STK2130 – Week 9

A. B. Huseby

Department of Mathematics
University of Oslo, Norway

Chapter 4.8 Time Reversible Markov Chains

Consider an ergodic Markov chain with transition probabilities P_{ij} and stationary probabilities π_i, $i, j \in \mathcal{S}$.

Then let n be so large that we have reached a stationary state, i.e. $P_{ij}^n \approx \pi_j$.

We then consider the backwards chain $X_n, X_{n-1}, X_{n-2}, \ldots$

The backwards chain is also a Markov chain with transitions probabilities Q_{ij}, $i, j \in \mathcal{S}$ given by:

$$Q_{ij} = P(X_m = j \mid X_{m+1} = i) = \frac{P(X_m = j \cap X_{m+1} = i)}{P(X_{m+1} = i)}$$

$$= \frac{P(X_m = j)P(X_{m+1} = i \mid X_m = j)}{P(X_{m+1} = i)} = \frac{\pi_j P_{ji}}{\pi_i}.$$

Time Reversible Markov Chains (cont.)

We say that $\{X_n\}$ is time reversible if $Q_{ij} = P_{ij}$ for all $i, j \in \mathcal{S}$.

Hence, $\{X_n\}$ is time reversible if and only if:

$$\frac{\pi_j P_{ji}}{\pi_i} = P_{ij}, \quad \text{for all } i, j \in \mathcal{S}.$$

or equivalently if and only if:

$$\pi_i P_{ij} = \pi_j P_{ji}, \quad \text{for all } i, j \in \mathcal{S}.$$

Doubly stochastic transition matrices

Let $\boldsymbol{P}$ be the transition probability matrix of a Markov chain $\{X_n\}$ with state space $\mathcal{S}$ where $|\mathcal{S}| = n < \infty$. We then know that:

$$\sum_{j \in \mathcal{S}} P_{ij} \cdot 1 = 1, \quad \text{for all } i \in \mathcal{S}.$$

If $\boldsymbol{e} = (1, \ldots, 1)$ is an n-dimensional vector, we can rewrite this as:

$$\boldsymbol{P}\boldsymbol{e}^T = \boldsymbol{e}^T.$$

We say that $\boldsymbol{P}$ is doubly stochastic if we also have:

$$\sum_{i \in \mathcal{S}} P_{ij} \cdot 1 = 1, \quad \text{for all } j \in \mathcal{S}.$$

In matrix notation this condition can be expressed as:

$$\boldsymbol{e}\boldsymbol{P} = \boldsymbol{e}.$$

Doubly stochastic transition matrices (cont.)

Proposition

Let $\mathbf{P}$ be the transition probability matrix of a Markov chain $\{X_n\}$ with state space $\mathcal{S}$ where $|\mathcal{S}| = n < \infty$. Then $\mathbf{P}$ is doubly stochastic if and only if the stationary distribution is uniform, i.e.: $\pi = \frac{1}{n}\mathbf{e}$.

PROOF:

$$\mathbf{P} \text{ is doubly stochastic}$$
$$\Updownarrow$$
$$\mathbf{e}\mathbf{P} = \mathbf{e}$$
$$\Updownarrow$$
$$\frac{1}{n}\mathbf{e}\mathbf{P} = \frac{1}{n}\mathbf{e}$$
$$\Updownarrow$$
$$\pi\mathbf{P} = \pi$$
$$\Updownarrow$$

The stationary distribution is uniform

Doubly stochastic transition matrices (cont.)

Proposition

Let $\boldsymbol{P}$ be the transition probability matrix of a Markov chain $\{X_n\}$ with state space $\mathcal{S}$ where $|\mathcal{S}| = n < \infty$. Assume that $\boldsymbol{P}$ is doubly stochastic. Then the Markov chain is time reversible if and only if $\boldsymbol{P}$ is symmetric.

