

Homework 2

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

Due on: Oct. 19, 2021, beginning of class

Instructions Please typeset your solution in the `.tex` file provided in the homework `.zip`. Attach a readable printout of your code at the end of the document. Turn in your printed solution at the beginning of class on October 19th. Don't forget to fill in your student ID above :-)

Bonus points This homework contains 2 problems (each split into smaller tasks), for a total of 120 points. The goal of the 20 bonus points is to help you in case you lost some points here and there. The maximum score for the homework will be capped at 100. In other words, your score will be set to the minimum between 100 and the sum of the points you scored on all questions, including the bonus questions.

1 Optimistic multiplicative weights update (40 + 10 points)

Multiplicative weights update (MWU) and its predictive variant called *optimistic multiplicative weights update (OMWU)* are popular regret minimization algorithms for the probability simplex

$$\Delta^n := \{(x_1, \dots, x_n) \in \mathbb{R}_{\geq 0}^n : x_1 + \dots + x_n = 1\}. \quad (1)$$

They enjoy many strong theoretical properties, and were involved in a series of important papers in game theory. In this question, you will derive and analyze MWU and OMWU from first principles.

1.1 The negative entropy regularizer (10 points)

OMWU is just a special name for the predictive FTRL algorithm when the regularizer $\omega : \Delta^n \rightarrow \mathbb{R}$ is chosen to be the *negative entropy regularizer*

$$\omega(\mathbf{x}) := \sum_{i=1}^n x_i \log(x_i).$$

To avoid annoying issues with the logarithm of 0, we will only ever evaluate and differentiate ω in the relative interior of Δ^n , that is the set

$$\text{relint } \Delta^n = \{(x_1, \dots, x_n) \in \mathbb{R}_{> 0}^n : x_1 + \dots + x_n = 1\}$$

(note the strict inequality $\mathbb{R}_{> 0}$, as opposed to $\mathbb{R}_{\geq 0}$ in (1)).

For ω to be a valid choice of regularizer in predictive FTRL, we need to check that ω is 1-strongly convex with respect to some norm. In particular, it turns out that ω is 1-strongly convex both with respect to the Euclidean norm

$$\|\mathbf{x}\|_2 := \sqrt{\sum_{i=1}^n x_i^2} \quad \forall \mathbf{x} \in \mathbb{R}^n$$

and with respect to the ℓ_1 norm

$$\|\mathbf{x}\|_1 := \sum_{i=1}^n |x_i| \quad \forall \mathbf{x} \in \mathbb{R}^n.$$

The easiest way to verify strong convexity in this case passes through the following well-known characterization.

Lemma 1. Let $\mathcal{X} \subseteq \mathbb{R}^n$ be a convex set, $f : \mathcal{X} \rightarrow \mathbb{R}$ be a twice-differentiable function with Hessian matrix $\nabla^2 f(\mathbf{x})$ at every $\mathbf{x} \in \mathcal{X}$, and $\|\cdot\|$ be a norm. If

$$\mathbf{u}^\top \nabla^2 f(\mathbf{x}) \mathbf{u} \geq \|\mathbf{u}\|^2 \quad \forall \mathbf{u} \in \mathbb{R}^n, \mathbf{x} \in \mathcal{X},$$

then f is 1-strongly convex on \mathcal{X} with respect to norm $\|\cdot\|$.

In the next two exercises you will use Lemma 1 to verify that ω is 1-strongly convex on $\text{relint } \Delta^n$ with respect to $\|\cdot\|_2$ and $\|\cdot\|_1$.

Problem 1.1 (5 points). Apply Lemma 1 for $\mathcal{X} = \text{relint } \Delta^n$, $f = \omega$, and $\|\cdot\| = \|\cdot\|_2$ and conclude that ω is 1-strongly convex with respect to the Euclidean norm on $\text{relint } \Delta^n$.

- ★ Hint: the Hessian matrix of ω is particularly nice. Start by working that out first.
- ★ Hint: at some point, it might be useful to argue that $1/x_i \geq 1$ for any $i \in \{1, \dots, n\}$ whenever $\mathbf{x} \in \text{relint } \Delta^n$.

