ORF 363/COS 323	PROBLEM SET 7	Fall 2025, Princeton University
Instructor: A.A. Ahmadi		
TAs: Budway, Hua, Shi, Tang, Yang		
Due on Tuesday, November 25, 2025, at 1:30pm EST, on Gradescope		

For all problems that involve coding, please include your code.

Problem 1: Setting the odds in your favor with semidefinite programming
As the CEO of TigerCasino in Vegas, you are introducing a new game on your floor. In this game, a player rolls a die twice and receives q_{ij} dollars from the casino if the die shows i on one roll and j on the other (the order does not matter). The matrix $Q = (q_{ij})_{1 \leq i,j \leq 6}$ is announced to the player:

$$Q = 100 \times \begin{pmatrix} 4 & -2 & 1 & -1 & -2 & 1 \\ -2 & 4 & 1 & -2 & -2 & -1 \\ 1 & 1 & 4 & -2 & 1 & -1 \\ -1 & -2 & -2 & 4 & 1 & -1 \\ -2 & -2 & 1 & 1 & 4 & -1 \\ 1 & -1 & -1 & -1 & -1 & 4 \end{pmatrix}.$$

Gamblers are leaving the Bellagios and rushing to your table because if the die was fair (i.e., had a probability of $\frac{1}{6}$ assigned to each outcome of a roll), they would make \$11.11 in expectation in every play. Little do they know, however, that you have been using your optimization knowledge to optimally bias the die and maximize the profit of TigerCasino. Let x_i be the probability that the die comes out i. The optimization problem of interest to you is:

$$\begin{aligned} \min_{x \in \mathbb{R}^6} \quad & x^T Q x \\ \text{s.t.} \quad & x \geq 0 \\ & \sum_{i=1}^{6} x_i = 1. \end{aligned} \tag{1}$$

The constraints make sure that x is a valid probability vector and the objective function is the expected payoff of the player in every play.

(a) Is problem (1) a convex optimization problem? Why or why not?

(b) Recall that a matrix $A \in \mathbb{S}^{n \times n}$ is copositive if $x^T A x \geq 0$ for all $x \geq 0$. Denote the set of $n \times n$ copositive matrices by $\mathcal{C}_n$. Show that the optimal value of (1) is equal to the optimal value of the following problem:

$$\begin{aligned} \max_{t \in \mathbb{R}} \quad & t \\ \text{s.t.} \quad & Q - tJ \in \mathcal{C}_6. \end{aligned} \tag{2}$$

Here, $J \in \mathbb{S}^{6\times 6}$ is the all-ones matrix.

(c) Denote the optimal value of (1) (or equivalently (2)) by OPT. Denote the optimal value of the semidefinite program

$$\begin{aligned} \max_{t \in \mathbb{R}, N \in \mathbb{S}^{6\times 6}} \quad & t \\ \text{s.t.} \quad & Q - tJ - N \succeq 0 \\ & N \geq 0 \end{aligned} \tag{3}$$

by SDP_{OPT}. (Here, "$\geq$" denotes an entrywise nonnegativity constraint and "$\succeq$" denotes a positive semidefiniteness constraint.) Show that $SDP_{OPT} \leq OPT$.

(d) Report SDP_{OPT} by solving (3) in CVX or CVXPY. Show that $SDP_{OPT} = OPT$ by presenting a vector $x^* \in \mathbb{R}^6$ that is feasible to (1) and makes the objective function of (1) equal to SDP_{OPT}. (Hint: you may wish to start with an eigenvector associated with the smallest eigenvalue of $Q - t^*J - N^*$, where (t^*, N^*) form an optimal solution to (3).) What probability does your optimal die assign to each of its six outcomes? What is the expected win/loss of TigerCasino in dollars every time a player plays this game?

Problem 2: Nearest correlation matrix

You are the Head of Quantitative Strategies at HoneyMoney Technologies LLC, a new hedge fund firm in NYC whose proprietary optimization algorithms has Wall Street raving. Your main competitor, RenaissancE Technologies[1], has sent in a spy, disguised as a summer intern, to interfere with your investments. The spy has gotten his hands on your *correlation matrix* of n important stocks[2], to which he has added some random noise, leaving you with a matrix $\hat{C}$. We remark that to be a valid correlation matrix, a matrix must be symmetric, positive semidefinite, and have all diagonal entries equal to one. The spy has been careful enough to make sure that the resulting matrix $\hat{C}$ is symmetric and has ones on the diagonal, but he hasn't noticed that his change has made $\hat{C}$ not positive semidefinite.

(a) Suppose we have

$$\hat{C} = \begin{pmatrix} 1.00 & -0.76 & 0.07 & -0.96 \\ -0.76 & 1.00 & 0.18 & 0.07 \\ 0.07 & 0.18 & 1.00 & 0.41 \\ -0.96 & 0.07 & 0.41 & 1.00 \end{pmatrix}.$$

Using CVX or CVXPY, recover the nearest (valid) correlation matrix to $\hat{C}$ in Frobenius norm (i.e., the correlation matrix C that minimizes $\|C - \hat{C}\|_F$). Give your optimal solution.

(b) Show that for any symmetric matrix $\hat{C}$, the problem of finding the closest correlation matrix to $\hat{C}$ in Frobenius norm has a unique solution. (You can use the fact that an optimal solution to this problem exists without proof.)