PROOF: Since $\boldsymbol{P}$ is doubly stochastic $\pi_i = \frac{1}{n}$ for all $i \in \mathcal{S}$. Hence, we get:

$$Q_{ij} = \frac{\pi_j P_{ji}}{\pi_i} = \frac{\frac{1}{n} P_{ji}}{\frac{1}{n}} = P_{ji} \quad i, j \in \mathcal{S}$$

Thus, $\boldsymbol{Q} = \boldsymbol{P}$ if and only if $P_{ji} = P_{ij}$ for all $i, j \in \mathcal{S}$.

That is, $\boldsymbol{Q} = \boldsymbol{P}$ if and only if $\boldsymbol{P}$ is symmetric.

Example

Let $\boldsymbol{P}$ be the transition probability matrix of a Markov chain $\{X_n\}$ with state space $\mathcal{S} = \{0, 1, 2\}$, where:

$$\boldsymbol{P} = \begin{bmatrix} 0.2 & 0.4 & 0.4 \\ 0.4 & 0.2 & 0.4 \\ 0.4 & 0.4 & 0.2 \end{bmatrix}$$

Since $\boldsymbol{P}$ is doubly stochastic and symmetric, the Markov chain is time reversible.

Example

Let $\boldsymbol{P}$ be the transition probability matrix of a Markov chain $\{X_n\}$ with state space $\mathcal{S} = \{0, 1, 2\}$, where:

$$\boldsymbol{P} = \begin{bmatrix} 0.0 & 0.1 & 0.9 \\ 0.9 & 0.0 & 0.1 \\ 0.1 & 0.9 & 0.0 \end{bmatrix}$$

Since $\boldsymbol{P}$ is doubly stochastic and not symmetric, the Markov chain is not time reversible.

Time Reversible Markov Chains (cont.)

Theorem (4.2)

A stationary Markov chain for which $P_{ij} = 0$ whenever $P_{ji} = 0$ is time reversible if and only if starting in state i, any path back to i has the same probability as the reversed path. That is, if:

$$P_{i,i_1} P_{i_1,i_2} \cdots P_{i_k,i} = P_{i,i_k} P_{i_k,i_{k-1}} \cdots P_{i_1,i}$$

for all states $i, i_1, \ldots, i_k$, $k = 1, 2, \ldots$.

Time Reversible Markov Chains (cont.)

Proposition (4.9)

Consider an irreducible Markov chain with transition probability matrix $\boldsymbol{P}$ *and state space* $\mathcal{S}$*. If we can find a stationary distribution* $\boldsymbol{\pi}$ *on* $\mathcal{S}$*, and a transition probability matrix* $\boldsymbol{Q}$ *such that:*

$$\pi_i P_{ij} = \pi_j Q_{ji}, \qquad \text{for all } i, j \in \mathcal{S},$$

then $\boldsymbol{Q}$ *is the transition probability matrix of the reversed chain, and* $\boldsymbol{\pi}$ *is the stationary distribution for both the original and reversed chain.*

PROOF: The result is an immediate consequence of the previously established formula:

$$Q_{ji} = \frac{\pi_i P_{ij}}{\pi_j}, \qquad \text{for all } i, j \in \mathcal{S}.$$

Example 4.40

Let L denote the lifetime of a type of light bulb, and assume that the distribution of L in days is:

$$P(L = i) = p_i, \quad i = 1, 2, \ldots$$

We have an infinite supply of light bulb of this kind, with lifetimes $L_1, L_2, \ldots$ being independent and with the same distribution as L. Each time a bulb fails, it is replaced by a new one the next day.

$$X_n = \text{The age of the light bulb at day } n, \quad n = 1, 2, \ldots$$

Then $\{X_n\}$ is a Markov chain with state space $\mathcal{S} = \{1, 2, \ldots\}$ and with transition probabilities:

$$P_{i,1} = \frac{P(L = i)}{P(L \geq i)} = \frac{p_i}{\sum_{j=i}^{\infty} p_j}, \quad P_{i,i+1} = 1 - P_{i,1}, \quad i = 1, 2, \ldots$$

Example 4.40 (cont.)

