

Homework 1

Student: [*** Your Andrew ID here ***]

Due on: Sep. 30, 2024, beginning of class

Instructions Please submit your homework on Gradescope. The Gradescope entry code can be found on Piazza. On Gradescope, please submit a single pdf file containing **both** your written solutions **and** your code (`stub.py`) attached to the end (for example, using a `verbatim` environment). You can (but are not required to) typeset your written solutions in the `.tex` file provided in the homework `.zip`.

1 A Regret-Based Proof of the Minmax Theorem (20 points)

Let $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ be convex and compact sets, and $\mathbf{A} \in \mathbb{R}^{n \times m}$. The minmax theorem asserts that

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

In this problem you will show that the fact that (external) regret minimization algorithms for \mathcal{X} exist is a powerful enough statement to imply the minmax theorem.

One direction (called *weak duality*) of the proof is easy and very general. Specifically, show the following.

Problem 1.1 (5 points). Show that

$$\max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \mathbf{x}^\top \mathbf{A} \mathbf{y} \leq \min_{\mathbf{y} \in \mathcal{Y}} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \mathbf{A} \mathbf{y}. \quad (1)$$

Solution. [*** Your solution here ***] □

To show the reverse inequality, we will interpret the bilinear saddle point $\min_{\mathbf{y} \in \mathcal{Y}} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \mathbf{A} \mathbf{y}$ as a repeated game. At each time t , we will let a regret minimizer $\mathcal{R}_{\mathcal{X}}$ pick actions $\mathbf{x}^t \in \mathcal{X}$, whereas we will always assume that $\mathbf{y}^t \in \mathcal{Y}$ is chosen by the environment to best respond to \mathbf{x}^t , that is,

$$\mathbf{y}^t \in \arg \min_{\mathbf{y} \in \mathcal{Y}} (\mathbf{x}^t)^\top \mathbf{A} \mathbf{y}.$$

The utility function observed by $\mathcal{R}_{\mathcal{X}}$ at each time t is set to

$$\ell_{\mathcal{X}}^t : \mathbf{x} \mapsto \mathbf{x}^\top \mathbf{A} \mathbf{y}^t = (\mathbf{A} \mathbf{y}^t)^\top \mathbf{x}.$$

By definition of regret minimizer, we will assume that $\mathcal{R}_{\mathcal{X}}$ guarantees *sublinear regret* in the worst case.

Let $\bar{\mathbf{x}}^T \in \mathcal{X}$ and $\bar{\mathbf{y}}^T \in \mathcal{Y}$ be the average strategies output up to time T , that is,

$$\bar{\mathbf{x}}^T := \frac{1}{T} \sum_{t=1}^T \mathbf{x}^t \quad \bar{\mathbf{y}}^T := \frac{1}{T} \sum_{t=1}^T \mathbf{y}^t.$$

Problem 1.2 (5 points). Argue that at all t ,

$$\max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \mathbf{x}^\top \mathbf{A} \mathbf{y} \geq \frac{1}{T} \min_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T (\mathbf{x}^t)^\top \mathbf{A} \mathbf{y} \geq \frac{1}{T} \sum_{t=1}^T (\mathbf{x}^t)^\top \mathbf{A} \mathbf{y}^t. \quad (2)$$

★ Hint: Substitute $\bar{\mathbf{x}}^T$ for \mathbf{x} . Then, use the information you have on \mathbf{y}^t .

Solution. ***** Your solution here ***** □

Problem 1.3 (5 points). Let $R_{\mathcal{X}}^T$ be the regret accumulated by $\mathcal{R}_{\mathcal{X}}$ up to time T . Using (2), argue that at all times T

$$\max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \mathbf{x}^\top \mathbf{A} \mathbf{y} \geq \min_{\mathbf{y} \in \mathcal{Y}} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \mathbf{A} \mathbf{y} - \frac{R_{\mathcal{X}}^T}{T}. \quad (3)$$

★ Hint: use the definition of regret $R_{\mathcal{X}}^T$.

Solution. ***** Your solution here ***** □

Problem 1.4 (5 points). Use (3) and (1) and to conclude that

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

Solution. ***** Your solution here ***** □

2 (Coarse-)correlated equilibria and dominated actions (30 points)

Recall the following definitions from lecture.

