DECISION MAKING

Decision making under conditions of uncertainty

Set of States of nature: $\mathbf{S_1}$, ..., $\mathbf{S_j}$, ..., $\mathbf{S_n}$
Set of decision alternatives: $\mathbf{d_1}$, ...,$\mathbf{d_i}$,...,$\mathbf{d_m}$

The outcome of the decision $\mathbf{C_{ij}}$ depends on the state of nature which is unknown when the decision is made. The outcome is usually measured as a monetary payoff.

The outcomes can be represented in a table:

	State of nature			
Decision	$\mathbf{S_1}$	$\mathbf{S_2}$	**...**	$\mathbf{S_n}$
$\mathbf{d_1}$	C_{11}	C_{12}	...	C_{1n}
$\mathbf{d_2}$	C_{21}	C_{22}	...	C_{2n}
. .				. .
$\mathbf{d_m}$	C_{m1}	C_{m2}	...	C_{mn}

Uncertainty versus risk

Sometimes it is useful to draw a distinction between *risk* and *uncertainty*. Risk is where probabilities can be assigned to the states of nature. If not, then it is called ‘uncertainty’.

A numerical example

	State of nature					
Decision	$\mathbf{S_1}$	$\mathbf{S_2}$	$\mathbf{S_3}$	**max**	**min**	**ave**
$\mathbf{d_1}$	780	900	1020	1020	780	900
$\mathbf{d_2}$	756	936	1056	1056	756	916
$\mathbf{d_3}$	732	912	1092	1092	732	912

Decision Criteria not using probabilities 1

Maximax criterion

Choose the alternative where the maximum payoff is highest:

$$\max\{i\} \max\{j\} \mathbf{C_{ij}}.$$

This is an *optimistic* criterion in that it only considers the best possible outcome. In the example it results in choice of d_3.

Maximin criterion

Choose the alternative where the minimum payoff is highest

$$\max\{i\} \min\{j\} \mathbf{C_{ij}}.$$

This is a *pessimistic* criterion in that it only considers the worst possible outcome. In the example it results in choice of d_1.

Hurwicz criterion

This is a middle ground between the maximax and maximin criteria. It involves choosing a coefficient of optimism α:

$$0 \leq \alpha \leq 1 .$$

Then the criterion selects the alternative with the highest weighted payoff:

$$\max\{i\} (\alpha \max\{j\} \mathbf{C_{ij}} + (1-\alpha) \min\{j\} \mathbf{C_{ij}}).$$

For the choice $\alpha=0.5$ we have:

d_1: $\alpha *1020 + (1-\alpha) *780 = 900$

d_2: $\alpha *1056 + (1-\alpha) *756 = 906$

d_3: $\alpha *1092 + (1-\alpha) *732 = 912$

so that the criterion results in the choice of d_3.

The interesting question here is how should α be chosen.

Decision Criteria not using probabilities 2

Laplace criterion

Choose the alternative where the average payoff is highest:

$$\max\{i\} \left((1/n) \text{ sum}\{j\} \mathbf{C_{ij}}\right).$$

This criterion treats all alternatives equally and can be regarded as making the assumption that all alternatives are equally likely. In the example, the criterion results in choice of d_2.

Savage (Minimax regret) criterion

Choose the alternative that minimises the opportunity loss or *regret* from making the wrong decision.

The opportunity loss $\mathbf{L_{ij}}$ is defined by

$$\mathbf{L_{ij}} = \max\{i\} \mathbf{C_{ij}} -- \mathbf{C_{ij}}.$$

For the numerical example, the Opportunity Loss or regret can be represented in a table:

Opportunity Loss Table

	Opportunity Loss			
Decision	**S_1**	**S_2**	**S_3**	**max**
d_1	0	36	72	72
d_2	24	0	36	36
d_3	48	24	0	48

For each state of nature, there is one (or more) decision that gives the best payoff. For this decision, the regret is zero since no better decision could have been taken given that state of nature. For all other decisions, the regret is the difference between the payoff for the decision and the best payoff for the state of nature.

The Savage criterion tries to minimise disappointment from making the wrong decision. In the example this results in choice of d_2.