We then claim that the reversed chain has transition probabilities:

$$Q_{i,i-1} = 1, \quad i > 1$$
$$Q_{1,i} = P(L = i) = p_i, \quad i \geq 1$$

To show this we need to find a stationary distribution $\pi_1, \pi_2, \ldots$ such that:

$$\pi_i P_{ij} = \pi_j Q_{ji}, \quad i, j = 1, 2, \ldots$$

We start out by letting $j = 1$ and determine $\pi_1, \pi_2, \ldots$ such that:

$$\pi_i P_{i,1} = \pi_i \frac{P(L = i)}{P(L \geq i)} = \pi_1 Q_{1,i} = \pi_1 P(L = i).$$

This is equivalent to:

$$\pi_i = \pi_1 P(L \geq i)$$

Example 4.40 (cont.)

Summing over all i and using that the π_i-s add up to 1, we get:

$$\begin{aligned}
1 = \sum_{i=1}^{\infty} \pi_i &= \pi_1 \sum_{i=1}^{\infty} P(L \geq i) \\
&= \pi_1 \sum_{i=1}^{\infty} \sum_{j=i}^{\infty} p_j = \pi_1 \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} p_j I(j \geq i) \\
&= \pi_1 \sum_{j=1}^{\infty} \sum_{i=1}^{\infty} p_j I(i \leq j) = \pi_1 \sum_{j=1}^{\infty} p_j \sum_{i=1}^{\infty} I(i \leq j) \\
&= \pi_1 \sum_{j=1}^{\infty} p_j \cdot j = \pi_1 E[L].
\end{aligned}$$

Hence, it follows that $\pi_1 = (E[L])^{-1}$, and thus:

$$\pi_i = \pi_1 P(L \geq i) = \frac{P(L \geq i)}{E[L]}, \quad i = 1, 2, \ldots$$

Example 4.40 (cont.)

We note that if $j > 1$, we have $P_{ij} = Q_{ji} = 0$ for $j \neq i + 1$. Hence, it remains to verify that:

$$\pi_i P_{ij} = \pi_j Q_{ji}, \quad i = 1, 2, \ldots, \quad j = i + 1.$$

Using the expressions for $\pi_1, \pi_2, \ldots$ and that $Q_{i+1,i} = 1$ for $i > 1$, this is equivalent to:

$$\frac{P(L \geq i)}{E[L]} \left(1 - \frac{P(L = i)}{P(L \geq i)}\right) = \frac{P(L \geq i + 1)}{E[L]}$$

By simplifying this equation we get:

$$P(L \geq i) - P(L = i) = P(L \geq i + 1)$$

which is trivially true.

Chapter 4.9 Markov Chain Monte Carlo Methods

Let Z be a discrete random variable with a state space $\mathcal{S}$, and assume that:

$$P(Z = i) = \pi_i = \frac{b_i}{B}, \quad i \in \mathcal{S}.$$

We assume that b_i is known for all $i \in \mathcal{S}$.

Since the probabilities must add up to 1, we obviously have:

$$\sum_{i \in \mathcal{S}} \frac{b_i}{B} = B^{-1} \sum_{i \in \mathcal{S}} b_i = 1,$$

Hence, it follows that the normalizing constant B is given by:

$$B = \sum_{i \in \mathcal{S}} b_i.$$

Thus, in principle B is known as well. However, if $|\mathcal{S}|$ is large, calculating B may be a time-consuming task.

MCMC Methods (cont.)

EXAMPLE: Let T and Z be two discrete random variables with state spaces $\mathcal{T}$ and $\mathcal{S}$ respectively. We assume that the marginal distribution of Z and the conditional distribution of T given Z are known.