- An n -player finite normal-form game is defined by n finite strategy sets A_1, \dots, A_n and n utility functions $u_i : A \rightarrow \mathbb{R}$, one per player, where $A := A_1 \times \dots \times A_n$ is the set of pure strategy profiles.
- A *coarse-correlated equilibrium* (CCE) is a distribution $\pi \in \Delta(A)$ such that, for every player i and every action $a'_i \in A_i$, we have

$$\mathbb{E}_{a \sim \pi} [u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i})] \leq 0.$$

- A *correlated equilibrium* (CE) is a distribution $\pi \in \Delta(A)$ such that, for every player i and every function $\phi : A_i \rightarrow A_i$, we have

$$\mathbb{E}_{a \sim \pi} [u_i(\phi(a_i), a_{-i}) - u_i(a_i, a_{-i})] \leq 0.$$

- An action $a_i \in A_i$ *strictly dominates* another action $a'_i \in A_i$ if $u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i})$ for every possible opponent action profile a_{-i} . In this case we call a'_i *strictly dominated*. If a_i strictly dominates every other action, then we call it *strictly dominant*.

Problem 2.1 (5 points). Prove that, if an action $\hat{a}_i \in A_i$ is **strictly dominant**, then Player i plays \hat{a}_i with probability 1 in every **CCE**.

Solution. ***** Your solution here *****

□

Problem 2.2 (10 points). Prove that, if an action $\hat{a}_i \in A_i$ is **strictly dominated** and the game is **two-player zero-sum**, then Player i plays \hat{a}_i with probability 0 in every **CCE**.

★ Hint: For any mixed strategy $x_1 \in \Delta(A_1)$, define the best-response value as $v^*(x_1) := \min_{a_2 \in A_2} u_1(x_1, a_2)$. What can you say about $v^*(x_1)$ if x_1 plays a dominated action with positive probability?

Solution. ***** Your solution here *****

□

Problem 2.3 (5 points). Prove that, if an action $\hat{a}_i \in A_i$ is **strictly dominated**, then Player i plays \hat{a}_i with probability 0 in every **CE**.

Solution. ***** Your solution here *****

□

Problem 2.4 (10 points). Find an explicit example of a normal-form game, and a CCE π of that normal-form game, in which Player 1 plays a **strictly dominated** action with strictly positive probability.

★ Hint: By the previous three parts:

1. P1 cannot have a strictly dominant action (and thus P1 must have at least three actions),
2. your game cannot be a two-player zero-sum game, and
3. your CCE cannot also be a CE.

Solution. ***** Your solution here *****

□

3 Linear programming for solving extensive-form zero-sum games, and application to low randomization in poker (50 points)

In many games, the optimal Nash equilibrium requires that all players randomize their moves. As an example, consider rock-paper-scissors: any deterministic strategy (for example, always playing rock) is heavily suboptimal. In this problem, you will quantify how much value is lost by insisting on playing deterministic strategies in three games: rock-paper-superscissors and two well-known poker variants—Kuhn poker [Kuhn, 1950] and Leduc poker [Southey et al., 2005]. A description of each game is given in the zip of this homework. The description is specified in Section 3.1. The zip of the homework also contains a stub Python file to help you set up your implementation.

3.1 Format of the game files

Each game is encoded as a json file with the following structure.

- At the root, we have a dictionary with three keys: `decision_problem_p11`, `decision_problem_p12`, and `utility_p11`. The first two keys contain a description of the tree-form sequential decision problems faced by the two players, while the third is a description of the bilinear utility function for Player 1 as a function of the sequence-form strategies of each player. Since both games are zero-sum, the utility for Player 2 is the opposite of the utility of Player 1.
- The tree of decision points and observation points for each decision problem is stored as a list of nodes. Each node has the following fields

id is a string that represents the identifier of the node. The identifier is unique among the nodes for the same player.

type is a string with value either `decision` (for decision points) or `observation` (for observation points).

actions (only for decision points). This is a set of strings, representing the actions available at the decision node.

signals (only for observation points). This is a set of strings, representing the signals that can be observed at the observation node.

parent_edge identifies the parent edge of the node. If the node is the root of the tree, then it is `null`. Else, it is a pair (`parent_node_id`, `action_or_signal`), where the first member is the `id` of the parent node, and `action_or_signal` is the action or signal that connects the node to its parent.

parent_sequence (only for decision points). Identifies the parent sequence p_j of the decision point, defined as the last sequence (that is, decision point-action pair) encountered on the path from the root of the decision process to j .

Remark 1. The list of nodes of the tree-form sequential decision process is given in top-down traversal order. The bottom-up traversal order can be obtained by reading the list of nodes backwards.

