0.4 when prob(r) = 0.6 as the minimum of the upper envelope of the possible responses T and B of the maximizer.

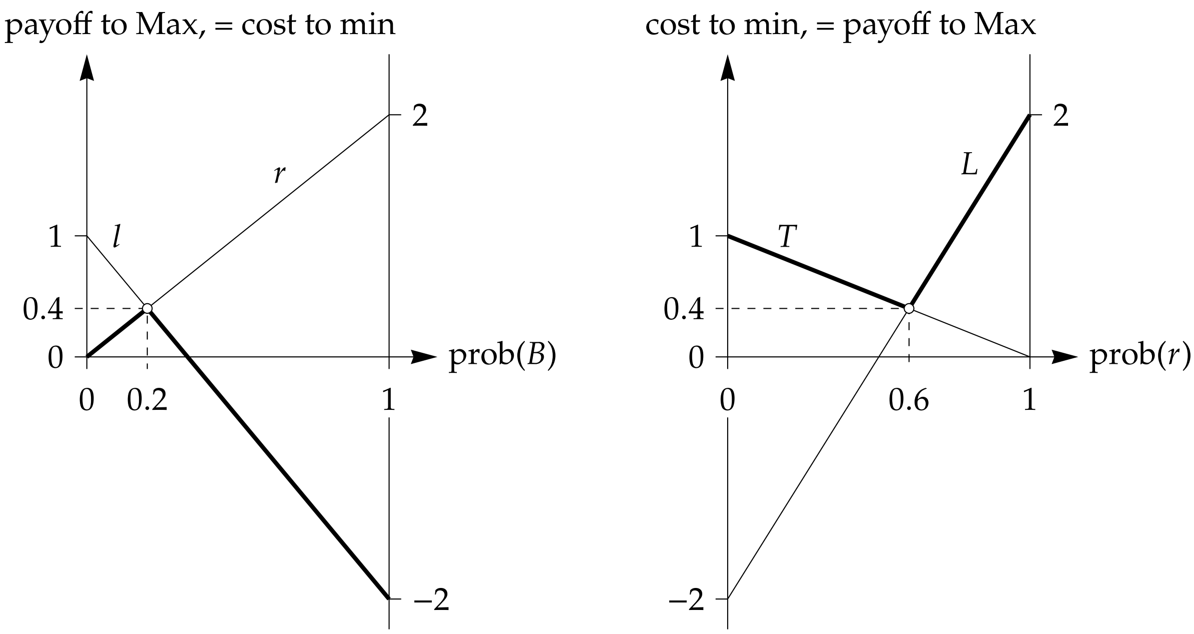

Solution to Exercise 2

(a) The following best responses show that the game does not have an equilibrium in pure strategies:

Max \ min	l	w	r
L	0.6	0.7	1.0
R	1.0	0.8	0.7

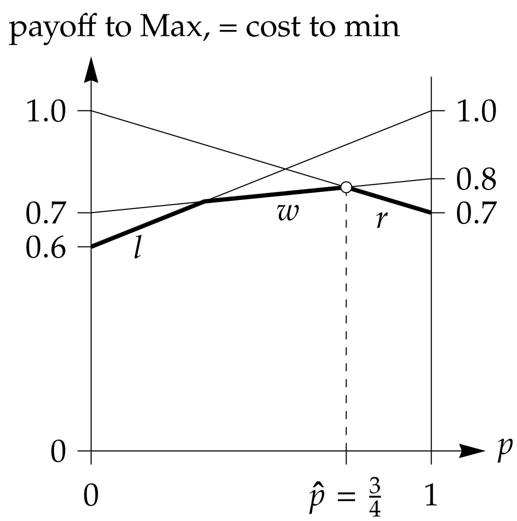

An equilibrium strategy of player I is a max-min strategy, with maximum worst-case expected payoff to player I, given by the respective best response of player II. The above goalpost diagram shows the expected payoff to player I if the strategy of player II is l, w, or r. For a max-min strategy of player I, one has to consider the minimum of his expected payoffs (the lower envelope), which is shown as the bold line.

