IE 5545: Practice Problems (Solution Key)

Prof. Krishnamurthy Iyer

Spring 2026

1. (Tadelis 8.10) Consider the following dynamic game: Player 1 can choose to play it safe (S), in which case both he and player 2 get a payoff of 3 each, or he can risk (R) playing a game with player 2. If he chooses R, the two players play the following simultaneous-move game:

	A	B
C	(8,0)	(0,2)
D	(6,6)	(2,2)

   (a) Draw a game tree that represents this game. How many proper subgames does it have?

   Solution: An extensive form that represents the game is shown in Figure 1: At the root, player 1 chooses S or R. If S, the game ends with payoff $(3,3)$. If R, players 1 and 2 simultaneously choose $\{C, D\}$ and $\{A, B\}$ and receive payoffs from the matrix.

   There is **one** proper subgame: the subgame starting after R is chosen, which is the simultaneous-move game with payoff matrix as shown above. □

   (b) Are there any other game trees that also represent this game? Explain why or why not.

   Solution: Yes, an alternative, but equivalent, representation is shown in Figure 2. This representation differs from the first one in how it represents the simultaneous-move game after player 1's chooses R. □

   (c) Provide the strategic form of the dynamic game.

   Solution: Player 1 must specify what he would do at the root and what he would do in the R-subgame. Thus player 1 has four pure strategies:

$$\mathcal{S}_1 = \{(S,C),(S,D),(R,C),(R,D)\},$$

   where the first component is player 1's action at the root, and the second component is his action in the simultaneous-move game after R is played. Player 2 acts only in the R-subgame, so

$$\mathcal{S}_2 = \{A, B\}.$$

   The payoff is given by the following matrix:

	A	B
(S,C)	$(3,3)$	$(3,3)$
(S,D)	$(3,3)$	$(3,3)$
(R,C)	$(8,0)$	$(0,2)$
(R,D)	$(6,6)$	$(2,2)$

□

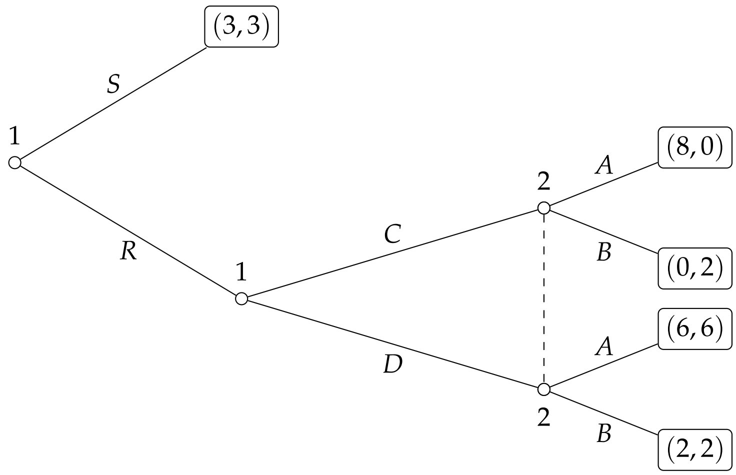

Figure 1: Extensive form for the two-player game in Problem 1.

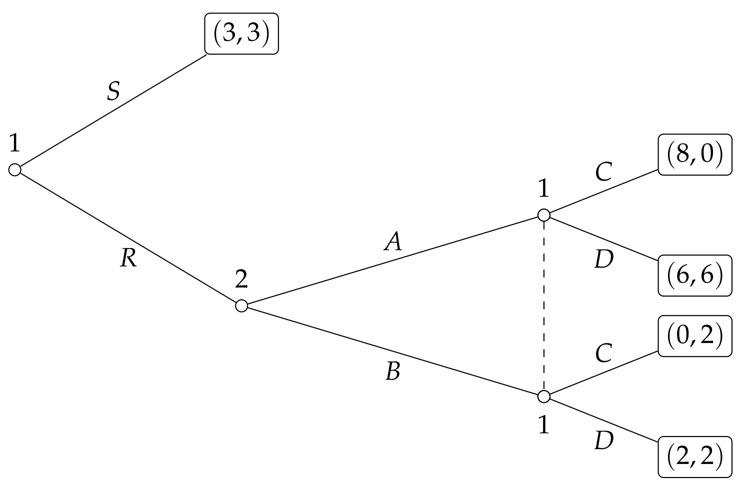

Figure 2: Alternative extensive form for the game in Problem 1.