- The bilinear utility function for Player 1 is given through the payoff matrix \mathbf{A} such that the (expected) utility of Player 1 can be written as

$$u_1(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top \mathbf{A} \mathbf{y},$$

where \mathbf{x} and \mathbf{y} are sequence-form strategies for Players 1 and 2 respectively. We represent \mathbf{A} in the file as a list of all non-zero matrix entries, storing for each the row index, column index, and value. Specifically, each entry is an object with the fields

sequence_p11 is a pair (`decision_pt_id_p11`, `action_p11`) which represents the sequence of Player 1 (row of the entry in the matrix).

sequence_p12 is a pair (`decision_pt_id_p12`, `action_p12`) which represents the sequence of Player 2 (column of the entry in the matrix).

value is the non-zero float value of the matrix entry.

Example: Rock-paper-superscissors In the case of rock-paper-superscissors the decision problem faced by each of the players has only one decision points with three actions: playing rock, paper, or superscissors. So, each tree-form sequential decision process only has a single node, which is a decision node. The payoff matrix of the game is

$$\begin{matrix} & \text{r} & \text{p} & \text{s} \\ \text{r} & \begin{pmatrix} 0 & -1 & 1 \end{pmatrix} \\ \text{p} & \begin{pmatrix} 1 & 0 & -2 \end{pmatrix} \\ \text{s} & \begin{pmatrix} -1 & 2 & 0 \end{pmatrix} \end{matrix}$$

So, the game file in this case has content:

```
{
  "decision_problem_pl1": [
    {"id": "d1_pl1", "type": "decision", "actions": ["r", "p", "s"],
      "parent_edge": null, "parent_sequence": null}
  ],
  "decision_problem_pl2": [
    {"id": "d1_pl2", "type": "decision", "actions": ["r", "p", "s"],
      "parent_edge": null, "parent_sequence": null}
  ],
  "utility_pl1": [
    {"sequence_pl1": ["d1_pl1", "r"], "sequence_pl2": ["d1_pl2", "p"], "value": -1},
    {"sequence_pl1": ["d1_pl1", "r"], "sequence_pl2": ["d1_pl2", "s"], "value": 1},
    {"sequence_pl1": ["d1_pl1", "p"], "sequence_pl2": ["d1_pl2", "r"], "value": 1},
    {"sequence_pl1": ["d1_pl1", "p"], "sequence_pl2": ["d1_pl2", "s"], "value": -2},
    {"sequence_pl1": ["d1_pl1", "s"], "sequence_pl2": ["d1_pl2", "r"], "value": -1},
    {"sequence_pl1": ["d1_pl1", "s"], "sequence_pl2": ["d1_pl2", "p"], "value": 2}
  ]
}
```

3.2 Computing the value of the game (15 points)

As a warmup, you will implement the linear program formulation of Nash equilibrium strategies seen in Lecture 10 using the commercial solver Gurobi (<https://www.gurobi.com/>). Gurobi is a powerful commercial solver for linear and non-linear optimization problems. You can download the solver and request a free license for academic use from their website.

Installing gurobipy Installation instructions for Gurobi’s python bindings are available on the Gurobi website, [here](#).¹

Linear programming formulation of Nash equilibrium strategies For your convenience, here are again the linear programs—for Player 1 and Player 2, respectively—that you need to implement:

$$\mathcal{P}_1 : \begin{cases} \max & \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ & \textcircled{3} \mathbf{x} \geq \mathbf{0}, \mathbf{v} \text{ free,} \end{cases} \quad \mathcal{P}_2 : \begin{cases} \max & \mathbf{f}_1^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} -\mathbf{A} \mathbf{y} - \mathbf{F}_1^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ & \textcircled{3} \mathbf{y} \geq \mathbf{0}, \mathbf{v} \text{ free,} \end{cases} \quad (4)$$

where $\{\mathbf{x} \in \mathbb{R}^{|\Sigma_1|} : \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1, \mathbf{x} \geq \mathbf{0}\}$ and $\{\mathbf{y} \in \mathbb{R}^{|\Sigma_2|} : \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2, \mathbf{y} \geq \mathbf{0}\}$ are the sequence-form polytopes of the two players, and \mathbf{A} is the payoff matrix for Player 1. Conveniently, the objective values of \mathcal{P}_1 and \mathcal{P}_2 will

¹https://www.gurobi.com/documentation/9.1/quickstart_linux/cs_python.html#section:Python

be the exact expected utility that each player can secure by playing against a perfectly rational opponent. Since all games are zero sum, the objective values of \mathcal{P}_1 and \mathcal{P}_2 will sum to 0 (if they don't, you must have a bug somewhere).