The picture shows that the max-min payoff is 0.775 when $p = \text{prob}(R) = \frac{3}{4}$, namely when player I receives equal expected payoff when player II plays r or w. This equality, namely the equation $1(1-p) + 0.7p = 0.7(1-p) + 0.8p$, determines $p = \frac{3}{4}$.

A quick way of finding the probabilities for L and R is given by the difference trick: In row L, the difference of the matrix entries for w and r is 0.3. In row R, the difference is 0.1. Hence, L and R have to be played with inversely proportional probabilities, namely ratio 1 to 3, that is, L has probability $\frac{1}{4}$ and R has probability $\frac{3}{4}$. The equal resulting cost of 0.775 for w and r is easily verified.

Furthermore, the strategy l of player II has expected cost 0.9, and thus will not be played by II as a best response.

In order to get an equilibrium, player II has to make player I indifferent between L and R by assigning suitable probabilities to w and r. Again, we use the difference trick: The difference among the payoffs for column w is 0.1, whereas for column r it is 0.3. Hence w and r have probabilities $\frac{3}{4}$ and $\frac{1}{4}$, respectively, giving player I expected payoff 0.775 in each case. Because only best responses have positive probability, this is an equilibrium.

This is a zero-sum game, so the equilibrium payoff is the unique value of the game.

(b) Now player I has M as an extra strategy. We first try the solution obtained so far. It works if player II can still keep his min-max cost of 0.775 by using the above probabilities $(0, \frac{3}{4}, \frac{1}{4})$ even when player I has the extra possible response. Indeed, by playing M player I would only get 0.25 and this is not a best response for player I, so he does not play M and the 2×3 solution carries over to the 3×3 game. This solves the game, because it is a zero-sum game and one equilibrium is sufficient for us because any other equilibrium would have the same payoff. (It is in fact the only equilibrium.)

Solution to Exercise 3

In fact, we show a stronger property: If $(\overline{x}, \overline{y})$ is an equilibrium of $(A, -A)$, then the minimax theorem holds. Namely, $(\overline{x}, \overline{y})$ is an equilibrium if and only if

$$\max_{x \in X} x^\top A \overline{y} = \overline{x}^\top A \overline{y} = \min_{y \in Y} \overline{x}^\top A y \,.$$

Then

$$\begin{aligned} \overline{x}^\top A \overline{y} &= \min_{y \in Y} \overline{x}^\top A y \\ &\le \max_{x \in X} \min_{y \in Y} x^\top A y \\ &\le \min_{y \in Y} \max_{x \in X} x^\top A y \\ &\le \max_{x \in X} x^\top A \overline{y} = \overline{x}^\top A \overline{y} \,. \end{aligned}$$

Hence, all inequalities are equalities, and we have $\max\min = \min\max$, and $\overline{x}$ is a max-min strategy and $\overline{y}$ is a min-max strategy.

Solution to Exercise 4

(a) The game matrix looks as follows, with payoffs to player I.

I \ II	1	2	3	4	5
1	0	1	−1	−1	−1
2	−1	0	1	−1	−1
3	1	−1	0	1	−1
4	1	1	−1	0	1
5	1	1	1	−1	0

(b) There is no strict domination. Both 4 and 5 weakly dominate 1 (strategy 1 only beats 2 which is also beaten by 4 and 5), and 5 weakly dominates 2 (because 2 only beats 3 which is also beaten by 5).

(c) A zero-sum game has a unique value. The game is symmetric, so when changing the players and changing the signs of the payoffs its value changes sign but must be the same, which is only possible if the value is zero. In more detail: The value is the max-min payoff to player I and min-max cost to player II. If the value was positive, then player I would get a positive payoff and player II could adopt player I's strategy and also gain a positive amount, which is not possible. If the value was negative, then player II would get a positive amount (in terms of her own payoff), which could similarly be obtained by player I, also a contradiction. So the value is zero.