The conditional distribution of Z given T is then:

$$P(Z = i|T = t) = \frac{P(Z = i)P(T = t|Z = i)}{\sum_{j \in \mathcal{S}} P(Z = j)P(T = t|Z = j)} = \frac{b_i(t)}{B(t)}, \quad i \in \mathcal{S},\ t \in \mathcal{T},$$

where we have introduced:

$$b_i(t) = P(Z = i)P(T = t|Z = i), \quad i \in \mathcal{S},\ t \in \mathcal{T},$$

$$B(t) = \sum_{j \in \mathcal{S}} P(Z = j)P(T = t|Z = j) = P(T = t), \quad t \in \mathcal{T}.$$

If $|\mathcal{S}|$ is large, we may want to avoid calculating $B(t)$.

MCMC Methods (cont.)

PROBLEM: Construct a Markov chain $\{X_n\}$ with state space $\mathcal{S}$ and stationary distribution equal to the distribution of Z.

SOLUTION (Hastings-Metropolis): Let $\boldsymbol{Q}$ be any given irreducible Markov chain transition probability matrix on $\mathcal{S}$, and define:

$$\alpha_{ij} = \min\left(\frac{b_j Q_{ji}}{b_i Q_{ij}}, 1\right), \quad i, j \in \mathcal{S}.$$

We then let the transition probability matrix of $\{X_n\}$, denoted $\boldsymbol{P}$, be defined as follows:

$$P_{ij} = Q_{ij}\alpha_{ij}, \quad i \neq j,$$

$$P_{ii} = 1 - \sum_{j \neq i} Q_{ij}\alpha_{ij}, \quad i \in \mathcal{S}$$

MCMC Methods (cont.)

We then claim that $\{X_n\}$ is time reversible and have a stationary distribution equal to the distribution of Z.

To show this it is sufficient to verify that:

$$\pi_i P_{ij} = \pi_j P_{ji}, \quad \text{for all } i, j \in \mathcal{S}. \tag{1}$$

Since (1) is trivially satisfied for $i = j$, we focus on the case where $i \neq j$, where (1) can be expressed as:

$$\frac{b_i}{B} Q_{ij}\alpha_{ij} = \frac{b_j}{B} Q_{ji}\alpha_{ji}, \quad i \neq j. \tag{2}$$

By eliminating B from these equations and inserting the expression for α_{ij} we get:

$$b_i Q_{ij} \cdot \min\left(\frac{b_j Q_{ji}}{b_i Q_{ij}}, 1\right) = b_j Q_{ji} \cdot \min\left(\frac{b_i Q_{ij}}{b_j Q_{ji}}, 1\right), \quad i \neq j. \tag{3}$$

MCMC Methods (cont.)

CASE 1: $b_i Q_{ij} \leq b_j Q_{ji}$

In this case $\alpha_{ij} = 1$ while $\alpha_{ji} = (b_i Q_{ij})/(b_j Q_{ji})$, and hence, (3) simplifies to:

$$b_i Q_{ij} = b_j Q_{ji} \cdot (b_i Q_{ij})/(b_j Q_{ji}), \quad i \neq j. \tag{4}$$

CASE 2: $b_i Q_{ij} \geq b_j Q_{ji}$

In this case $\alpha_{ij} = (b_j Q_{ji})/(b_i Q_{ij})$ while $\alpha_{ji} = 1$, and hence, (3) simplifies to:

$$b_i Q_{ij} \cdot (b_j Q_{ji})/(b_i Q_{ij}) = b_j Q_{ji}, \quad i \neq j. \tag{5}$$

Since obviously both (4) and (5) hold true, we conclude that (1) holds true as well.

MCMC Methods (cont.)

We recall that:

$$P_{ij} = Q_{ij}\alpha_{ij}, \quad i \neq j,$$

$$P_{ii} = 1 - \sum_{j \neq i} Q_{ij}\alpha_{ij}, \quad i \in \mathcal{S}$$

Assume that $X_n = i$. Then X_{n+1} can be generated using the following two-step Monte Carlo simulation procedure:

STEP 1. Generate $J \in \mathcal{S}$ such that $P(J = j) = Q_{ij}$, $j \in \mathcal{S}$.