Problem 3.1 (15 points). Implement the linear program for computing Nash equilibrium strategies for both Player 1 and Player 2.

For each of the three games (rock-paper-superscissors, Kuhn poker, and Leduc poker), and for each of the two player, report Gurobi's output.

★ Hint: make sure to take a look at the "Gurobi tips and tricks" at the end of this document. It includes some tips as to how to debug common issues.

★ Hint: start from rock-paper-superscissors, and only then move to the more complex games.

★ Hint: since all games are zero-sum, the objective values of \mathcal{P}_1 and \mathcal{P}_2 must sum to 0.

★ Hint: the objective value for \mathcal{P}_1 should be 0 for rock-paper-superscissors, -0.0555 for Kuhn poker, and -0.0856 for Leduc poker.

Solution. ***** Your solution here. You should include six Gurobi outputs (3 games, 2 players per game). Feel free to use the `verbatim` environment in Latex to simply dump the output. Make sure to specify what game and what player each Gurobi output refers to. Don't forget to include your code at the end of your homework. For example, your output in the case of rock-paper-superscissors for Player 1 should look roughly like this ***]**

```
Gurobi Optimizer version 9.1.1 build v9.1.1rc0 (linux64)
Thread count: 8 physical cores, 16 logical processors, using up to 16 threads
Optimize a model with 4 rows, 4 columns and 12 nonzeros
Model fingerprint: 0x5264c0a3
Coefficient statistics:
  Matrix range      [1e+00, 2e+00]
  Objective range   [1e+00, 1e+00]
  Bounds range      [0e+00, 0e+00]
  RHS range         [1e+00, 1e+00]
Presolve removed 1 rows and 0 columns
Presolve time: 0.01s
Presolved: 3 rows, 4 columns, 11 nonzeros

Iteration    Objective          Primal Inf.    Dual Inf.      Time
     0         1.0000000e+00    1.000000e+00   0.000000e+00   0s
     2        -0.0000000e+00    0.000000e+00   0.000000e+00   0s

Solved in 2 iterations and 0.01 seconds
Optimal objective -0.000000000e+00
```

□

3.3 Computing optimal deterministic strategies (15 points)

In this subsection we study how much worse each player is if they (but not the opponent) are restricted to playing deterministic strategies only. To model this, we will add a constraint saying that all entries of the sequence-form strategy vectors \mathbf{x} and \mathbf{y} in (4) can only take values in $\{0, 1\}$. The resulting *integer* linear

programs—which we call $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$ —are given next.

$$\tilde{\mathcal{P}}_1 : \begin{cases} \max & \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ & \textcircled{3} \mathbf{x} \in \{0, 1\}^{|\Sigma_1|}, \mathbf{v} \text{ free,} \end{cases} \quad \tilde{\mathcal{P}}_2 : \begin{cases} \max & \mathbf{f}_1^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} -\mathbf{A} \mathbf{y} - \mathbf{F}_1^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ & \textcircled{3} \mathbf{y} \in \{0, 1\}^{|\Sigma_2|}, \mathbf{v} \text{ free.} \end{cases} \quad (5)$$

Problem 3.2 (15 points). Implement the integer linear programs given in (5) for computing optimal deterministic strategies for both Player 1 and Player 2.

For each of the three games (rock-paper-superscissors, Kuhn poker, and Leduc poker), and for each of the two player, report Gurobi’s output.

★ Hint: make sure to take a look at the “Gurobi tips and tricks” at the end of this document. It includes some tips as to how to debug common issues.

★ Hint: start from rock-paper-superscissors, and only then move to the more complex games.

★ Hint: here there are *no guarantees* that the value of $\tilde{\mathcal{P}}_1$ and the value of $\tilde{\mathcal{P}}_2$ sum to 0 anymore! In fact, that will be false in all games.