(d) Find all Nash and subgame-perfect Nash equilibria of the dynamic game.

Solution:

Equilibria in the R-subgame. In the 2×2 game: player 1's best response is C to A and D to B. Player 2's best response is B to C and A to D. Hence there is no pure Nash equilibrium. We can also conclude that there is no mixed Nash equilibrium where only one player mixes. Thus, we must look for a fully mixed Nash equilibrium, which can be obtained by indifference of players payoffs:

$$8q = 2 + 4q \Rightarrow q = \frac{1}{2}, \qquad 6 - 6p = 2 \Rightarrow p = \frac{2}{3},$$

where $p \in (0,1)$ is prob. player 1 plays C and $q \in (0,1)$ is prob. player 2 plays A The equilibrium payoffs in the R-subgame are

$$\pi_1 = 8q = 4, \qquad \pi_2 = 2.$$

SPNE of the dynamic game. In any SPNE, play in the R-subgame must be a Nash equilibrium of that subgame, so the continuation after R is the mixed equilibrium above. At the root, player 1 compares S (payoff 3) to R (payoff 4) and chooses R. Therefore the (unique) SPNE outcome has player 1 choose R, and then the players play the mixed NE in the following subgame. The strategies are given by

$$\begin{aligned} &\text{player 1:} \quad && \left(R, \left[\frac{2}{3}, \frac{1}{3}\right]^\top\right) \\ &\text{player 2:} \quad && \left[\frac{1}{2}, \frac{1}{2}\right]^\top \end{aligned}$$

Nash equilibria.

Recall the strategic form:

	A	B
(S,C)	$(3,3)$	$(3,3)$
(S,D)	$(3,3)$	$(3,3)$
(R,C)	$(8,0)$	$(0,2)$
(R,D)	$(6,6)$	$(2,2)$

From the payoff matrix, we can see that there are also two pure-strategy Nash equilibria in which player 1 chooses (S,C) or (S,D) and player 2 chooses B (since B makes player 1 prefer S over R):

$$((S,C),B) \quad \text{and} \quad ((S,D),B).$$

These are Nash equilibria but not subgame perfect (because the prescribed play is not a Nash equilibrium of the R-subgame). These equilibria involve player 2 making the non-credible threat of playing B in the subgame after R.

There are also additional mixed Nash equilibria where player 1 mixes between (S,C) and (S,D) whereas player 2 mixes between A and B (with appropriate probabilities) with the resulting payoff also equal to $(3,3)$. $\square$

2. Consider a zero sum game with the following payoff matrix for player I:

$$A = \begin{bmatrix} -1 & 7 & -4 & 2 \\ 3 & -3 & 5 & -1 \end{bmatrix}.$$

(a) Restricting attention to pure strategies, compute the maximin strategy and the maximin value for both players I and II.

Solution:

Player I (row player). Row minima:

$$\min(-1, 7, -4, 2) = -4, \qquad \min(3, -3, 5, -1) = -3.$$

Thus player I's pure maximin value is $m_p = \max\{-4, -3\} = -3$, achieved by playing row 2.

Player II (column player). In a zero-sum game, player II's maximin value corresponds to choosing a column that minimizes player I's best response: compute column maxima:

$$\max(-1, 3) = 3, \quad \max(7, -3) = 7, \quad \max(-4, 5) = 5, \quad \max(2, -1) = 2.$$

So $M_p = \min\{3, 7, 5, 2\} = 2$, achieved by playing column 4. This guarantees player I's payoff is at most 2, i.e., player II's payoff is at least -2. □

(b) Write down an optimization problem in p whose optimal value is the maximin value of player I, and whose solution p^* corresponds to her maximin strategy.

Solution: Let player I mix as $\sigma_1 = [p \;\; 1-p]^\top$ with $p \in [0, 1]$. Against each column j, player I's expected payoff is:

$$\pi_1(p, 1) = 3 - 4p, \quad \pi_1(p, 2) = -3 + 10p, \quad \pi_1(p, 3) = 5 - 9p, \quad \pi_2(p, 4) = -1 + 3p.$$

Player I's maximin problem is

$$\max_{p \in [0,1]} \min\{\pi_1(p, 1), \pi_1(p, 2), \pi_1(p, 3), \pi_1(p, 4)\}$$

This is equivalent to

$$\max_{p \in [0,1]} \min\{3 - 4p, -3 + 10p, 5 - 9p, -1 + 3p\}.$$

□

(c) Solve the preceding problem and find p^* and the maximin value for player I.