STEP 2. Generate $B \in \{0, 1\}$ such that $P(B = 1 \mid J = j) = \alpha_{ij}$, and let:

$$X_{n+1} = B \cdot j + (1 - B) \cdot i$$

Thus, a transition from state i to state j where $i \neq j$ happens if and only if $J = j$ and $B = 1$. If not, the process stays in state i.

MCMC Methods (cont.)

The Monte Carlo simulation procedure can be used to estimate some unknown parameter in the distribution of Z, e.g.:

$$\theta = E[h(Z)] = \sum_{i \in \mathcal{S}} h(i) P(Z = i),$$

where h is some function of interest, and the normalizing constant B of the distribution of Z is too time-consuming to calculate.

By simulating the Markov chain $\{X_n\}$, having a stationary distribution which is equal to the distribution of Z, we may estimate θ by:

$$\hat{\theta}_n = \frac{1}{n} \sum_{m=1}^{n} h(X_m).$$

By the law of large numbers it follows that $\hat{\theta}_n \to \theta$ when $n \to \infty$.

MCMC Methods (cont.)

NOTE: $X_1, X_2, \ldots$ are not independent samples.

Moreover, the chain may converge slowly towards its stationary distribution.

Both these issues tend to have a negative effect on the convergence rate of the estimator $\hat{\theta}_n$.

If many of the α_{ij}-s are small, the Markov chain tends to get stuck for a long time before eventually transiting to another state. In such cases the estimator $\hat{\theta}_n$ will converge very slowly.

MCMC Methods (cont.)

For optimal performance, i.e., fast convergence, the matrix $\boldsymbol{Q}$ should ideally be chosen so that:

$$b_i Q_{ij} = b_j Q_{ji}, \qquad \text{for all } i, j \in \mathcal{S}.$$

Then it follows that:

$$\alpha_{ij} = \min\left(\frac{b_j Q_{ji}}{b_i Q_{ij}}, 1\right) = 1, \qquad \text{for all } i, j \in \mathcal{S}.$$

Hence, $\boldsymbol{Q} = \boldsymbol{P}$, i.e., $\boldsymbol{Q}$ is itself the transition probability matrix of $\{X_n\}$.

Finding the optimal matrix $\boldsymbol{Q}$ implies finding a transition probability matrix with a stationary distribution which is equal to the distribution of Z. In real-life applications, this can be difficult.

Instead we may think of $\boldsymbol{Q}$ as our best guess, while the α_{ij}-s are correction factors which are used to generate a Markov chain with the correct stationary distribution.

Gibbs sampling

Assume that $\boldsymbol{Z} = (Z_1, \ldots, Z_r)$ is a discrete random vector with values in $\mathcal{S}$ where:

$$P(\boldsymbol{Z} = \boldsymbol{z}) = p(\boldsymbol{z}) = g(\boldsymbol{z})/B, \qquad \text{for all } \boldsymbol{z} \in \mathcal{S},$$

where the $g(\boldsymbol{z})$ is known for all $\boldsymbol{z} \in \mathcal{S}$ and B is an unknown normalizing constant.

We then consider the first step of the Hastings-Metropolis algorithm, and assume that $X_n = \boldsymbol{z} = (z_1, \ldots, z_r)$. The candidate for the next state, X_{n+1}, is generated as follows:

1. Generate $K = k$ uniformly from the set $\{1, \ldots, r\}$.
2. Generate $Z_k = z$ conditional on $Z_i = z_i$, $i = 1, \ldots, (k-1), (k+1) \ldots, r$.