Solution. [*** Your solution here. You should include six Gurobi outputs (3 games, 2 players per game). Feel free to use the `verbatim` environment in Latex to simply dump the output. Make sure to specify what game and what player each Gurobi output refers to. Don’t forget to include your code at the end of your homework. For example, your output in the case of Kuhn poker for Player 1 should look roughly like this ***]

```
Gurobi Optimizer version 9.1.1 build v9.1.1rc0 (linux64)
Thread count: 8 physical cores, 16 logical processors, using up to 16 threads
Optimize a model with 18 rows, 18 columns and 57 nonzeros
Model fingerprint: 0x57532587
Variable types: 6 continuous, 12 integer (12 binary)
Coefficient statistics:
  Matrix range      [2e-01, 1e+00]
  Objective range   [1e+00, 1e+00]
  Bounds range      [1e+00, 1e+00]
  RHS range         [1e+00, 1e+00]
Found heuristic solution: objective -0.3333335
Presolve removed 11 rows and 10 columns
Presolve time: 0.00s
Presolved: 7 rows, 8 columns, 22 nonzeros
Found heuristic solution: objective -0.3333333
Variable types: 0 continuous, 8 integer (5 binary)

Root relaxation: objective -5.555556e-02, 9 iterations, 0.00 seconds
```

Nodes		Current Node			Objective Bounds			Work	
Expl	Unexpl	Obj	Depth	Int/Inf	Incumbent	BestBd	Gap	It/Node	Time
	0	0	-0.05556	0	3	-0.33333	-0.05556	83.3%	- 0s
H	0	0				-0.1666667	-0.05556	66.7%	- 0s
	0	0	-0.05556	0	3	-0.16667	-0.05556	66.7%	- 0s

Explored 1 nodes (9 simplex iterations) in 0.00 seconds
 Thread count was 16 (of 16 available processors)

Solution count 3: -0.166667 -0.333333 -0.333333
 No other solutions better than -0.166667

Optimal solution found (tolerance 1.00e-04)
 Best objective -1.666666666667e-01, best bound -1.666666666667e-01, gap 0.0000%

□

3.4 Controlling the amount of determinism (20 points)

In Problem 3.1 no determinism constraint was present. At the other extreme, in Problem 3.2 we insisted that at all decision points a deterministic strategy be followed. In this last subsection we will explore intermediate cases: for each value of k , we will study how much value each player can secure if they are constrained to play deterministically in at least k decision points. When $k = 0$, we will recover the objective values seen in Problem 3.1. When k is equal to the number of decision points of the player in the game, we will recover the objective values seen in Problem 3.2.

Integer programming model An optimal strategy for Player 1 subject to the constraint that at least k decision points must prescribe a deterministic strategy can be obtained as the solution to the integer linear program $\mathcal{P}_1(k)$ given in (6). Understanding the details is not important for this problem, though we included a description of the meaning of each constraint just in case you are curious.

$$\mathcal{P}_1(k) : \left\{ \begin{array}{l} \max \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t. } \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ \textcircled{3} \mathbf{x}[ja] \geq \mathbf{z}[ja] \quad \forall j \in \mathcal{J}_1 : p_j = \emptyset, a \in A_j \\ \textcircled{4} \mathbf{x}[ja] \geq \mathbf{x}[p_j] + \mathbf{z}[ja] - 1 \quad \forall j \in \mathcal{J}_1 : p_j \neq \emptyset, a \in A_j \\ \textcircled{5} \sum_{a \in A_j} \mathbf{z}[ja] \leq 1 \quad \forall j \in \mathcal{J}_1 \\ \textcircled{6} \sum_{j \in \mathcal{J}_1} \sum_{a \in A_j} \mathbf{z}[ja] \geq k \\ \textcircled{7} \mathbf{x} \geq \mathbf{0}, \mathbf{z} \in \{0, 1\}^{|\Sigma_1|}, \mathbf{v} \text{ free,} \end{array} \right. \quad (6)$$

- $\mathbf{z} \in \{0, 1\}^{|\Sigma_1|}$ is a binary vector that decides, for each strategy $ja \in \Sigma_1$ of Player 1, whether to pick action a at j with probability 1. Since the strategy vector \mathbf{x} is expressed in sequence-form, picking action a with probability 1 at j is expressed through constraints $\textcircled{3}$ and $\textcircled{4}$.
- Constraint $\textcircled{5}$ asserts that no more than one action at each decision point can be forced to be played with probability 1.
- Constraint $\textcircled{6}$ asserts that in at least k decision point, exactly one of the actions will be forced to be played with probability 1.

The integer linear program $\mathcal{P}_2(k)$ for Player 2 is analogous.

Problem 3.3 (15 points). Implement the integer linear programs $\mathcal{P}_1(k)$ and $\mathcal{P}_2(k)$, described above, for computing optimal strategies with a given lower bound on the amount of determinism.

For each of the three games (rock-paper-superscissors, Kuhn poker, and Leduc poker), and for each of the two player i , plot the objective value of $\mathcal{P}_i(k)$ as a function of $k \in \{0, \dots, |\mathcal{J}_i|\}$ (number of decision points of Player i).