Solution: The easiest way is to plot the four lines (in p) inside the minimization in the optimization problem above, and taking the lower envelope of the four lines. This lower envelope will be the plot of the interior minimization. Then we compute the maximum of this function. Upon doing so, we can see that the maximum is attained at the point where $\pi_1(p, 3)$ and $\pi_1(p, 4)$ intersect:

$$5 - 9p = -1 + 3p \implies 12p = 6 \implies p^* = \frac{1}{2}.$$

The maximin value for player I is

$$m = \min\{\pi_1(p^*, 1), \pi_1(p^*, 2), \pi_1(p^*, 3), \pi_1(p^*, 4)\} = \min\{1, 2, \tfrac{1}{2}, \tfrac{1}{2}\} = \frac{1}{2}.$$

So player I's maximin strategy is $\sigma_1^* = [\frac{1}{2} \;\; \frac{1}{2}]^\top$ and her maximin value is $\frac{1}{2}$. □

(d) Show that the maximin strategy for player II is of the form $\sigma_2 = [0 \ \ 0 \ \ q \ \ 1-q]^\mathsf{T}$ for some $q \in [0,1]$.

Solution: At player I's maximin strategy $p^* = \frac{1}{2}$, the expected payoffs to player I for each action of player II are

$$\pi_1(\tfrac{1}{2}, 1) = 1, \quad \pi_1(\tfrac{1}{2}, 2) = 2, \quad \pi_1(\tfrac{1}{2}, 3) = \tfrac{1}{2}, \quad \pi_1(\tfrac{1}{2}, 4) = \tfrac{1}{2}.$$

Since $m = \frac{1}{2}$, we see that the constraints corresponding to strategies 1 and 2 in the maximin LP hold with strict inequality:

$$\begin{aligned} 3 - 4p^* &> \frac{1}{2} \\ -3 + 10p^* &> \frac{1}{2}. \end{aligned}$$

Hence, by *complementary slackness*, the corresponding optimal variables for the dual minimax LP must be zero. Thus $q_1^* = q_2^* = 0$. Hence, any optimal solution of minimax LP (and hence the maximin strategy of player II) must be of the form $[0 \ \ 0 \ \ q \ \ 1-q]^\top$ for some $q \in [0,1]$.

Alternative solution: Since we have a zero-sum game, we have

$$\pi_2(\tfrac{1}{2}, 1) = -1, \quad \pi_2(\tfrac{1}{2}, 2) = -2, \quad \pi_2(\tfrac{1}{2}, 3) = -\tfrac{1}{2}, \quad \pi_2(\tfrac{1}{2}, 4) = -\tfrac{1}{2}.$$

Thus, any best response for player II against p^* should only mix between strategies 3 and 4 (since both those strategies attain the highest payoff for player II. Recall that both players maximin strategies form a Nash equilibrium, and hence the maximin strategy for player II is a best response to the maximin strategy for player I. Hence, we conclude that the maximin strategy for player II must be of the form $[0 \ \ 0 \ \ q \ \ 1-q]^\top$ for some $q \in [0,1]$. □

3. Recall the infinitely repeated Prisoner's Dilemma $G(\infty, \delta)$, with discount factor $\delta < 1$, where G is given by the following strategic form:

	S	B
S	$-1, -1$	$-10, 0$
B	$0, -10$	$-5, -5$

The two players play G repeatedly at times $t = 0, 1, 2, \ldots$. The payoff of player i is given by

$$\pi_i(s_i, s_{-i}) = \sum_{t=0}^{\infty} \delta^t \pi_i^G(a_i^t, a_{-i}^t).$$

Consider the following modification of the grim-trigger strategy for player 1. Player 1 starts by playing S (stay silent) and continues doing so until player 2 first plays B (betray). Thereafter, player 1 continues playing B, until the first time player 2 plays S, at which point player 1 forgives player 2 and starts playing S (until the next betrayal by player 2 when the same cycle of punishment and forgiveness continues). We will call such a player 1 a *forgiver*. We are interested in the case where both players are forgivers and ask when such strategies form an SPNE.