The resulting candidate for the next state, denoted $\boldsymbol{y}$, is then:

$$\boldsymbol{y} = (z_1, \ldots, z_{k-1}, z, z_{k+1}, \ldots, z_r)$$

Gibbs sampling (cont.)

This implies that we have the following transition probabilities:

$$
\begin{aligned}
Q_{\boldsymbol{z},\boldsymbol{y}} &= \tfrac{1}{r} P(Z_k = z \mid Z_i = z_i, i \neq k) \\
&= \frac{g(\boldsymbol{y})/B}{r \cdot \sum_{z_k} g(\boldsymbol{y})/B} = \frac{g(\boldsymbol{y})}{r \cdot \sum_{z_k} g(\boldsymbol{y})}
\end{aligned}
$$

By the same type of argument, we also have:

$$
Q_{\boldsymbol{y},\boldsymbol{z}} = \frac{g(\boldsymbol{z})}{r \cdot \sum_{z_k} g(\boldsymbol{z})}.
$$

However, since $\sum_{z_k} g(\boldsymbol{y}) = \sum_{z_k} g(\boldsymbol{z})$, this implies that:

$$
g(\boldsymbol{z}) Q_{\boldsymbol{z},\boldsymbol{y}} = g(\boldsymbol{y}) Q_{\boldsymbol{y},\boldsymbol{z}}, \quad \text{for all } \boldsymbol{z}, \boldsymbol{y} \in \mathcal{S}.
$$

Hence, $\alpha_{\boldsymbol{z},\boldsymbol{y}} = 1$ for all $\boldsymbol{z}, \boldsymbol{y} \in \mathcal{S}$, and thus, $\boldsymbol{Q}$ is an optimal transition probability matrix.

Chapter 5

The Exponential Distribution and the Poisson Process

Chapter 5.2 The Exponential Distribution

A continuous random variable X is said to have an exponential distribution with parameter $\lambda > 0$, denoted as $X \sim exp(\lambda)$, if its probability density function is given by:

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

If $X \sim exp(\lambda)$, then the cdf of X is given by:

$$F(x) = P(X \leq x) = \int_0^x f(t)dt = \begin{cases} 1 - e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

Moreover, the ccdf of X is given by:

$$\bar{F}(x) = P(X > x) = 1 - F(x) = \begin{cases} e^{-\lambda x} & x \geq 0 \\ 1 & x < 0 \end{cases}$$

The Exponential Distribution (cont.)

The exponential distribution is a special case of the gamma distribution with parameters $\alpha > 0$ and $\lambda > 0$, denoted as $X \sim gamma(\alpha, \lambda)$ with probability density function:

$$f(x) = \begin{cases} \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

where $\Gamma(\alpha)$, defined for all $\alpha > 0$, is the gamma function given by:

$$\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} e^{-x} dx, \qquad \Gamma(n) = (n-1)!, \quad n = 1, 2, \ldots$$

By substituting $u = \lambda x$ and $du = \lambda dx$, we find that:

$$\int_0^\infty f(x) dx = \frac{1}{\Gamma(\alpha)} \int_0^\infty u^{\alpha-1} e^{-u} du = 1.$$

Thus, $f(x)$ is indeed a proper probability density.

The Exponential Distribution (cont.)

Assume that $X \sim exp(\lambda)$, and let $p > -1$. We then have:

$$E[X^p] = \int_0^\infty x^p f(x) dx = \int_0^\infty \lambda x^p e^{-\lambda x} dx$$
$$= \frac{\Gamma(p+1)}{\lambda^p} \int_0^\infty \frac{\lambda^{p+1}}{\Gamma(p+1)} x^{(p+1)-1} e^{-\lambda x} dx$$
$$= \frac{\Gamma(p+1)}{\lambda^p}.$$

In particular:

$$E[X] = \frac{\Gamma(2)}{\lambda^1} = \frac{(2-1)!}{\lambda} = \frac{1}{\lambda}, \qquad E[X^2] = \frac{\Gamma(3)}{\lambda^2} = \frac{(3-1)!}{\lambda^2} = \frac{2}{\lambda^2},$$

$$\mathrm{Var}[X] = E[X^2] - (E[X])^2 = \frac{2}{\lambda^2} - \frac{1}{\lambda^2} = \frac{1}{\lambda^2}.$$

The Exponential Distribution (cont.)

