ORF 418 - Optimal Learning - Fall 2025
H. Mete Soner

Homework assignment 2 – Solution

Exercise 1. LQ problem [10pt]

Consider the general Linear-Quadratic problem seen in class (or Section 2.1.1 of the lecture notes), with infinite horizon and the following parameters: $d = 2$, $\ell = 1$, $N = 1$, $\rho = 1/2$ and

$$A := \begin{pmatrix} 0 & 1 \\ -2 & 0 \end{pmatrix}, \; B := \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \; M := \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

In the following, we will denote by $x_k := (y_k, z_k)^\top \in \mathbb{R}^2$ the state of the system at time $k \in \mathbb{N}$, with initial condition $x_0 \in \mathbb{R}^2$ known. We also denote by $v(x)$ the value function associated to this LQ problem, for any $x \in \mathbb{R}^2$.

1. **[7pt]** *Compute the controllability* matrix and decide whether this system is *controllable* or not.

   Solution. The controllability matrix is

   $$\mathcal{C} = \begin{bmatrix} B & AB \end{bmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

   which has rank 2**. Hence the pair** (A, B) **is controllable.**

2. **[10pt]** We know that $v(y, z) = ay^2 + bz^2$ for some constants a, b. Use the dynamic programming equation to *compute* the constants a, b.

   Solution. The discounted DPE is

   $$v(x) = x^\top M x + \inf_{u \in \mathbb{R}} \Big\{ u^2 + \rho\, v(Ax + Bu) \Big\}, \qquad \rho = \tfrac{1}{2}. \tag{1.1}$$

   Assume a quadratic form $v(x) = x^\top V x$ **with**

   $$V = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, \qquad a, b > 0.$$

   Writing $x = (y, z)^\top$**, we have** $Ax + Bu = (z,\, -2y + u)^\top$**. Substituting into (??) gives**

   $$ay^2 + bz^2 = y^2 + 2z^2 + \inf_{u \in \mathbb{R}} \Big\{ u^2 + \tfrac{1}{2}\big(az^2 + b(-2y + u)^2\big) \Big\}.$$

   Define the convex function

   $$f(u) := u^2 + \tfrac{1}{2}\big(az^2 + b(-2y + u)^2\big), \qquad f''(u) = 2 + b > 0 \;\; \textbf{if } b > -2.$$

   The first-order condition $f'(u) = (2 + b)u - 2by = 0$ **yields the unique minimizer**

   $$u^\star = \frac{2by}{2 + b}.$$

   A short calculation yields the minimal value

   $$\inf_{u \in \mathbb{R}} f(u) = \tfrac{1}{2}\, az^2 + \frac{4b}{2 + b}\, y^2. \tag{1.2}$$

Plugging (??) back gives the identity (for all y, z)

$$\big(a - 1 - \tfrac{4b}{2+b}\big)y^2 \;+\; \big(b - 2 - \tfrac{a}{2}\big)z^2 \;=\; 0,$$

hence the coefficient system

$$a - 1 - \frac{4b}{2+b} = 0, \qquad b - 2 - \frac{a}{2} = 0. \tag{1.3}$$

From the second equation, $a = 2b - 4$**. Substituting into the first:**

$$(2b-5) - \frac{4b}{2+b} = 0 \iff (2b-5)(2+b) = 4b \iff 2b^2 - 5b - 10 = 0.$$

Thus

$$b = \frac{5+\sqrt{105}}{4} \approx 3.8118, \qquad a = 2b - 4 = \frac{-3+\sqrt{105}}{2} \approx 3.6235,$$

and $V = \operatorname{diag}(a, b) \succ 0$**. Therefore** $v(x) = x^\top V x$ **solves the DPE.**

3. **[10pt]** Consider the following minimization problem:

$$\widetilde{v}(r_0, r_1) := \inf_u \sum_{k=0}^{\infty} \frac{1}{2^k}\big(r_k^2 + 2r_{k+1}^2 + u_k^2\big),$$

where $r_{k+2} + 2r_k = u_k$ for $k = 0, 1, \dots,$ and the initial conditions r_0, r_1 are given. *Formulate* this problem as a Linear Quadratic problem, *i.e.* write down A, B, M, N and ρ.

Solution. Consider

$$\tilde{v}(r_0, r_1) := \inf_{\{u_k\}} \sum_{k=0}^{\infty} \frac{1}{2^k}\Big(r_k^2 + 2r_{k+1}^2 + u_k^2\Big), \qquad r_{k+2} + 2r_k = u_k.$$

Let $x_k := (r_k, r_{k+1})^\top$**; then**

$$x_{k+1} = \begin{pmatrix} 0 & 1 \\ -2 & 0 \end{pmatrix} x_k + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u_k = Ax_k + Bu_k,$$

and the stage cost is $x_k^\top M x_k + u_k^\top N u_k$ **with the same** (M, N) **as above and discount** $\rho = \frac{1}{2}$**. Hence** $\tilde{v}$ **is exactly the same discounted LQ with data** (A, B, M, N, ρ)**.**

4. **[3pt]** *Compute* $v(1, 2)$.