(d) By Proposition 8.7 of *Game Theory Basics* (and also easy to see), eliminating weakly dominated strategies does not change the value of a zero-sum game, so we can omit strategies 1 and 2 for both players when we look for the value of the game. The resulting 3×3 matrix game is the rock–scissors–paper game (with 3 for rock, 4 for scissors, 5 for paper) where the optimal mixed strategy is to randomize uniformly. Hence, in the original 5×5 game an equilibrium strategy for each player is to play the strategies $1, 2, 3, 4, 5$ with probabilities $0, 0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}$. Then the unplayed strategies get negative expected payoff (-1 for 1 and $-\frac{1}{3}$ for 2) and the three used strategies get payoff 0, which is optimal. Although not asked, this is in fact the unique equilibrium because strategies 1 and 2 are not best responses and cannot be played in any equilibrium of this zero-sum game, and the remaining strategies 3, 4, and 5 are uniquely played with probabilities $\frac{1}{3}$ to give payoff zero for each. There is no equilibrium in rock–scissors–paper where one of the three strategies is not played with positive probability. For example, if the column player never uses paper, then for the row player rock strictly dominates scissors (tying instead of losing against rock, and winning instead of tying against scissors), and the mixed equilibrium of the resulting 2×2 game does not extend to the 3×3 game.

Solution to Exercise 5

The game describes a kind of "duel" with time reversed where both players have an incentive to act early but after the other. The inspector can clearly garantee a detection time of $\frac{1}{2}$ by choosing $t = \frac{1}{2}$, but can in fact do better. The following considerations are best understood by considering a discretized version of the game where s and t can only be chosen as multiples of 0.01, for example. By the best-response condition, the inspectee's payoff is too small if he violates too late, so that he will select s with a certain probability distribution from an interval $[0, b]$ where $b < 1$. Consequently, the inspector will not inspect later than b. We prove that the game has the following solution: The inspector chooses the inspection time t from $[0, b]$ according to the density function $p(t) = 1/(1 - t)$, where $b = 1 - 1/e$. The inspectee chooses his violation time s according to the distribution function $Q(s) = 1/e(1 - s)$ for $s \in [0, b]$, and $Q(s) = 1$ for $s > b$. The value of the game is $1/e$.
Proof: The inspector chooses $t \in [0, b]$ according to a density function p, where

$$\int_0^b p(t)dt = 1. \tag{1}$$

The *expected* payoff $V(s)$ to the inspectee for $s \in [0, b]$ is then given as follows, taking into account that the inspection may take place before or after the violation:

$$\begin{aligned} V(s) &= \int_0^s (1-s)p(t)dt + \int_s^b (t-s)p(t)dt \\ &= \int_0^s p(t)dt + \int_s^b t\,p(t)dt - s\int_0^b p(t)dt. \end{aligned}$$

If the inspectee randomizes as described, then this payoff must be constant (this can also be proved formally, see S. Karlin (1959), Mathematical Methods and Theory in Games, Programming, and Economics, Vol. II: The Theory of Infinite Games. Addison-Wesley, Reading, Dover reprint 1992, Lemma 2.2.1, p. 27). That is, its derivative with respect to s, which by (1) is given by $p(s) - s\,p(s) - 1$, should be zero. The density for the inspection time is therefore given by

$$p(t) = \frac{1}{1-t},$$

with b in (1) given by $b = 1 - 1/e$. The constant expected payoff to the violator is then $V(s) = 1/e$. For $s > b$, the inspectee's payoff is $1 - s$ which is smaller than $1/e$ so that he will indeed not violate that late.
The optimal distribution of the violation time has an atom at 0 and a density q on the remaining interval $(0, b]$. We consider its distribution function $Q(s)$ denoting the probability that a violation takes place at time s or earlier. The mentioned atom is $Q(0)$, the derivative of Q is q. The resulting expected payoff $-V(t)$ for $t \in [0, b]$ to the inspector

is given by

$$\begin{aligned} V(t) &= t \cdot Q(0) \;+\; \int_0^t (t-s)q(s)ds \;+\; \int_t^b (1-s)q(s)ds \\ &= t \cdot Q(0) \;+\; t\int_0^t q(s)ds \;+\; \int_t^b q(s)ds \;-\; \int_0^b s\,q(s)ds \\ &= t \cdot Q(t) \;+\; Q(b) \;-\; Q(t) \;-\; \int_0^b s\,q(s)ds\,. \end{aligned}$$

Again, the inspector randomizes only if this is a constant function of t, which means that $(t-1)\cdot Q(t)$ is a constant function of t. Thus, because $Q(b) = Q(1-1/e) = 1$, the distribution function of the violation time s is for $s \in [0,b]$ given by

$$Q(s) = \frac{1}{e}\cdot\frac{1}{1-s},$$

and for $s > b$ by $Q(s) = 1$. The nonzero atom is given by $Q(0) = 1/e$.