Assume that $X \sim gamma(\alpha, \lambda)$. Then the moment generating function of X is given by:

$$
\begin{aligned}
M_X(t) = E[e^{tX}] &= \int_0^\infty e^{tx} \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} dx \\
&= \int_0^\infty \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-(\lambda - t)x} dx \\
&= \frac{\lambda^\alpha}{(\lambda - t)^\alpha} \int_0^\infty \frac{(\lambda - t)^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-(\lambda - t)x} dx \\
&= \frac{\lambda^\alpha}{(\lambda - t)^\alpha}, \qquad \text{for all } t < \lambda.
\end{aligned}
$$

In particular, if $X \sim exp(\lambda)$, we have:

$$
M_X(t) = \frac{\lambda}{\lambda - t}, \qquad \text{for all } t < \lambda.
$$

The Exponential Distribution (cont.)

Proposition (5.1)

Assume that $X_1, \ldots, X_n$ are independent and $X_i \sim exp(\lambda)$, $i = 1, \ldots, n$, and let:

$$Y = X_1 + \cdots + X_n$$

Then $Y \sim gamma(n, \lambda)$.

PROOF: Using moment generating functions we get:

$$M_Y(t) = E[e^{tY}] = E[e^{tX_1 + \cdots + tX_n}] = M_{X_1}(t) \cdots M_{X_n}(t)$$

$$= \frac{\lambda}{\lambda - t} \cdots \frac{\lambda}{\lambda - t} = \frac{\lambda^n}{(\lambda - t)^n}$$

Hence, $Y \sim gamma(n, \lambda)$.

The Exponential Distribution (cont.)

Proposition 5.1 is a special case of the following more general result:

Proposition (5.1b)

Assume that $X_1, \ldots, X_n$ are independent and $X_i \sim gamma(\alpha_i, \lambda)$, $i = 1, \ldots, n$, and let:

$$Y = X_1 + \cdots + X_n$$

Then $Y \sim gamma(\alpha, \lambda)$, where $\alpha = \sum_{i=1}^{n} \alpha_i$.

PROOF: Using moment generating functions we get:

$$M_Y(t) = E[e^{tY}] = E[e^{tX_1 + \cdots + tX_n}] = M_{X_1}(t) \cdots M_{X_n}(t)$$

$$= \frac{\lambda^{\alpha_1}}{(\lambda - t)^{\alpha_1}} \cdots \frac{\lambda^{\alpha_n}}{(\lambda - t)^{\alpha_n}} = \frac{\lambda^{\alpha}}{(\lambda - t)^{\alpha}}$$

Hence, $Y \sim gamma(\alpha, \lambda)$.

The Exponential Distribution (cont.)

A random variable X is said to be memoryless if:

$$P(X > s + t | X > t) = P(X - t > s | X > t) = P(X > s), \quad \text{for all } s, t \geq 0.$$

Thus, X is memoryless if $(X - t)|(X > t)$ has the same distribution as X.

If $X \sim exp(\lambda)$, we have:

$$\begin{aligned} P(X > s + t | X > t) &= \frac{P(X > s + t \cap X > t)}{P(X > t)} \\ &= \frac{P(X > s + t)}{P(X > t)} \\ &= \frac{e^{-\lambda(s+t)}}{e^{-\lambda(t)}} = e^{-\lambda s} = P(X > s) \end{aligned}$$

Hence, we conclude that X is memoryless.

Example 5.2

The amount of time one spends in a bank, denoted X, is exponentially distributed with mean ten minutes. That is, $X \sim exp(\lambda) = exp(\frac{1}{10})$.