Since $v(x) = x^\top V x$ **with** $V = \operatorname{diag}(a, b)$ **as above,**

$$v(1,2) = \begin{pmatrix} 1 & 2 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} \begin{pmatrix} 1 \\ 2 \end{pmatrix} = a + 4b = \frac{-3+\sqrt{105}}{2} + 4 \cdot \frac{5+\sqrt{105}}{4} = \frac{7+3\sqrt{105}}{2} \approx 18.871.$$

Exercise 2. [30pt]

Suppose that $Y_1, \dots, Y_n$ are i.i.d. from the *Poisson distribution* with mean $\theta > 0$, i.e., for any $k = 1, \dots,$ conditioned on the value of θ,

$$\mathbb{P}(Y_k = y \,|\, \theta) \;=\; \frac{\theta^y}{y!}\, e^{-\theta}, \qquad y = 0, 1, \dots$$

The parameter θ is not known and hence, is a random variable. Suppose that the *prior distribution of θ is the gamma distribution* with parameters $a, b > 0$, i.e., the p.d.f. is given by,

$$f_\theta(\theta) = \frac{b^a}{\Gamma(a)}\, \theta^{a-1}\, e^{-b\,\theta}, \qquad \theta > 0,$$

where Γ is the *gamma function*:

$$\Gamma(a) := \int_0^\infty z^{a-1}\, e^{-z}\, dz.$$

In this exercise you do not need the properties of the Γ function.

1. Show that the posterior distribution of θ based on the observation that $Y_1 = y_1$ is again gamma-distributed with parameters,

$$a + y_1, \qquad b + 1.$$

Solution. By Bayes' Theorem, we have

$$\begin{aligned} f_{\theta|Y}(\theta \mid Y_1 = y_1) &= \frac{\mathbb{P}(\,Y_1 = y_1 \ \mid\ \theta)\, f_\theta(\theta)}{\mathbb{P}(\,Y_1 = y_1\,)} \\ &= C\,\mathbb{P}(Y_1 = y_1 \mid \theta) f_\theta(\theta), \end{aligned}$$

where C a suitable normalization constant that makes $f_{\theta|Y}$ a pdf, i.e., C is the unique constant so that

$$\int f_{\theta|Y}(\theta \mid Y_1 = y_1)\ \mathrm{d}\theta = 1.$$

Given that $Y \mid \theta \sim$ Poi(θ) and the prior $\theta \sim$ Gamma(a, b), we have

$$\begin{aligned} f_{\theta|Y}(\theta \mid Y_1 = y_1) &= C\ (\frac{\theta^{y_1}}{y_1!} e^{-\theta})\left(\theta^{a-1} e^{-b\theta}\right) \\ &= C'\ \theta^{a+y_1-1} e^{-(b+1)\theta}, \qquad \theta > 0, \end{aligned}$$

where C' absorbs all multiplicative factors that don't depend on θ. By matching the terms to the known form of the gamma distribution, we see that $\theta \mid Y_1 = y_1 \sim$ Gamma$(a + y_1, b + 1)$, as desired.

2. By iterating part a, show that the posterior distribution of θ based on the observations $Y_1 = y_1$, $Y_2 = y_2, \ldots, Y_n = y_n$ is again gamma-distributed with parameters,

$$a + \sum_{i=1}^n y_i, \qquad b + n.$$

Solution. Proceed by induction on n. The base case $n = 1$ is shown in part a. Assuming the inductive hypothesis holds for $n - 1$ observations, that is

$$\theta \mid Y_1 = y_1, \cdots Y_{n-1} = y_{n-1} \sim \textbf{Gamma}(a + \sum_{i=1}^{n-1} y_i, b + n - 1)$$

Since the observations are i.i.d., we can apply Bayes Theorem using the posterior distribution after the first $n - 1$ observations as the prior before including the n-th observation. This gives

$$\begin{aligned} f_{\theta|Y}(\theta \mid Y_1 = y_1, \cdots, Y_n = y_n) &= C\ (\frac{\theta^{y_n}}{y_n!} e^{-\theta})\ (\theta^{[a+\sum_{i=1}^{n-1} y_1 - 1]}\ e^{-(b+n-1)}) \\ &= C'\ \theta^{a+\sum_{i=1}^{n} y_i - 1}\ e^{-(b+n)\theta} \end{aligned}$$

Again, we recognize this form as the pdf of a Gamma distributed random variable. So, by induction, we have

$$\theta \mid Y_1 = y_1, \cdots Y_n = y_n \sim \textbf{Gamma}(a + \sum_{i=1}^{n} y_i, b + n)$$

as desired.