★ Hint: make sure to take a look at the “Gurobi tips and tricks” at the end of this document. It includes some tips as to how to debug common issues.

★ Hint: Gurobi can be pretty verbose by default. For this problem, if you would like to silence Gurobi you can use `m.setParam("OutputFlag", 0)`

★ Hint: For Leduc poker, if Gurobi is taking too long to optimize when k is large, you can lower the solution precision by calling `m.setParam("MIPGap", 0.01)` before `m.optimize()`. Expect a runtime of up to one-five hours for Leduc poker, depending on how powerful the machine you are using is.

Solution. ***** Your solution here. You should include six plots (3 games, 2 players per game). Make sure to specify what game and what player each plot refers to. Don't forget to include your code at the end of your homework. ***** □

Problem 3.4 (5 points). Comment on the results you obtained in this problem: do highly-deterministic strategies for the three small games exist? Are the results what you expected? If yes, what did he results confirm? If not, how do you think you can reconcile your previous intuition with the experimental findings?

Solution. ***** Your solution here ***** □

A Appendix: Gurobi tips and tricks

Basic notation Let `m` denote the Gurobi model object. Then, here is a quick cookbook.

- Add a continuous free variable:
`m.addVar(-GRB.INFINITY, GRB.INFINITY, vtype=GRB.CONTINUOUS, name="human_var_name_here")`
- Add a continuous nonnegative variable:
`m.addVar(0.0, GRB.INFINITY, vtype=GRB.CONTINUOUS, name="human_var_name_here")`
- Add a binary variable:
`m.addVar(0.0, 1.0, vtype=GRB.BINARY, name="human_var_name_here")`
- Add an equality constraint:
`m.addConstr(lhs == rhs)`
- Add an inequality (\geq) constraint:
`m.addConstr(lhs >= rhs)`
- Set a maximization objective:
`m.setObjective(obj, sense=GRB.MAXIMIZE)`

Accessing the solution After calling `m.optimize()`, you can obtain the objective value by calling

```
m.getAttr(GRB.Attr.ObjVal)
```

If you want to inspect the solution, given a variable object `v` (the object returned by `m.addVar`), you can access the value of `v` in the current solution by calling

```
v.getAttr(GRB.Attr.X)
```

Troubleshooting First of all, if you are having a problem with Gurobi, the first thing you should try to do is to ask Gurobi to dump the model that it thinks you are asking to solve to a file in a human readable format. *Reading the model file will be so much easier if you gave names to the variables in your model, using the ‘name’ optional argument of `addVar`.*

To have Gurobi dump the model, you can use something like this:

```
m.write("/tmp/model.lp")
```

Of course, you can specify a different path. However, it is important that you keep the ‘.lp’ extension: there are multiple format that Gurobi can output, and it uses the file extension to guess which format you want.

Beyond the general rule of thumb above, make sure of the following:

- Start from rock-paper-superscissors. There, the `/tmp/model.lp` file for Player 1 for Problem 3.1 should look something like this (probably with different variable names):

```
\ Model game_value_p1
\ LP format - for model browsing. Use MPS format to capture full model detail.
Maximize
    v[d1_p12]
Subject To
    R0: x[('d1_p11','p')] - x[('d1_p11','s')] - v[d1_p12] >= 0
    R1: - x[('d1_p11','r')] + 2 x[('d1_p11','s')] - v[d1_p12] >= 0
    R2: x[('d1_p11','r')] - 2 x[('d1_p11','p')] - v[d1_p12] >= 0
    R3: x[('d1_p11','r')] + x[('d1_p11','p')] + x[('d1_p11','s')] = 1
Bounds
    v[d1_p12] free
End
```

Note: Gurobi omits listing nonnegative variables in the **Bounds** section.

- Did you remember to specify that you want a *maximization* problem? (Gurobi’s default is minimization) If Gurobi says that the model is unbounded, chances are you forgot.
- Check that the number of variables and constraints is what you expect. Are the sense of the constraints (equality, \leq , \geq) what you wanted?

References

- H. W. Kuhn. A simplified two-person poker. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1 of *Annals of Mathematics Studies*, 24, pages 97–103. Princeton University Press, Princeton, New Jersey, 1950.
- Finnegan Southey, Michael Bowling, Bryce Larson, Carmelo Piccione, Neil Burch, Darse Billings, and Chris Rayner. Bayes’ bluff: opponent modelling in poker. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, pages 550–558, 2005.