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

□

Problem 1.2 (5 points). Apply Lemma 1 for $\mathcal{X} = \text{relint } \Delta^n$, $f = \omega$, and $\|\cdot\| = \|\cdot\|_1$ and conclude that ω is 1-strongly convex with respect to the ℓ_1 norm on $\text{relint } \Delta^n$.

- ★ Hint: The Cauchy-Schwarz inequality asserts that for any pair of vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$,

$$\left(\sum_{i=1}^n a_i b_i \right)^2 \leq \left(\sum_{i=1}^n a_i^2 \right) \left(\sum_{i=1}^n b_i^2 \right). \quad (2)$$

Now, let $\mathbf{x} \in \text{relint } \Delta^n$ and $\mathbf{u} \in \mathbb{R}^n$, and consider the vectors $\mathbf{a} := (u_1/\sqrt{x_1}, \dots, u_n/\sqrt{x_n})$ and $\mathbf{b} := (\sqrt{x_1}, \dots, \sqrt{x_n})$. What happens if you plug them into (2)? Don't forget that $x_1 + \dots + x_n = 1$ since $\mathbf{x} \in \text{relint } \Delta^n$.

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

□

1.2 Gradient of ω and of its conjugate (15 points)

In this subsection, you will derive a formula for the gradient of ω and for the gradient of its convex conjugate. Let's start with the gradient.

Problem 1.3 (5 points). Give an expression for the gradient of ω at any point $\mathbf{x} \in \text{relint } \Delta^n$.

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

□

Now, let's focus on the gradient of the convex conjugate of ω , that is, the solution to the optimization problem

$$\nabla \omega^*(\mathbf{g}) := \arg \max_{\hat{\mathbf{x}} \in \text{relint } \Delta^n} \{\mathbf{g}^\top \hat{\mathbf{x}} - \omega(\hat{\mathbf{x}})\}. \quad (3)$$

Problem (3) is a *constrained* optimization problem, since the optimization variable $\hat{\mathbf{x}}$ is constrained to satisfy $\hat{\mathbf{x}} \in \text{relint } \Delta^n$. Call \mathbf{x}^* the optimal solution to (3). As a result of an important theorem in optimization theory (the Lagrange multiplier theorem), there exists a constant (called *Lagrange multiplier*) $\alpha \in \mathbb{R}$ such that

$$\mathbf{g} - \nabla\omega(\mathbf{x}^*) = \alpha\mathbf{1}, \quad (4)$$

where $\mathbf{1} \in \mathbb{R}^n$ is the vector of all ones.

Problem 1.4 (5 points). Plug in the expression for the gradient of ω that you developed in Problem 1.3 into (4). Note that (4) is a vector equation, and therefore it is equivalent to a system of n scalar equations. Isolate and solve for x_i^* for every $i \in \{1, \dots, n\}$, and show that

$$x_i^* = e^{-1-\alpha} \cdot e^{g_i} \quad \forall i \in \{1, \dots, n\}. \quad (5)$$

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

Problem 1.5 (5 points). Use Equation (5) together with the fact that the sum of the entries of $\mathbf{x}^* \in \text{relint } \Delta^n$ must be 1 to solve for the value of the Lagrange multiplier α . Then, plug in the value of α to conclude that for any $\mathbf{g} \in \mathbb{R}^n$, $\nabla\omega^*(\mathbf{g})$ —that is, the solution to the argmax in (3)—satisfies

$$x_i^* = \frac{e^{g_i}}{\sum_{j=1}^n e^{g_j}} \quad \forall i \in \{1, \dots, n\}.$$

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

1.3 OMWU as predictive FTRL (15 points)

Now that we verified that ω is 1-strongly convex, we can safely run predictive FTRL with ω as a regularizer. As a reminder, predictive FTRL is the algorithm recalled in Algorithm 1. In our case, \mathcal{X} will be the relative interior $\text{relint } \Delta^n$ of the probability simplex Δ^n , the regularizer φ will be the negative entropy function ω , and $\eta > 0$ will be a generic stepsize. The resulting algorithm is called OMWU.

