MS&E 111X & 211X
Introduction to Optimization (Accelerated)
Homework 3
Course Instructor: Yinyu Ye
Due Date: 11:59 pm Nov 4, 2021

Please submit your homework through Gradescope. If you haven't already been added to Gradescope, you can use the entry code **2RJNKV** to join. Please note: late homework will not be accepted. Each problem will be graded out of 10 points.
Some problems allow group work. Groups should be no larger than 4. If you decide to work together, provide the names of those you worked with.

Problem 1

The Cobb-Douglas production function is widely used in economics to represent the relationship between inputs and outputs of a firm. It takes on the form $Y = AL^\alpha K^\beta$ where Y represents output, L labor and K capital. α and β are constants that determine how production is scaled. We find that the Cobb-Douglas function can be applied to model firm utility. Consider the following utility maximization problem:

$$\begin{aligned} \max \quad & u(x) = x_1^\alpha x_2^{1-\alpha} \\ \text{subject to} \quad & p_1x_1 + p_2x_2 \leq w \\ & x_1, x_2 \geq 0 \end{aligned}$$

where $0 < \alpha < 1$ is fixed and p and w represent a given price and budget respectively. This is a particular instance of Cobb-Douglas utility. Assume that $w > 0$ and $p > 0$.

a) Perform a logarithmic transformation to find an equivalent maximization problem. Explain why this transformation leads to an equivalent problem.

b) Write the KKT conditions for the transformed problem and find an explicit solution for x as a function of p, w and α. Are these conditions sufficient for optimality?

c) Find the corresponding Lagrangian multiplier, λ, in terms of p, w and α. Describe an interpretation for λ.

d) Suppose $w = 100$, $p_1 = 1$, $p_2 = 2$, and $\alpha = 0.2$, find the optimal consumption bundle x_1 and x_2.

Solution:

a) First, note that with $w > 0$ and $p > 0$, a zero solution is clearly not optimal so we can take the log excluding potential solutions at zero. The log-transformed problem is

$$\begin{aligned} \max \quad & u(x) = \alpha \log x_1 + (1-\alpha) \log x_2 \\ \text{subject to} \quad & p_1 x_1 + p_2 x_2 \leq w \\ & x_1, x_2 \geq 0 \end{aligned}$$

Logarithmic functions are monotonous. Therefore, the optimal solution to the transformed problem is the same as the optimal solution to the original problem.

b) Lagrange multipliers:

$$\begin{aligned} \lambda : \quad & p_1 x_1 + p_2 x_2 \leq w \\ s_1 : \quad & x_1 \geq 0 \\ s_2 : \quad & x_2 \geq 0 \end{aligned}$$

Lagrangian:

$$L(x_1, x_2, \lambda, s_1, s_2) = \alpha \log x_1 + (1-\alpha) \log x_2 - \lambda(p_1 x_1 + p_2 x_2 - w) - s_1 x_1 - s_2 x_2$$

KKT conditions:

$$\begin{gathered} p_1 x_1 + p_2 x_2 \leq w \\ x_1, x_2 \geq 0 \\ \frac{\alpha}{x_1} - p_1 \lambda - s_1 = 0 \\ \frac{1-\alpha}{x_2} - p_2 \lambda - s_2 = 0 \\ \lambda(p_1 x_1 + p_2 x_2 - w) = 0 \\ s_1 x_1 = 0 \\ s_2 x_2 = 0 \\ \lambda \geq 0 \\ s_1, s_2 \leq 0 \end{gathered}$$

Since the optimal consumption bundle (x_1, x_2) has to be positive and $s_i x_i = 0$, we have

$$\begin{aligned} x_1 &= \frac{\alpha}{p_1 \lambda} \\ x_2 &= \frac{1-\alpha}{p_2 \lambda} \end{aligned}$$

Therefore $\lambda > 0$. From $\lambda(p_1x_1 + p_2x_2 - w) = 0$, we have $p_1x_1 + p_2x_2 = w$. Using the equations above, we yield

$$x_1 = \frac{\alpha w}{p_1}$$
$$x_2 = \frac{(1-\alpha)w}{p_2}$$

The objective function is concave, hence these conditions are necessary and sufficient for optimality.

c)

$$\lambda = \frac{1}{w}$$

The Lagrangian Multiplier represents the increase in utility that corresponds with a marginal increase in wealth.

d) Plugging in the numbers, we yield the optimal consumption bundle $(x_1, x_2) = (20, 40)$.