PROBLEM 1: What is the probability that a customer will spend more than fifteen minutes in the bank?

SOLUTION:

$$P(X > 15) = e^{-15\lambda} = e^{-15/10} \approx 0.223$$

PROBLEM 2: What is the probability that a customer will spend more than fifteen minutes in the bank given that she is still in the bank after ten minutes?

SOLUTION:

$$P(X > 15 \mid X > 10) = e^{-(15-10)\lambda} = e^{-5/10} \approx 0.607$$

The Exponential Distribution (cont.)

Assume that X_1, X_2 are independent and that $X_i \sim exp(\lambda_i)$, $i = 1, 2$. We want to calculate the probability of the event that $X_1 < X_2$.

$$
\begin{aligned}
P(X_1 < X_2) &= \int_0^\infty P(X_1 < X_2 | X_1 = x) \lambda_1 e^{-\lambda_1 x} dx \\
&= \int_0^\infty P(X_2 > x) \lambda_1 e^{-\lambda_1 x} dx \\
&= \int_0^\infty e^{-\lambda_2 x} \lambda_1 e^{-\lambda_1 x} dx \\
&= \frac{\lambda_1}{\lambda_1 + \lambda_2} \int_0^\infty (\lambda_1 + \lambda_2) e^{-(\lambda_1 + \lambda_2) x} dx \\
&= \frac{\lambda_1}{\lambda_1 + \lambda_2}.
\end{aligned}
$$

The Exponential Distribution (cont.)

Assume that $X_1, \ldots, X_n$ are independent and that $X_i \sim exp(\lambda_i)$, $i = 1, \ldots, n$.

$$
\begin{aligned}
P(\min_{1 \leq i \leq n} X_i > x) &= P(\bigcap_{i=1}^{n} X_i > x) \\
&= \prod_{i=1}^{n} P(X_i > x) \qquad \text{(by independence)} \\
&= \prod_{i=1}^{n} e^{-\lambda_i x} \\
&= e^{-(\sum_{i=1}^{n} \lambda_i)x}
\end{aligned}
$$

Thus, we have shown that $\min_{1 \leq i \leq n} X_i \sim exp(\sum_{i=1}^{n} \lambda_i)$.

The Exponential Distribution (cont.)

The following result combines the two previous results:

Assume that $X_1, \ldots, X_n$ are independent and that $X_i \sim exp(\lambda_i)$, $i = 1, \ldots, n$. We want to calculate the probability that X_i is the smallest of all the variables, i.e., that $X_i < X_j$ for all $j \neq i$.

$$P(\bigcap_{j \neq i}[X_i < X_j]) = P(X_i < \min_{j \neq i} X_j)$$

$$= \frac{\lambda_i}{\lambda_i + \sum_{j \neq i} \lambda_j}, \qquad \text{since } \min_{j \neq i} X_j \sim exp(\textstyle\sum_{j \neq i} \lambda_j)$$

$$= \frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j}$$

The Exponential Distribution (cont.)

Proposition (5.2)

Assume that $X_1, \ldots, X_n$ are independent and that $X_i \sim exp(\lambda_i)$, $i = 1, \ldots, n$. Then $\min_i X_i \sim exp(\sum_{i=1}^n \lambda_i)$. Moreover, $\min_i X_i$ and the rank order of $X_1, \ldots, X_n$ are independent.

PROOF: Since the exponential distribution is memoryless, we get that:

$$P(X_{i_1} < \cdots < X_{i_n} \mid \min_{1 \le i \le n} X_i > t)$$

$$= P(X_{i_1} < \cdots < X_{i_n} \mid \bigcap_{i=1}^n X_i > t)$$

$$= P(X_{i_1} - t < \cdots < X_{i_n} - t \mid \bigcap_{i=1}^n X_i > t)$$

$$= P(X_{i_1} < \cdots < X_{i_n})$$