In order to define the strategy formally, we need to define a (common) state of the game X_t at time t to denote whether the game is in cooperation mode ($X_t = 0$), or whether player 1 is being punished ($X_t = 1$) or whether player 2 is being punished ($X_t = 2$). The state evolves as follows: First, the game starts in cooperation mode: $X_0 = 0$. For any $t \geq 1$, we have

$$X_{t+1} = \begin{cases} 0, & \text{if } X_t = 0, \text{ and } (S,S) \text{ or } (B,B) \text{ was played at time } t; \\ 1, & \text{if } X_t = 0, \text{ and } (B,S) \text{ was played at time } t; \\ 2, & \text{if } X_t = 0, \text{ and } (S,B) \text{ was played at time } t; \\ & \\ 0, & \text{if } X_t = 1, \text{ and } (S,B) \text{ was played at time } t; \\ 1, & \text{if } X_t = 1, \text{ and } (S,B) \text{ was } \textbf{not played} \text{ at time } t; \\ & \\ 0, & \text{if } X_t = 2, \text{ and } (B,S) \text{ was played at time } t; \\ 2, & \text{if } X_t = 2, \text{ and } (B,S) \text{ was } \textbf{not played} \text{ at time } t. \end{cases}$$

(In words, if the game is in cooperation mode, and (S,S) or (B,B) is played, then the game stays in cooperation mode. Otherwise, if (B,S) is played, then game transitions to punishing player 1 or if (S,B) is played, the game transitions to punishing player 2. If player 1 is being punished, that remains the case until (S,B) is played, whereupon the game transitions to cooperation mode. Similarly, if player 2 is being punished, that remains the case until (B,S) is played, whereupon the game transitions to cooperation mode.)

Define $\mathsf{forgiver}_1$ to be the following strategy for player 1: play S if $X_t = 0$ or $X_t = 1$, and play B if $X_t = 2$.

Similarly, define $\mathsf{forgiver}_2$ to be the following strategy for player 2: play S if $X_t = 0$ or $X_t = 2$, and play B if $X_t = 1$.

(a) State the path of play (the sequence of action profiles at each time) under the strategy profile $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$. What is player 1's expected payoff under this strategy profile? (Your answer will be in terms of δ.)

Solution: In Figure 3, we show how the state X_t changes depending on which action profile is played. (For instance, if the state is 1, and action profile (S,B) is played, then the state changes to 0.) Using this state diagram, we can easily identify the sequence of plays for any strategy profile (or one-step deviations).

Since $X_0 = 0$, both players play S under $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$. After (S,S) the state remains 0, so the play in the next period is also (S,S). From this, we conclude that the path of play is (S,S) forever. Player 1's payoff is

$$\sum_{t=0}^{\infty} \delta^t(-1) = -\frac{1}{1-\delta}.$$

□

(b) Suppose the state at time t is given by $X_t = 0$. What is the path of play under the one-step deviation for player 1? What is player 1's expected payoff under this deviation in this subgame?

Solution: At $X_t = 0$, the $\mathsf{forgiver}_1$ strategy prescribes player 1 plays S. A one-step deviation is to play B at time t.

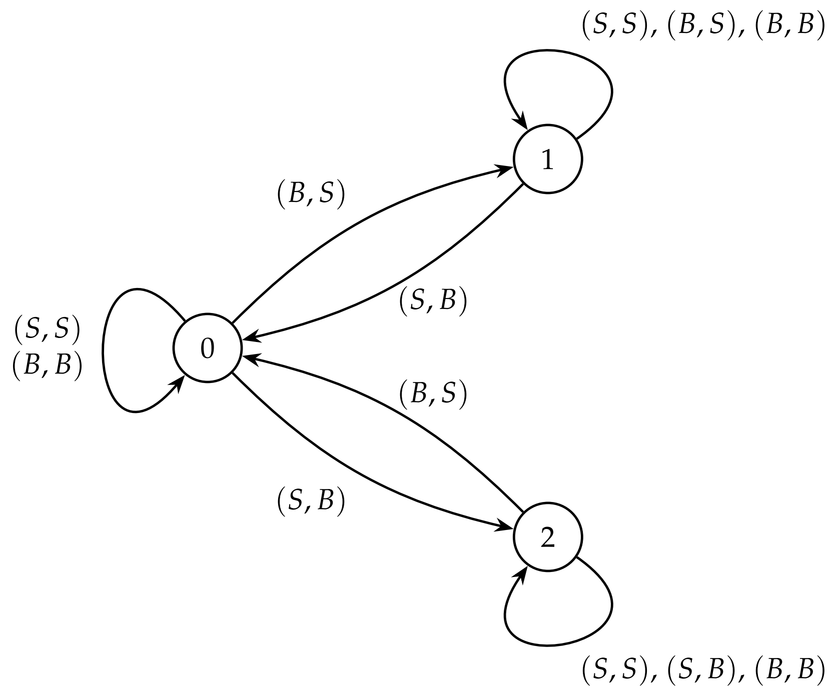

Figure 3: State diagram for the $\mathsf{forgiver}$ strategies.