Problem 2

Consider the optimization problem

$$\begin{aligned} \min \quad & x_1^2 + 4x_2^2 \\ \text{subject to} \quad & x_1^2 + 2x_2^2 \geq 4 \end{aligned}$$

a) Find all points that satisfy the KKT conditions.

b) Apply the second order condition to determine the whether or not the KKT solutions are local minimizers or maximizers or neither.

Solution:

a) The Lagrangian of this function is

$$L(x_1, x_2, y) = x_1^2 + 4x_2^2 - y(x_1^2 + 2x_2^2 - 4)$$

where $y \geq 0$. The KKT conditions specify that the local optima satisfy the following:

stationarity

$$\nabla L(x_1, x_2, y) = \begin{pmatrix} \frac{\partial L}{\partial x_1} \\ \frac{\partial L}{\partial x_2} \end{pmatrix} = \begin{pmatrix} 2x_1(1-y) \\ 4x_2(2-y) \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

primal feasibility

$$x_1^2 + 2x_2^2 - 4 \geq 0$$

dual feasibility

$$y \geq 0$$

and complementary slackness

$$y(x_1^2 + 2x_2^2 - 4) = 0.$$

The points that satisfy these necessary conditions are

$$\begin{aligned}
&P_1 : y = 1, x_1 = 2, x_2 = 0 \\
&P_2 : y = 1, x_1 = -2, x_2 = 0 \\
&P_3 : y = 2, x_1 = 0, x_2 = \sqrt{2} \\
&P_4 : y = 2, x_1 = 0, x_2 = -\sqrt{2}.
\end{aligned}$$

Note that $x_1 = 0$ and $x_2 = 0$ is not a solutions even though it is a stationary point because it does not satisfy primal feasibility.

b) The Hessian of the Lagrangian is

$$\nabla^2 L(x_1, x_2, y) = \begin{pmatrix} 2(1-y) & 0 \\ 0 & 4(2-y) \end{pmatrix}$$

This matrix is positive semi-definite if and only if $y \leq 1$. Therefore, points P_1 and P_2 are local minima of the Lagrangian and thus of the original program. At points P_3 and P_4, the Hessian is negative semi-definite which makes them local maxima for the unconstrained Lagrangian. Note however that these points are not local maxima if the original program were set to be a maximization problem with the same constraints.

Problem 3

Consider the maze run MDP problem in Problem 4 of Homework 1, also shown here in Figure 1. In Homework 1, we formulated this problem as a linear program where the decision variables are the cost-to-go-values of the decision states.

a) Write down the dual problem of this linear program, solve it using your favorite solver (in teams), and give some interpretations about these dual variables.

b) Give an optimal policy for this Reinforcement Learning problem (Hint: use the dual optimal solution and the complementarity conditions).

Solution:

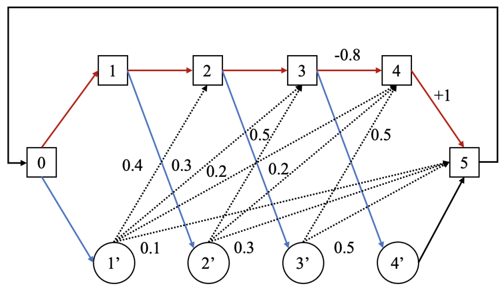

Figure 1: Modified Maze Run

Recall the linear program formulation of the problem:

$$
\begin{aligned}
\max \quad & \sum_{i=0}^{5} y_i \\
\text{subject to} \quad & y_5 \leq \gamma y_0 \\
& y_4 \leq 1 + \gamma y_5 \\
& y_3 \leq -0.8 + \gamma y_4 \\
& y_3 \leq \gamma y_5 \\
& y_2 \leq \gamma y_3 \\
& y_2 \leq \gamma(0.5 y_4 + 0.5 y_5) \\
& y_1 \leq \gamma y_2 \\
& y_1 \leq \gamma(0.5 y_3 + 0.2 y_4 + 0.3 y_5) \\
& y_0 \leq \gamma y_1 \\
& y_0 \leq \gamma(0.4 y_2 + 0.3 y_3 + 0.2 y_4 + 0.1 y_5)
\end{aligned}
$$

where y_i is the cost-to-go for state i for $i = 0, \ldots, 5$ and $\gamma = 0.7$ is the discount factor. Associate to each inequality in the primal problem a Lagrange multiplier

$$
\begin{aligned}
x_5: \quad & y_5 \leq \gamma y_0 \\
x_4: \quad & y_4 \leq 1 + \gamma y_5 \\
x_{3r}: \quad & y_3 \leq -0.8 + \gamma y_4 \\
x_{3b}: \quad & y_3 \leq \gamma y_5 \\
x_{2r}: \quad & y_2 \leq \gamma y_3 \\
x_{2b}: \quad & y_2 \leq \gamma(0.5 y_4 + 0.5 y_5) \\
x_{1r}: \quad & y_1 \leq \gamma y_2 \\
x_{1b}: \quad & y_1 \leq \gamma(0.5 y_3 + 0.2 y_4 + 0.3 y_5) \\
x_{0r}: \quad & y_0 \leq \gamma y_1 \\
x_{0b}: \quad & y_0 \leq \gamma(0.4 y_2 + 0.3 y_3 + 0.2 y_4 + 0.1 y_5)
\end{aligned}
$$

where x_{ir} represents the multiplier for the constraint arising from taking the red action from state i and x_{ib} represents the multiplier for the constraint arising from taking the blue action from state i for $i = 0, \ldots, 3$. States 4 and 5 only have one action so they only have one multiplier each associated to them. These multipliers will be the decision variables of the dual problem.

The primal problem can be cast in standard form:

$$\begin{aligned} \max \quad & c^T y \\ \text{subject to} \quad & Ay \leq b \end{aligned}$$

where

$$c = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \qquad A = \begin{bmatrix} -\gamma & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & -\gamma \\ 0 & 0 & 0 & 1 & -\gamma & 0 \\ 0 & 0 & 0 & 1 & 0 & -\gamma \\ 0 & 0 & 1 & -\gamma & 0 & 0 \\ 0 & 0 & 1 & 0 & -.5\gamma & -.5\gamma \\ 0 & 1 & -\gamma & 0 & 0 & 0 \\ 0 & 1 & 0 & -.5\gamma & -.2\gamma & -.3\gamma \\ 1 & -\gamma & 0 & 0 & 0 & 0 \\ 1 & 0 & -.4\gamma & -.3\gamma & -.2\gamma & -.1\gamma \end{bmatrix}, \qquad b = \begin{bmatrix} 0 \\ 1 \\ -0.8 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Then, the dual of this linear program is

$$\begin{aligned} \min \quad & b^T x \\ \text{subject to} \quad & A^T x = c \\ & x \geq 0 \end{aligned}$$

where

$$x = \begin{bmatrix} x_5 \\ x_4 \\ x_{3r} \\ x_{3b} \\ x_{2r} \\ x_{2b} \\ x_{1r} \\ x_{1b} \\ x_{0r} \\ x_{0b} \end{bmatrix}.$$

Dual variable x_i represents the discounted number of times action i is taken by a policy. The solution to the dual problem is

$$x^* = \begin{bmatrix} 3.3333 \\ 3.3333 \\ 3.3333 \\ 0.0000 \\ 3.3333 \\ 0.0000 \\ 3.3333 \\ 0.0000 \\ 3.3333 \\ 0.0000 \end{bmatrix}$$