Problem 1.6 (10 points). Use the characterization of $\nabla\omega^*(\mathbf{g})$ given in the statement of Problem 1.5 to prove that at times $t = 2, 3, \dots$, for all $i \in \{1, \dots, n\}$, the strategies $\mathbf{x}^t \in \Delta^n$ produced by OMWU satisfy^a

$$x_i^t = \frac{x_i^{t-1} \exp\{\eta(\ell_i^{t-1} - m_i^{t-1} + m_i^t)\}}{\sum_{j=1}^n x_j^{t-1} \exp\{\eta(\ell_j^{t-1} - m_j^{t-1} + m_j^t)\}},$$

^aFor readability we used the notation $\exp\{\bullet\}$ to mean e^\bullet .

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

Since OMWU is just predictive FTRL, we can use the known regret bound for predictive FTRL we saw in class—and the following proposition recalls—to give a regret guarantee for OMWU.

Algorithm 1: Predictive FTRL

Data: $\mathcal{X} \subseteq \mathbb{R}^n$ convex domain
 $\varphi : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ strongly convex regularizer
 $\eta > 0$ step-size parameter

```
1  $\mathbf{L}^0 \leftarrow \mathbf{0} \in \mathbb{R}^n$ 
2 function NEXTSTRATEGY( $\mathbf{m}^t$ )
   [ $\triangleright$  Set  $\mathbf{m}^t = \mathbf{0}$  for non-predictive version]
3   return  $\arg \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ (\mathbf{L}^{t-1} + \mathbf{m}^t)^\top \hat{\mathbf{x}} - \frac{1}{\eta} \varphi(\hat{\mathbf{x}}) \right\}$ 
4 function OBSERVEUTILITY( $\ell^t$ )
5    $\mathbf{L}^t \leftarrow \mathbf{L}^{t-1} + \ell^t$ 
```

Algorithm 2: Predictive OMD

Data: $\mathcal{X} \subseteq \mathbb{R}^n$ convex domain
 $\varphi : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ strongly convex regularizer
 $\eta > 0$ step-size parameter

```
1  $\mathbf{z}^0 \leftarrow \arg \min_{\mathbf{z} \in \mathcal{X}} \omega(\mathbf{z})$ 
2 function NEXTSTRATEGY( $\mathbf{m}^t$ )
   [ $\triangleright$  Set  $\mathbf{m}^t = \mathbf{0}$  for non-predictive version]
3   return  $\arg \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ (\mathbf{m}^t)^\top \hat{\mathbf{x}} - \frac{1}{\eta} D_\varphi(\hat{\mathbf{x}} \parallel \mathbf{z}^{t-1}) \right\}$ 
4 function OBSERVEUTILITY( $\ell^t$ )
5    $\mathbf{z}^t \leftarrow \arg \max_{\hat{\mathbf{z}} \in \mathcal{X}} \left\{ (\ell^t)^\top \hat{\mathbf{z}} - \frac{1}{\eta} D_\varphi(\hat{\mathbf{z}} \parallel \mathbf{z}^{t-1}) \right\}$ 
```

Proposition 1. Consider the predictive FTRL algorithm given in Algorithm 1. Let Ω denote the range of φ over \mathcal{X} , that is, $\Omega := \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} \{\varphi(\mathbf{x}) - \varphi(\mathbf{x}')\}$. For any T , the regret cumulated up to time T satisfies

$$R^T \leq \frac{\Omega}{\eta} + \eta \sum_{t=1}^T \|\ell^t - \mathbf{m}^t\|_*^2 - \frac{1}{4\eta} \sum_{t=2}^T \|\mathbf{x}^t - \mathbf{x}^{t-1}\|^2,$$

where $\|\cdot\|$ is any norm with respect to which φ is 1-strongly convex, and $\|\cdot\|_*$ is the dual of the norm $\|\cdot\|$.