Solution to Exercise 6

We first show (as in the lecture) the payoff matrix C with strategy labels ①,②,③ for both players. The corresponding triangle subdivided into pure best responses (rows) against a mixed strategy strategy z (for playing the columns, which represent the *same* strategies) is constructed with an upper envelope as shown on the right, which is particularly easy to visualize because of the row ② with constant payoffs 2.

	①	②	③
①	**0**	**3**	**0**
②	**2**	**2**	**2**
③	**3**	**0**	**0**

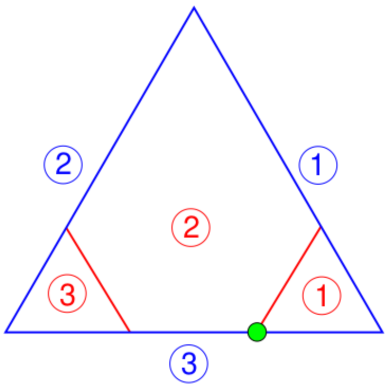

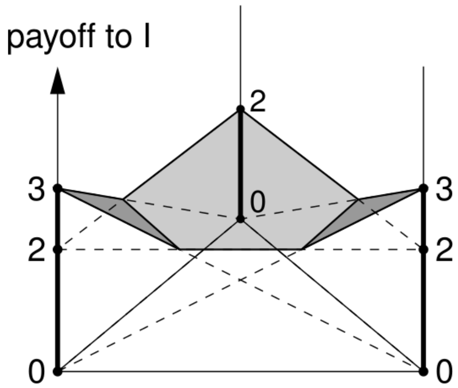

The small circle indicates the unique symmetric equilibrium (z,z) of this symmetric game $(C, C^\top)$ given by $z = (\frac{1}{3}, \frac{2}{3}, 0)^\top$.

We can use the same diagram by now considering the bimatrix game $(C, C^\top)$ with players mixing independently with a strategy pair (x,y). Then the columns have new labels as pure strategies ④,⑤,⑥ of player 2, with C shown on the left:

	④	⑤	⑥
①	**0**	**3**	**0**
②	**2**	**2**	**2**
③	**3**	**0**	**0**

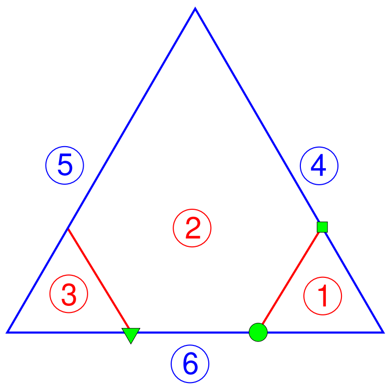

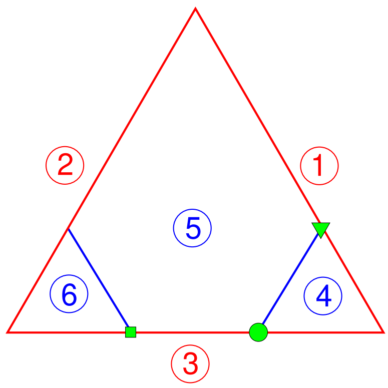

Here the first triangle shows the mixed strategy set Y of the column player 2 (now with the sides of the triangle labeled with ④, ⑤, ⑥), subdivided by the best responses of player 1. The second triangle is the same with the players's roles exchanged. We obtain two further mixed equilibria (x, y) (pair of small triangles) with $x = (0, \frac{2}{3}, \frac{1}{3})^\top$ and $y = (\frac{2}{3}, \frac{1}{3}, 0)^\top$, and (y, x) (pair of small squares). They are distinct because the equilibria are not symmetric, that is, $x \neq y$. These non-symmetric equilibria of a symmetric game obviously come in pairs, and the number of symmetric equilibria is odd.