[1] Not to be confused with Renaissance Technologies that would never do such a thing.

[2] If you are curious, the correlation matrix is an $n \times n$ symmetric matrix used frequently in investment banking. Its (i, j)-th entry is a number between -1 and 1, with numbers close to 1 meaning that stocks i and j are likely to move up together, close to -1 meaning that the two stocks are likely to move in opposite directions, and close to zero meaning that they are likely uncorrelated. The problem of finding the closest correlation matrix to a given matrix is an important problem in financial engineering; see e.g. this article.

Problem 3: Bitcoin fever

Decentralized exchanges (DEXs) enable users to trade cryptocurrencies without relying on a central authority. We consider a DEX which offers exchanges of n assets. Let $\mathbb{R}^n_+$ denote the set of entrywise nonnegative vectors in $\mathbb{R}^n$. The DEX has a reserve vector $r \in \mathbb{R}^n_+$, where entry r_i denotes the quantity of asset i in the reserves. A proposed trade by a user consists of two vectors $x \in \mathbb{R}^n_+, y \in \mathbb{R}^n_+$, where x_i denotes the quantity of asset i that the trader proposes to give to the DEX, and y_i denotes the quantity of asset i that the trader proposes to receive from the DEX. The DEX has a trading function $T : \mathbb{R}^n_+ \to \mathbb{R}$ that needs to remain constant through every exchange. More specifically, the DEX accepts a proposed trade (x, y) by a user if and only if

$$T(r + \gamma x - y) = T(r),$$

where $\gamma \in (0, 1)$ accounts for a small trading fee. The user wants to find a trade (x, y) that is accepted by the DEX and maximizes their utility function $U : \mathbb{R}^n \to \mathbb{R}$. This can be achieved by solving the following optimization problem:

$$\begin{array}{ll} \max\limits_{x,y\in\mathbb{R}^n} & U(y - x) \\ \text{s.t.} & T(r + \gamma x - y) = T(r) \\ & r + \gamma x - y \geq 0 \\ & 0 \leq x_i \leq x_{max}, \quad i = 1, \dots, n \\ & 0 \leq y_i \leq y_{max}, \quad i = 1, \dots, n. \end{array} \tag{4}$$

Here, the scalars $x_{\max}$ and $y_{\max}$ denote bounds on the quantity of assets that the DEX is willing to receive and give out. In practice, the trading function T and the utility function U are both concave. Nevertheless, optimization problem (4) is generally nonconvex due to the equality constraint. Thus, we consider the following convex relaxation of (4):

$$\begin{array}{ll} \max\limits_{x,y\in\mathbb{R}^n} & U(y - x) \\ \text{s.t.} & T(r + \gamma x - y) \geq T(r) \\ & r + \gamma x - y \geq 0 \\ & 0 \leq x_i \leq x_{max}, \quad i = 1, \dots, n \\ & 0 \leq y_i \leq y_{max}, \quad i = 1, \dots, n. \end{array} \tag{5}$$

(a) Show that problem (5) is indeed a convex optimization problem.

(b) We say a function $f : \Omega \to \mathbb{R}$, where Ω is a subset of $\mathbb{R}^n$, is *non-decreasing* if for all vectors $u, v \in \Omega$, $u \leq v$ entrywise implies $f(u) \leq f(v)$. A non-decreasing function f is *increasing* if whenever $u \leq v$ entrywise and $u_i < v_i$ for some $i \in \{1, \ldots, n\}$, we have $f(u) < f(v)$. Suppose the functions U and T are both continuous, with T being non-decreasing and U being increasing. Show that the set of optimal solutions to (4) is the same as the set of optimal solutions to (5).

(c) Consider the setting of exchanging two assets, i.e., when $n = 2$. Let the utility function $U : \mathbb{R}^2 \to \mathbb{R}$ be given as $U(z) = c^T z$, where $c \in \mathbb{R}^2$ is a given vector with positive entries. Let the trading function $T : \mathbb{R}^2_+ \to \mathbb{R}$ be defined as $T(z) = \sqrt{z_1 z_2}$. Show that in this setting, the set of optimal solutions to (4) is the same as the set of optimal solutions to the following semidefinite program:

$$\begin{aligned}
\max_{x,y \in \mathbb{R}^2} \quad & U(y - x) \\
\text{s.t.} \quad & \begin{bmatrix} r_1 + \gamma x_1 - y_1 & T(r) \\ T(r) & r_2 + \gamma x_2 - y_2 \end{bmatrix} \succeq 0 \\
& 0 \leq x_1 \leq x_{max}, \quad 0 \leq x_2 \leq x_{max} \\
& 0 \leq y_1 \leq y_{max}, \quad 0 \leq y_2 \leq y_{max}.
\end{aligned} \tag{6}$$

(d) In the setting of part (c), let $\gamma = 0.9$, $r = (1, 2)^T$, $x_{max} = y_{max} = 1$, and $c = (2t, 1)^T$, where t is a parameter. Take $t = 0.5, 0.6, 0.7, \ldots, 2.0$ and for each value of t solve the optimization problem (6). For $i = 1, 2$, plot the net change $y_i - x_i$ of the optimal trade versus t. Briefly comment on what you observe.