At time t under one-step deviation for player 1, (B,S) is played, so player 1 gets 0 and the state becomes $X_{t+1} = 1$. At time $t+1$ in state 1, $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ plays (S,B), so player 1 gets -10 and the state returns to 0. Thereafter play is (S,S) forever. In other words, the path of play under one-step deviation is $(B,S), (S,B), (S,S), (S,S), \ldots$. Thus player 1's deviation payoff (from time t onward) is

$$0 + \delta(-10) + \delta^2(-1) + \delta^3(-1) + \cdots = 0 + \delta(-10) + \delta^2\left(-\frac{1}{1-\delta}\right).$$

□

(c) Suppose the state at time t is given by $X_t = 1$. What is the path of play under the strategy profile $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ and under the one-step deviation for player 1? What is player 1's expected payoff for each case in this subgame?

Solution: If $X_t = 1$, then $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ plays (S,B) at time t, giving player 1 payoff -10 and returning the state to 0. Thus, the path of play is $(S,B), (S,S), (S,S), \ldots$. Player 1's payoff is

$$V_1^{\text{follow}} = -10 + \delta\left(-\frac{1}{1-\delta}\right).$$

A one-step deviation for player 1 is to play B at time t (while player 2 plays B in state 1). Then at time t the action is (B,B), giving player 1 payoff -5, and the state remains 1 (since (S,B) was not played). At time $t+1$, play reverts to (S,B) (follow the strategy thereafter), giving -10 and returning to 0, and then (S,S) forever. In other words, the path of play is $(B,B), (S,B), (S,S), (S,S), \ldots$. Player 1's payoff is

$$V_1^{\text{dev}} = -5 + \delta(-10) + \delta^2\left(-\frac{1}{1-\delta}\right).$$

□

(d) Suppose the state at time t is given by $X_t = 2$. What is the path of play under the strategy profile $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ and under the one-step deviation for player 1? What is player 1's expected payoff for each case in this subgame?

Solution: If $X_t = 2$, then $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ plays (B, S) at time t, so player 1 gets 0 and the state returns to 0. Thus, path of play is $(B, S), (S, S), (S, S), \ldots$. Player 1's payoff is

$$V_2^{\text{follow}} = 0 + \delta\left(-\frac{1}{1-\delta}\right).$$

A one-step deviation is to play S at time t (while player 2 plays S in state 2). Then (S, S) is played, giving player 1 payoff -1, and the state remains 2 (since (B, S) was not played). At time $t+1$, follow the strategy again: (B, S) is played, giving payoff 0 and returning to 0, then (S, S) forever. Thus, the path of play is $(S, S), (B, S), (S, S), (S, S), \ldots$. Player 1's payoff is

$$V_2^{\text{dev}} = -1 + \delta(0) + \delta^2\left(-\frac{1}{1-\delta}\right).$$

□

(e) For what values of $\delta \in (0, 1)$ is the strategy profile $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ an SPNE of $G(\infty, \delta)$? (Be sure to state the *principle* you are using to conclude the result.)

Solution:

Principle: Use the *one-step deviation principle* for infinitely repeated games: a strategy profile is SPNE if and only if no player can profitably deviate in any subgame by changing her action at a single history and then returning to the prescribed strategy.

We check player 1 (player 2 is symmetric).

State $X_t = 0$: No profitable deviation requires

$$-\frac{1}{1-\delta} \geq -10\delta - \frac{\delta^2}{1-\delta} \iff 9\delta^2 - 10\delta + 1 \leq 0 \iff \delta \geq \frac{1}{9}.$$

State $X_t = 1$: No profitable deviation requires $V_1^{\text{follow}} \geq V_1^{\text{dev}}$, which simplifies to

$$9\delta - 5 \geq 0 \iff \delta \geq \frac{5}{9}.$$

State $X_t = 2$: $V_2^{\text{follow}} \geq V_2^{\text{dev}}$ holds for all $\delta \in (0, 1)$.

Therefore $(\mathsf{forgiver}_1, \mathsf{forgiver}_2)$ is an SPNE if and only if

$$\delta \in \left[\frac{5}{9}, 1\right).$$

□

4. Consider a two player static game of incomplete information, where each player $i \in \{1, 2\}$ has a type $\theta_i \in [0, 1]$ that is drawn independently and identically from a uniform distribution over $[0, 1]$. Each player i, after learning her type, chooses an action $x_i \geq 0$. The payoff function of player i is given by

$$\pi_i(x_i, x_j, \theta_i, \theta_j) = \theta_i x_i x_j - x_i^3,$$

where $j \neq i$. Your goal is to find a Bayes-Nash equilibrium for this game.