Solution to Exercise 7

(a) The strategic form is the following, with the min-max costs plotted on the right:

	c	d
(l, L)	-2	4
(l, R)	1	1
(r, L)	-1	3
(r, R)	2	0

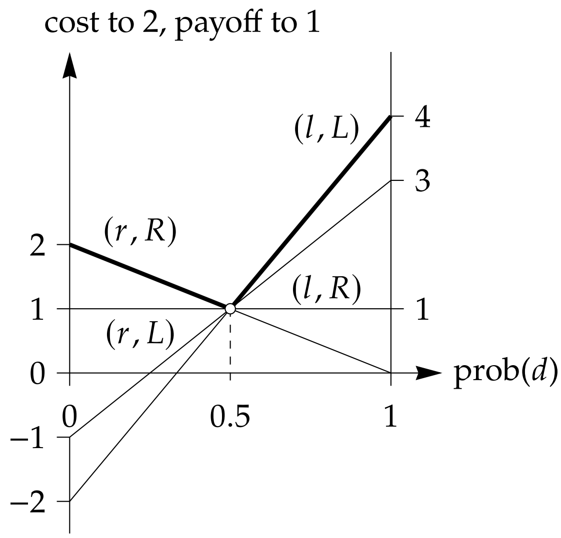

The min-max strategy of player 2 is clearly to play $(\frac{1}{2}, \frac{1}{2})$ with min-max value 1. An obvious max-min strategy of player 1 is the pure strategy (l, R) with constant expected payoff 1. There are further optimal mixed strategies. Any of them that uses the first or third row also needs to play the fourth row. One further max-min strategy is $(0, 0, \frac{1}{3}, \frac{2}{3})$ which mixes between (r, L) and (r, R) (and thus effectively only the moves L and R while playing r for sure), and another one is $(\frac{1}{4}, 0, 0, \frac{3}{4})$; this mixed strategy is *not* a behavior strategy (why?). Any other optimal mixed strategy of player 1 is a convex combination of these two and the pure strategy (l, R), which as a mixed strategy is the probability vector $(0, 1, 0, 0)$.

(b) Player 1 has the sequences $\emptyset$, l, r, L and R, and player 2 has the sequences $\emptyset$, c and d. These index the components of the realization plans $x = (x_\emptyset, x_l, x_r, x_L, x_R)$ and $y = (y_\emptyset, y_c, y_d)$. The constraint matrices E and F in $Ex = e$ and $Fy = f$ are

$$E = \begin{pmatrix} 1 & & & & \\ -1 & 1 & 1 & & \\ -1 & & & 1 & 1 \end{pmatrix}, \quad e = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \qquad \text{and} \qquad F = \begin{pmatrix} 1 & & \\ -1 & 1 & 1 \end{pmatrix}, \quad f = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

With the sequences $\emptyset, l, r, L, R$ and $\emptyset, c, d$ indicating rows and columns, respectively, the (sparsely represented) payoff matrix A is

$$A = \begin{pmatrix} 0 & & & & \\ & & & 1 & -1 \\ & & & -2 & 4 \\ 1 & & & & \end{pmatrix}.$$

Note that the payoffs at the leaves are multiplied by the chance probability of getting to the leaf to obtain the entries of A. The first row of A is zero since the sequence $\emptyset$ is not defined by a leaf. Because this game is zero-sum, $B = -A$.
We now solve the LP for the sequence form given by

$$\begin{array}{lrcrcl} \text{minimize} & e^\top u & & & & \\ \text{subject to} & E^\top u & - & Ay & \geq & \mathbf{0} \\ & & & Fy & = & f \\ & & & y & \geq & \mathbf{0} \end{array}$$