Exercise 3. Bayesian estimation and bias [40pt]

The first part of this exercise is theoretical, while the second part (from question 5 onwards) involves numerical implementation. You do not have to upload your code, you should just answer the questions by giving appropriate numerical values and relevant graphs.

This exercise is related to a court case *Castaneda vs. Partida* 430 U.S. 482 (1977) deciding whether there was bias in juror selection in a particular county. The local population of the given county was 79% Mexican-American. However, over an 11-year period, among the 870 persons summoned to serve as grand jurors, only 339 were Mexican-American.

Court assumed that the selection of jurors was random with P being the probability of any selected juror being Mexican-American. If Y denotes the number of Mexican-American jurors, then Y is a Binomial random variable with $n = 870$ and success probability P. Namely, the conditional distribution of Y given $P = p$ is:

$$\mathbb{P}(Y = y \mid P = p) = \binom{n}{y} p^y (1-p)^{n-y}, \qquad \text{for } y = 0, 1, \ldots, n.$$

Based on many scientific experiments, the Court assumed that the distribution of P prior to the observation is a Beta distribution with parameters a and b.

1. **[4pt]** Compute $\mathbb{E}[Y \mid P = 0.79]$. By comparing this value with the observation, comment on potential bias in the jury selection procedure.

   Solution. The expectation of a Binomial with parameter n and success probability P is given by $n \times P$. Therefore, $\mathbb{E}[Y \mid P = 0.79] = 870 \times 0.79 = 687.3$. This value is substantially bigger than the actually observed value of $Y = 339$. Therefore, based on the observation $Y = 339$, one might guess that there is bias in the jury selection procedure in this county.

2. **[6pt]** Write the (prior) density function of P and give the formula for its expectation.

   Solution. The prior distribution of P is a Beta distribution with parameters a and b. Therefore, its density is given by:

   $$f_P(p) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1}, \qquad p \in (0,1).$$

   Its expectation is $\mathbb{E}[P] = a/(a+b)$.

3. **[10pt]** Show that the posterior distribution of P, given an observation $Y = y$ for $y \in \{0, \ldots, n\}$ satisfies the following, for some function C independent of p (*you do not need to compute the function C*):

   $$f_{P|Y=y}(p) = C(y) p^{a+y-1}(1-p)^{b+n-y-1}, \qquad p \in (0,1).$$

   Solution. Given the framework of the problem, we should use the following (mixed) Bayes' formula:

   $$f_{P|Y=y}(p) = \frac{\mathbb{P}(Y = y \mid P = p) f_P(p)}{\mathbb{P}(Y = y)}.$$

We know that for $y = 0, 1, \ldots, n$ and $p \in (0,1)$,

$$\mathbb{P}(Y = y \mid P = p) = \binom{n}{y} p^y (1-p)^{n-y}, \quad \textbf{and} \quad f_P(p) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1}.$$

Replacing in the previous formula, we obtain for $y = 0, 1, \ldots, n$ and $p \in (0,1)$:

$$\begin{aligned} f_{P|Y=y}(p) &= \frac{1}{\mathbb{P}(Y=y)} \binom{n}{y} p^y (1-p)^{n-y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1} \\ &= \frac{1}{\mathbb{P}(Y=y)} \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a+y-1}(1-p)^{b+n-y-1}. \end{aligned}$$

Defining the function C as:

$$C(y) := \frac{1}{\mathbb{P}(Y=y)} \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)},$$

which is indeed independent of p, we obtain the required equality for $y = 0, 1, \ldots, n$ and $p \in (0,1)$:

$$f_{P|Y=y}(p) = C(y) p^{a+y-1}(1-p)^{b+n-y-1}.$$

4. **[10pt]** Compute that

$$C(y) = \frac{\Gamma(a+b+n)}{\Gamma(a+y)\Gamma(b+n-y)},$$

and conclude that the posterior distribution of P, given the observation $Y = 339$, is again a Beta distribution with parameters $a + 339$ and $b + 531$.