(a) Give a formal description of a pure strategy for player i.

Solution: A pure strategy for player i is a function

$$s_i : [0,1] \to [0,\infty),$$

mapping each type θ_i to an action $x_i = s_i(\theta_i)$. $\square$

(b) Suppose player $j \neq i$ is playing a fixed pure strategy s_j. What is the *interim* expected payoff $\Phi_i(x_i|\theta_i)$ of player i when her type is θ_i and her action is x_i?

Solution: Fix θ_i and choose action $x_i \geq 0$. Since θ_j is independent of θ_i and uniform,

$$\Phi_i(x_i \mid \theta_i) = \mathbf{E}_{\theta_j}\left[\theta_i x_i s_j(\theta_j) - x_i^3\right] = \theta_i x_i \mathbf{E}[s_j(\theta_j)] - x_i^3.$$

Let $m_j \triangleq \mathbf{E}[s_j(\theta_j)]$. Then $\Phi_i(x_i \mid \theta_i) = \theta_i m_j x_i - x_i^3$. $\square$

(c) State what conditions Φ_1 and Φ_2 must satisfy in a Bayes-Nash equilibrium.

Solution: A profile (s_1, s_2) is a Bayes-Nash equilibrium if and only if for each player i and each type $\theta_i \in [0,1]$,

$$\Phi_i(s_i(\theta_i) \mid \theta_i) \geq \Phi_i(x_i \mid \theta_i), \quad \text{for all } x_i \geq 0.$$

where Φ_i is computed given the other player's strategy. $\square$

(d) Suppose player $j \neq i$ is playing a fixed pure strategy s_j. What is the best response strategy for bidder i?

Solution: Let $m_j = \mathbf{E}[s_j(\theta_j)]$. For a given type θ_i, player i's interim payoff for choosing x_i is given by

$$\Phi_i(x_i \mid \theta_i) = \theta_i m_j x_i - x_i^3.$$

To find the optimal action, player i solves

$$\max_{x_i \geq 0} \theta_i m_j x_i - x_i^3.$$

The objective is strictly concave in x_i (since $-x_i^3$ dominates), so the first-order necessary condition gives the unique maximizer:

$$\frac{\partial \Phi_i(x_i \mid \theta_i)}{\partial x_i} = \theta_i m_j - 3x_i^2 = 0 \quad \Longrightarrow \quad x_i^*(\theta_i) = \sqrt{\frac{\theta_i m_j}{3}}.$$

Thus the best response strategy is $BR_i[s_j](\theta_i) = \sqrt{\theta_i m_j / 3}$. $\square$

(e) Find a Bayes-Nash equilibrium for this game.

Solution: We look for a symmetric BNE with $s_1 = s_2 = s$ and $m = \mathbf{E}[s(\theta)]$ where $\theta \sim \text{Unif}[0,1]$.

From the best response formula, a symmetric equilibrium must satisfy

$$s(\theta) = \sqrt{\frac{\theta m}{3}}.$$

Taking expectations over θ gives

$$m = \mathbf{E}[s(\theta)] = \sqrt{\frac{m}{3}}\,\mathbf{E}[\sqrt{\theta}] = \sqrt{\frac{m}{3}} \int_0^1 \sqrt{\theta}\, d\theta = \sqrt{\frac{m}{3}} \cdot \frac{2}{3}.$$

If $m > 0$, divide by $\sqrt{m}$ to obtain

$$\sqrt{m} = \frac{2}{3\sqrt{3}} \quad \Longrightarrow \quad m = \frac{4}{27}.$$

Substitute back:

$$s(\theta) = \sqrt{\frac{\theta}{3} \cdot \frac{4}{27}} = \sqrt{\frac{4\theta}{81}} = \frac{2}{9}\sqrt{\theta}.$$

Therefore a (nontrivial) symmetric Bayes-Nash equilibrium is

$$s_1(\theta) = s_2(\theta) = \frac{2}{9}\sqrt{\theta}.$$

Another solution corresponds to $m = 0$, yielding a trivial Bayes-Nash equilibrium $s_1(\theta) \equiv 0$ and $s_2(\theta) \equiv 0$. □