where we know a min-max realization plan for player 2 is given by $y = (y_\emptyset, y_c, y_d) = (1, \frac{1}{2}, \frac{1}{2})$ because y_c and y_d are the probabilities for the moves c and d. Then $Ay = (0,0,0,1,1)^\top$. A corresponding optimal dual solution $x = (x_\emptyset, x_l, x_r, x_L, x_R)$ to the dual LP has to fulfill (as a best response) $Ex = e$ and, by complementary slackness,

$$x \geq \mathbf{0} \quad \perp \quad E^\top u \geq Ay\,.$$

Because $x_\emptyset = 1$, this means that we have to have equality in the first row and thus, with $u = (u_0, u_1, u_2)^\top$, that $u_0 - u_1 - u_2 = 0$. Given that $Ay = (0,0,0,1,1)^\top$, we have in fact $E^\top u = Ay$ with $(u_0, u_1, u_2) = (1, 0, 1)$ where $u_1 = 0$ is the partial payoff to player 1 at the first information set with moves l, r and and $u_2 = 1$ is the partial payoff to player 1 at the second information set with moves L, R and $u = 1$ the payoff to player 1 overall.
Of interest is obtaining the realization plans for the optimal strategies of player 1. For the pure strategy (l, R) this is clearly just the realization plan $(1,1,0,0,1)$. For the mixed strategy $(r, L)\frac{1}{3} + (r, R)\frac{2}{3}$ this is the convex combination of the "pure" realization plans given by $(1,0,1,1,0)^\top\frac{1}{3} + (1,0,1,0,1)^\top\frac{2}{3} = (1,0,1,\frac{1}{3},\frac{2}{3})$. More interestingly, for the mixed (non-behavior) strategy $(l, L)\frac{1}{4} + (r, R)\frac{3}{4}$ this is given by $(1,1,0,1,0)^\top\frac{1}{4} + (1,0,1,0,1)^\top\frac{3}{4} = (1,\frac{1}{4},\frac{3}{4},\frac{1}{4},\frac{3}{4})$, which is a *convex combination* of the previous two realization plans. Note that the set of realization plans is convex, whereas the set of behavior strategies is not – their convex combinations include in particular the convex combinations of pure strategies, which are mixed strategies.

Solution to Exercise 8

(a) This game does not have perfect recall because at the second information set of player 1, the left node is preceded by the sequence with move $\overline{x}$ where is the right node is preceded by the empty sequence of own earlier moves. Hence, this information set indicates that the player has forgotten *whether* they have made a move or not.

(b) The strategic form is the following:

	c_1	c_2	c_3
(x, y)	0	1	0
$(x, \overline{y})$	0	1	1
$(\overline{x}, y)$	1	1	0
$(\overline{x}, \overline{y})$	1	0	1

The value of the game is $\frac{2}{3}$, which can be seen with the optimal strategies $(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for player 1 and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for player 2.

(c) The formula is not satisfiable because the first clause $(\overline{x})$ states that x must be set to false, and similarly the last clause $(\overline{y})$ that y must be set to false, but then the middle clause $(x \vee y)$ is not satisfied. This can be told from the value of the game because the strategies of player 1 are the truth assignments to the Boolean variables, and the columns are the clauses of the formula and how they are satisfied. A satisfying assignment corresponds to an all-1 row, and only then the game has value 1 (otherwise the uniform distribution on the columns gives a value less than 1 for each row).

For a general Boolean formula in conjunctive normal form, player 2 moves once and chooses a clause. Each Boolean variable is represented by one information set of player 1, with nodes reached (along a path of successive moves of player 1) by the clauses that contain that variable. Player 1 chooses whether to set that variable to true or false, and the eventual payoff is 1 if the clause is true, otherwise 0. Then the strategic form of the game is essentially the (exponentially long) truth table for the formula as in this example. This is the construction (on page p. 534) of D. Koller and N. Megiddo (1992), The complexity of two-person zero-sum games in extensive form, Games and Economic Behavior 4, 528–552, that finding the value in an extensive zero-sum game without perfect recall is NP-hard.