Proposition 1 was stated in general for any instantiation of FTRL. In the particular case of OMWU, the negative entropy function ω was proven to be 1-strongly convex with respect to both the Euclidean norm (Problem 1.1) and the ℓ_1 norm (Problem 1.2). So, in principle, either norm can be used in Proposition 1. However, one choice dominates the other.

Problem 1.7 (2 points). The negative entropy function ω is 1-strongly convex with respect to both the Euclidean norm (Problem 1.1) and the ℓ_1 norm (Problem 1.2). So, in principle, either norm can be used when invoking Proposition 1. Which norm do you think leads to a stronger regret bound, and why?

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

□

Problem 1.8 (3 points). Prove that the range $\Omega = \sup_{\mathbf{x}, \mathbf{x}' \in \text{relint } \Delta^n} \{\omega(\mathbf{x}) - \omega(\mathbf{x}')\}$ of ω on $\text{relint } \Delta^n$ is $\Omega = \log n$. Then, use Proposition 1—which was stated in general for any instantiation of FTRL—to argue that OMWU for the simplex Δ^n satisfies the regret bound

$$R^T \leq \frac{\log n}{\eta} + \eta \sum_{t=1}^T \|\ell^t - \mathbf{m}^t\|_\infty^2 - \frac{1}{4\eta} \sum_{t=2}^T \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_1^2.$$

-
- ★ Hint: The minimizer \mathbf{x}^* of ω over $\text{relint } \Delta^n$ is $\nabla \omega^*(\mathbf{0})$. You already know how to compute this from Problem 1.5.
 - ★ Hint: The supremum of ω over $\text{relint } \Delta^n$ is 0 (you should prove this).

★ Hint: You can take for granted the fact that $\|\cdot\|_\infty$ is the dual norm of $\|\cdot\|_1$.

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

□

1.4 OMWU as predictive OMD (10 bonus points)

It turns out that OMWU—which was defined as the instance of predictive FTRL in which the regularizer is set the negative entropy function—is equivalent to predictive OMD with negative entropy function, in the sense that the two algorithms produce the same iterates at every time t . Predictive OMD is recalled in Algorithm 2. As a reminder, the Bregman divergence $D_\varphi(\cdot\|\cdot)$ is defined with respect to any regularizer φ and any two points \mathbf{x}, \mathbf{c} as

$$D_\varphi(\mathbf{x}\|\mathbf{c}) := \varphi(\mathbf{x}) - \varphi(\mathbf{c}) - \nabla\varphi(\mathbf{c})^\top(\mathbf{x} - \mathbf{c}).$$

Problem 1.9 (10 bonus points). Consider the predictive OMD algorithm (Algorithm 2), where \mathcal{X} is set to be the relative interior $\text{relint } \Delta^n$ of the n -simplex, the regularizer φ is set to be the negative entropy function ω , and $\eta > 0$ is a generic stepsize. Prove that the iterates produced by that algorithm coincide with those produced by OMWU as defined in Section 1.3.

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

□

2 Linear programming for solving zero-sum games, and application to low randomization in poker (60 + 10 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 the three games of Homework 1: 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, according to the same format of Homework 1, recalled in Section 2.1. The zip of the homework also contains a stub Python file to help you set up your implementation.

2.1 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

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

So, the game file in this case has content:

```
{
  "decision_problem_p11": [
    {"id": "d1_p11", "type": "decision", "actions": ["r", "p", "s"],
      "parent_edge": null, "parent_sequence": null}
  ],
  "decision_problem_p12": [
```

```

    {"id": "d1_p12", "type": "decision", "actions": ["r", "p", "s"],
     "parent_edge": null, "parent_sequence": null}
  ],
  "utility_p11": [
    {"sequence_p11": ["d1_p11", "r"], "sequence_p12": ["d1_p12", "p"], "value": -1},
    {"sequence_p11": ["d1_p11", "r"], "sequence_p12": ["d1_p12", "s"], "value": 1},
    {"sequence_p11": ["d1_p11", "p"], "sequence_p12": ["d1_p12", "r"], "value": 1},
    {"sequence_p11": ["d1_p11", "p"], "sequence_p12": ["d1_p12", "s"], "value": -2},
    {"sequence_p11": ["d1_p11", "s"], "sequence_p12": ["d1_p12", "r"], "value": -1},
    {"sequence_p11": ["d1_p11", "s"], "sequence_p12": ["d1_p12", "p"], "value": 2}
  ]
}