The complementarity conditions state that either the constraint is active or the multiplier associated to that constraint is equal to zero. Since the constraints in the linear program came from maximization procedures of different actions, a constraint being active means that the maximum is achieved at the associated action. Therefore, looking at the solution of the dual program, we can determine the optimal policy to be

$$\pi_i = \arg\max\{x_{ir}, x_{ib}\}$$

for $i = 0, 1, 2, 3$, that is,

$$\pi = \begin{bmatrix} \pi_0 \\ \pi_1 \\ \pi_2 \\ \pi_3 \end{bmatrix} = \begin{bmatrix} r \\ r \\ r \\ r \end{bmatrix}$$

where r is the red action and b is the blue action. And at states 4 and 5 there is only one action.

Problem 4

Consider a variant of the Two-Person Zero-Sum Matrix Game in Slides 2-4 of Lecture Note #8, where the payoff matrix becomes:

$$P = \begin{bmatrix} 4 & -1 & -4 & -2 \\ -2 & 1 & 4 & 2 \\ 1 & -2 & 2 & -4 \end{bmatrix}$$

a) Write down the linear program for Player Row.

b) Write down the dual of the above linear program.

c) Give interpretations of the dual problem (with respect to the meaning of the dual variables and dual objective).

Solution:

a)

$$\begin{aligned}
\min \quad & u \\
\text{subject to} \quad & u - (4y_1 - 2y_2 + y_3) \geq 0 \\
& u - (-y_1 + y_2 - 2y_3) \geq 0 \\
& u - (-4y_1 + 4y_2 + 2y_3) \geq 0 \\
& u - (-2y_1 + 2y_2 - 4y_3) \geq 0 \\
& y_1 + y_2 + y_3 = 1 \\
& y_1, y_2, y_3 \geq 0
\end{aligned}$$

b)

$$\begin{aligned}
\max \quad & v \\
\text{subject to} \quad & v - (4x_1 - x_2 - 4x_3 - 2x_4) \leq 0 \\
& v - (-2x_1 + x_2 + 4x_3 + 2x_4) \leq 0 \\
& v - (x_1 - 2x_2 + 2x_3 - 4x_4) \leq 0 \\
& x_1 + x_2 + x_3 + x_4 = 1 \\
& x_1, x_2, x_3, x_4 \geq 0
\end{aligned}$$

c) x_i represents the probability of choosing column i. These probabilities induce an expected payoff to the row players. So v represents the minimum expected payoff among the row players. Therefore, the dual problem maximizes the minimum expected payoff of the row players.

Problem 5

Consider a variant of the Robust Portfolio Management Problem in Slides 19-23 of Lecture Note #8, where constraints on x_1, x_2 become:

$$x_1 + x_2 = 1$$

$$3x_1 - x_2 \geq 0$$

and constraints on μ_1, μ_2 become

$$\mu_1 + 3\mu_2 = 2$$

$$|\mu_1 - \mu_2| \leq 1$$

a) Write its inner problem as a linear program.

b) For fixed x_1 and x_2 (under the constraints of $x_1 + x_2 = 1, 3x_1 - x_2 \geq 0$), find the dual of the inner problem, and solve dual problem.

c) Combine the objectives of the outer and inner problem into a joint single layer problem.

Solution:

a) The inner problem given these constraints is

$$\begin{aligned} \max \quad & -\mu_1 x_1 - \mu_2 x_2 \\ \text{subject to} \quad & \mu_1 + 3\mu_2 = 2 \\ & |\mu_1 - \mu_2| \leq 1 \end{aligned}$$