(d) For player 1's behavior strategy, let p be the probability of player 1 choosing $\overline{x}$, and let q be the probability of choosing $\overline{y}$. Then the expected cost to player 2 if she chooses c_1, c_2, c_3 is p, $(1-p)+p(1-q)$, and q, respectively, and player 2 will choose the minimum of these expressions. This is best countered by player 1 by making these expressions all equal, which means $p = q$ and then $p = 1 - p + p - p^2 = 1 - p^2$, with solution $p = (\sqrt{5} - 1)/2 \approx 0.618\ldots$. This is less the max-min value of 2/3 when player 1 is allowed to play a mixed strategy.

An interesting question that arose exactly from this example is the following: Suppose player 1 plays, in a zero-sum game with possibly imperfect recall, a max-min behavior strategy, like with p and q here. That is, each information set is controlled by an independent agent that acts on behalf of player 1. Is there a strategy of player 2 (who has perfect recall) that extends this to an *equilibrium* of this "agent form" of the game, with no individual agent having an incentive to deviate unilaterally? The answer is yes, as shown in B. von Stengel and D. Koller (1997), Team-maxmin equilibria, Games and Economic Behavior 21, 309–321.

(e) If a zero-sum game with imperfect recall could be solved in polynomial time, in the sense of finding its value, then we could solve the NP-complete satisfiability problem in polynomial time and thus have P = NP. This is unlikely.

Solution to Exercise 9

We use the inequalities stated in the question, where the strict inequalities hold because exactly m many equalities hold for y and z, respectively.

$$\begin{array}{rclcrcl} c^\top y &=& \beta_0 & \qquad & c^\top z &<& \beta_0 \\ d^\top y &<& \beta_1 & & d^\top z &=& \beta_1 \\ a_2^\top y &=& \beta_2 & & a_2^\top z &=& \beta_2 \\ &\vdots& & & &\vdots& \\ a_m^\top y &=& \beta_m & & a_m^\top z &=& \beta_m \end{array}$$

Let $(\gamma, \delta, \alpha_2, \ldots, \alpha_m) \neq (0, 0, 0, \ldots, 0)$ with

$$\gamma c^\top + \delta d^\top + \alpha_2 a_2 +^\top \cdots + \alpha_m a_m^\top = \mathbf{0}$$

where, moreover, $\gamma \neq 0$ and $\delta \neq 0$ by the linear independence for the facet vectors of y and z. We now multiply the above sum once with y and z, respectively, and compare, which gives $(\gamma c^\top + \delta d^\top)y = (\gamma c^\top + \delta d^\top)z$ and thus

$$\gamma(c^\top y - c^\top z) = \delta(d^\top z - d^\top y)\,.$$

Here, $c^\top y - c^\top z = \beta_0 - c^\top z > 0$ and $d^\top z - c^\top y = \beta_1 - d^\top y > 0$ and hence γ and δ have the same sign. Using the bilinearity of determinants we have

$$|(\gamma c + \delta d)\; a_2 \cdots a_m| = \gamma\; |c\; a_2 \cdots a_m| + \delta\; |d\; a_2 \cdots a_m| = 0\,.$$

Hence, $|c\; a_2 \cdots a_m|$ and $|d\; a_2 \cdots a_m|$ have opposite sign.

Solution to Exercise 10

If $Y = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and $x^\top = (x_1, x_2)$, then the stationarity condition $x^\top Y = x^\top$ can be written as

$$\begin{array}{rcrcl} (1-b)x_1 &+& cx_2 &=& x_1 \\ bx_1 &+& (1-c)x_2 &=& x_2 \end{array}$$

using $a = 1 - b$ and $d = 1 - c$. Both equations are equivalent to $cx_2 = bx_1$ (note that Y does not have full rank because $Y\mathbf{1} = \mathbf{1}$, so the two equations are linearly dependent). If $c = b = 0$, then *any* probability distribution x is stationary (which is always the case if Y is the identity matrix), and x is not unique. Otherwise, $c + b > 0$ and x is unique, because with $x_1 = 1 - x_2$ we can rewrite $cx_2 = bx_1$ as $x_2 = \frac{b}{c+b}$, which implies $x^\top = (\frac{c}{c+b}, \frac{b}{c+b})$, as claimed.