```

2.2 Computing the value of the game (20 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:

$$\mathcal{P}_1 : \begin{cases} \max & \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ & \textcircled{3} \mathbf{x} \geq \mathbf{0}, \mathbf{v} \text{ free,} \end{cases} \quad \mathcal{P}_2 : \begin{cases} \max & \mathbf{f}_1^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} -\mathbf{A} \mathbf{y} - \mathbf{F}_1^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ & \textcircled{3} \mathbf{y} \geq \mathbf{0}, \mathbf{v} \text{ free,} \end{cases} \quad (6)$$

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 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 2.1 (20 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.

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

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
```

□

2.3 Computing optimal deterministic strategies (20 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 (6) 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.

$$\tilde{\mathcal{P}}_1 : \begin{cases} \max & \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ & \textcircled{3} \mathbf{x} \in \{0, 1\}^{|\Sigma_1|}, \mathbf{v} \text{ free,} \end{cases} \quad \tilde{\mathcal{P}}_2 : \begin{cases} \max & \mathbf{f}_1^\top \mathbf{v} \\ \text{s.t.} & \textcircled{1} -\mathbf{A} \mathbf{y} - \mathbf{F}_1^\top \mathbf{v} \geq \mathbf{0} \\ & \textcircled{2} \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ & \textcircled{3} \mathbf{y} \in \{0, 1\}^{|\Sigma_2|}, \mathbf{v} \text{ free.} \end{cases} \quad (7)$$

Problem 2.2 (20 points). Implement the integer linear programs given in (7) 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 IntInf | Incumbent    BestBd   Gap | It/Node Time
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----
    0     0 -0.05556   0   3  -0.33333  -0.05556  83.3%   -   0s
H    0     0          0   0          -0.166667  -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%
```

□

2.4 Controlling the amount of determinism (20 + 10 bonus points)

In Problem 2.1 no determinism constraint was present. At the other extreme, in Problem 2.2 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 2.1. 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 2.2.

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 (8). 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.

$$\mathcal{P}_1(k) : \left\{ \begin{array}{l} \max \mathbf{f}_2^\top \mathbf{v} \\ \text{s.t. } \textcircled{1} \mathbf{A}^\top \mathbf{x} - \mathbf{F}_2^\top \mathbf{v} \geq \mathbf{0} \\ \textcircled{2} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ \textcircled{3} \mathbf{x}[ja] \geq \mathbf{z}[ja] \quad \forall j \in \mathcal{J}_1 : p_j = \emptyset, a \in A_j \\ \textcircled{4} \mathbf{x}[ja] \geq \mathbf{x}[p_j] + \mathbf{z}[ja] - 1 \quad \forall j \in \mathcal{J}_1 : p_j \neq \emptyset, a \in A_j \\ \textcircled{5} \sum_{a \in A_j} \mathbf{z}[ja] \leq 1 \quad \forall j \in \mathcal{J}_1 \\ \textcircled{6} \sum_{j \in \mathcal{J}_1} \sum_{a \in A_j} \mathbf{z}[ja] \geq k \\ \textcircled{7} \mathbf{x} \geq \mathbf{0}, \mathbf{z} \in \{0, 1\}^{|\Sigma_1|}, \mathbf{v} \text{ free,} \end{array} \right. \quad (8)$$

- $\mathbf{z} \in \{0, 1\}^{|\Sigma_1|}$ 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 2.3 (20 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 2.4 (10 bonus 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 the 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 formats 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 2.1 should look something like this (probably with different variable names):

```

\ Model game_value_p11
\ 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.