Solution. Given the result of the previous question, we have, for $y = 0, 1, \ldots, n$ and $p \in (0,1)$,

$$f_{P|Y=y}(p) = C(y) p^{a+y-1}(1-p)^{b+n-y-1},$$

with

$$C(y) := \frac{1}{\mathbb{P}(Y=y)} \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}.$$

To compute C, we first compute

$$\begin{aligned} \mathbb{P}(Y=y) &= \int_0^1 \mathbb{P}(Y=y, P=p)\mathrm{d}p = \int_0^1 \mathbb{P}(Y=y \mid P=p) f_P(p)\mathrm{d}p \\ &= \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \int_0^1 p^y(1-p)^{n-y} p^{a-1}(1-p)^{b-1}\mathrm{d}p \\ &= \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \int_0^1 p^{a+y-1}(1-p)^{b+n-y-1}\mathrm{d}p. \end{aligned}$$

Since the integral of any density is equal to 1, we have:

$$\frac{\Gamma(a+b+n)}{\Gamma(a+y)\Gamma(b+n-y)} \int_0^1 p^{a+y-1}(1-p)^{b+n-y-1}\mathrm{d}p = 1,$$

and therefore

$$\int_0^1 p^{a+y-1}(1-p)^{b+n-y-1}\mathrm{d}p = \frac{\Gamma(a+y)\Gamma(b+n-y)}{\Gamma(a+b+n)}.$$

Replacing in the computations for $\mathbb{P}(Y=y)$, we have:

$$\mathbb{P}(Y=y) = \binom{n}{y}\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\frac{\Gamma(a+y)\Gamma(b+n-y)}{\Gamma(a+b+n)}.$$

Using this in the definition of the function C, we obtain the desired result for C, namely:

$$C(y) := \frac{1}{\mathbb{P}(Y=y)}\binom{n}{y}\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} = \frac{\binom{n}{y}\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}}{\binom{n}{y}\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\frac{\Gamma(a+y)\Gamma(b+n-y)}{\Gamma(a+b+n)}} = \frac{\Gamma(a+b+n)}{\Gamma(a+y)\Gamma(b+n-y)}.$$

Using this formula for $C(y)$, we obtain for all $y=0,1,\ldots,n$ and $p \in (0,1)$:

$$f_{P|Y=y}(p) = \frac{\Gamma(a+b+n)}{\Gamma(a+y)\Gamma(b+n-y)}p^{a+y-1}(1-p)^{b+n-y-1}.$$

Denoting by $\tilde{a} = a+y$ and $\tilde{b} = b+n-y$ and noticing that $\tilde{a}+\tilde{b} = a+y+b+n-y = a+b+n$, we can write, for $p \in (0,1)$,

$$f_{P|Y=y}(p) = \frac{\Gamma(\tilde{a}+\tilde{b})}{\Gamma(\tilde{a})\Gamma(\tilde{b})}p^{\tilde{a}-1}(1-p)^{\tilde{b}-1}.$$

We recognise the density of a Beta distribution with parameters $\tilde{a} = a+y$ and $\tilde{b} = b+n-y$. Therefore, the posterior distribution of P, given the observation $Y=y$, is a Beta distribution with parameters $\tilde{a} = a+y$ and $\tilde{b} = b+n-y$. In this example, we observe $Y = 339$, and recalling that $n = 870$, we obtain that the posterior distribution of P, given the observation $Y = 339$, is a Beta distribution with parameters $\tilde{a} = a+339$ and $\tilde{b} = b+870-339 = b+531$.

In the following, as the prior must be consistent with the population, we set $\mathbb{E}[P] = 0.79$, or equivalently $a = 79b/21$. Mathematically, the bias in juror selection would mean that the random variable P is substantially smaller than 0.79. Say we define the bias to be P being 80% or less than the population percentage of 79%. In other words, the event $P \leq 0.8 \times 0.79 = 0.632$ means bias.

5. **[4pt]** Remark that as $a = 79b/21$, the conditional probability $\mathbb{P}(P \leq 0.632 \mid Y = 339)$ is only a function b. Numerically compute and plot this probability as a function of $b \in [0.1, n]$ with step size 0.1.
6. **[4pt]** Determine the values of b for which the previous probability is less than 0.5. For the minimum admissible value of b, give the corresponding value of a and the resulting parameters for the posterior distribution.
7. **[2pt]** Numerically compute and plot the unconditional probability $\mathbb{P}(P \leq 0.632)$ as a function of $b \in [0.1, n]$ with step size 0.1. What do you obtain for the admissible values for b found in the previous question? Comment briefly.

Solution. The solution of Questions 5 to 7 can be found in the associated python file.

One additional remark can be made about the results. The calculations prove that in order to believe that there is no-bias, initially one must give almost zero probability to the possibility that there is bias. By playing a bit with the code, one can see that if initially one accepts even with a small probability that there is bias, then after the observation of $Y = 339$ the probability of bias is overwhelming.