Note that $|\mu_1 - \mu_2| \leq 1$ is equivalent to having $\mu_1 - \mu_2 \leq 1$ and $\mu_2 - \mu_1 \leq 1$. Thus the inner problem is

$$\begin{aligned} \max \quad & -\mu_1 x_1 - \mu_2 x_2 \\ \text{subject to} \quad & \mu_1 + 3\mu_2 = 2 \\ & \mu_1 - \mu_2 \leq 1 \\ & \mu_2 - \mu_1 \leq 1 \end{aligned}$$

b) The Lagrangian of this problem is

$$\begin{aligned} L(\mu_1, \mu_2, y_1, y_2, y_3) &= -\mu_1 x_1 - \mu_2 x_2 - y_1(\mu_1 + 3\mu_2 - 2) - y_2(\mu_1 - \mu_2 - 1) - y_3(\mu_2 - \mu_1 - 1) \\ &= \mu_1(-x_1 - y_1 - y_2 + y_3) + \mu_2(-x_2 - 3y_1 + y_2 - y_3) + 2y_1 + y_2 + y_3 \end{aligned}$$

where y_1 is free and $y_2, y_3 \geq 0$. The first-order conditions require that

$$\nabla_\mu L(\mu_1, \mu_2, y_1, y_2, y_3) = \begin{bmatrix} -x_1 - y_1 - y_2 + y_3 \\ -x_2 - 3y_1 + y_2 - y_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

and

$$\begin{aligned} y_2(\mu_1 - \mu_2 - 1) = 0 \\ y_3(\mu_2 - \mu_1 - 1) = 0. \end{aligned}$$

Solving for the optimum gives the dual function

$$\max_{\mu_1, \mu_2} L(\mu_1, \mu_2, y_1, y_2, y_3) = \phi(y_1, y_2, y_3) = 2y_1 + y_2 + y_3$$

and the dual problem is

$$\begin{aligned} \min \quad & \phi(y_1, y_2, y_3) = 2y_1 + y_2 + y_3 \\ \text{subject to} \quad & y_1 + y_2 - y_3 = -x_1 \\ & 3y_1 - y_2 + y_3 = -x_2 \\ & y_2, y_3 \geq 0 \end{aligned}$$

Next, we solve this dual problem. First, note that if the primal problem is feasible, a solution will have one of $\mu_1 - \mu_2 - 1 = 0$ or $\mu_2 - \mu_1 - 1 = 0$ but not both since these lines are parallel. Therefore, by the complementary slackness conditions, one of y_2 or y_3 must equal to zero which will depend on the values of x_1 and x_2. From the outer problem, we know that

$x_1 + x_2 = 1$ and $3x_1 - x_2 \geq 0$. This implies that $x_1 \geq 1/4$ and $\frac{x_2}{x_1} \leq 3$. From this, one can show that the level set curves of the primal problem $-x_1\mu_1 - x_2\mu_2 = c$ towards increasing direction lastly intersect with the feasible region at the extreme point at the intersection of $\mu_2 - \mu_1 = 1$ and $\mu_1 + 3\mu_2 = 2$. Therefore, the constraint $\mu_1 - \mu_2 = 1$ is not active and so y_2 must be equal to zero at the optimal solution. If y_2 is zero, then the optimal solution to the dual problem is $y^* = (-x_1/4 - x_2/4, 0, 3x_1/4 - x_2/4)$ with optimal objective value $\phi^* = x_1/4 - 3x_2/4$.

c) Therefore, the combined problem is

$$\begin{aligned} \min \quad & x_1^2 + 2x_2^2 - 2x_1x_2 + \frac{1}{4}x_1 - \frac{3}{4}x_2 \\ \text{subject to} \quad & x_1 + x_2 = 1 \\ & 3x_1 - x_2 \geq 0 \end{aligned}$$

Problem 6

For parts a)-c) below, label them as True or False. If true, provide a short reason; if false, provide reasoning or a counter example.

a) The shadow price of a non-binding constraint can be non-zero.

b) For a LP problem, it is possible that the primal problem has an unbounded objective value, while the dual problem has a non-empty feasible region.

c) For a LP problem, it is possible that the primal problem has a finite optimal value, while the dual problem has no feasible solution.

Solution:

a) False. Complementary slackness requires that at least the shadow price is zero or its associated constraint is binding.

b) False. The weak duality theorem implies that if the primal problem has unbounded objective value, then the dual problem is infeasible.

c) False. The strong duality theorem implies that if the primal problem is feasible and bounded, then the dual is also feasible and bounded.