Decision criteria using probabilities

We will now assume that we have the information to assign a probability $\mathbf{p_j}$ to each state of nature. For the numerical example these probabilities are represented in the table:

Probability Table

	State of nature		
Probability	**S_1**	**S_2**	**S_3**
p_j	0.2	0.5	0.3

We now consider two criteria that make use of this probability information about the relative likelihood of the different states of nature.

Expected Monetary Value (EMV) Criterion

Choose the alternative that maximises the expected monetary value of the payoff defined by:

$$\max\{i\} \text{ sum}\{j\}\ \mathbf{p_j}\ \mathbf{C_{ij}}.$$

For the numerical example we have:

d_1: $EMV_1 = 780 * 0.2 + 900 * 0.5 + 1020 * 0.3 = 912$

d_2: $EMV_2 = 756 * 0.2 + 936 * 0.5 + 1056 * 0.3 = 936$

d_3: $EMV_3 = 732 * 0.2 + 912 * 0.5 + 1092 * 0.3 = 930$

so that the alternative that maximises EMV_i is d_2.

Minimise Expected Regret Criterion

Choose the alternative that minimises *expected* loss

$$\min\{i\} \text{sum}\{j\} \mathbf{p_j L_{ij}}.$$

Expected Opportunity Loss Table

	Opportunity Loss			
Decision	**S_1**	**S_2**	**S_3**	**$E(L_i)$**
d_1	0	36	72	39.6
d_2	24	0	36	15.6
d_3	48	24	0	21.6

For the numerical example we have:

d_1: $E(L_1) = 0 * 0.2 + 36 * 0.5 + 72 * 0.3 = 39.6$

d_2: $E(L_2) = 24 * 0.2 + 0 * 0.5 + 36 * 0.3 = 15.6$

d_3: $E(L_3) = 48 * 0.2 + 24 * 0.5 + 0 * 0.3 = 21.6$

so that the alternative that minimises $E(L_i)$ is d_2.

It can be shown that the decision that minimises expected regret is *always* the same as the decision that maximises EMV value so that the two criteria can be regarded as equivalent.

The Expected Value of Perfect Information

Suppose that we had perfect information about the state of nature that was going to occur. Then we would always choose the decision to maximise payoff (this also corresponds to the decision with zero regret). The resulting payoffs are shown shaded in the diagram:

Payoffs with Perfect Information

	State of nature		
Decision	S_1	S_2	S_3
d_1	780	900	1020
d_2	756	936	1056
d_3	732	912	1092

The Expected Payoff from Perfect Information (EPPI) given the probability of each state of nature is defined by:

$$EPPI = \text{sum}\{j\}\ \mathbf{p}_j \max\{i\}\ \mathbf{C}_{ij}.$$

For the example:

$$EPPI = 780 * 0.2 + 936 * 0.5 + 1092 * 0.3 = 951.6.$$

The Expected Value of Perfect Information (EVPI) is the difference between the expected payoff with perfect information (EPPI) and the expected payoff without any information (EMV). This is the most that a decision maker would be willing to pay for the information.

$$EVPI = EPPI - EMV$$

$$= \text{sum}\{j\}\ \mathbf{p}_j \max\{i\}\ \mathbf{C}_{ij} - \max\{i\}\ \text{sum}\{j\}\ \mathbf{p}_j\ \mathbf{C}_{ij}.$$

For the numerical example EVPI = 951.6 - 936 = 15.6.

Attitudes to Risk

The Expected Monetary Value criterion implicitly assumes that the decision maker has a neutral attitude to risk. This can easily be seen from another simple numerical example:

Numerical example 2

	State of nature		
Decision	**S_1**	**S_2**	**EMV**
d_1	500	-250	50
d_2	35	60	50
d_3	50	50	50
p_j	0.4	0.6	

In this example, there are three decisions and two states of nature with probabilities 0.4 and 0.6 respectively. It can be seen that the expected monetary value of all three decisions is equal to 50 and so, using the EMV criterion, the decision maker should be indifferent between them.

However, decision $\mathbf{d_1}$ involves a 60% probability of making a large monetary loss whereas decision $\mathbf{d_3}$ carries a certain payoff of 50 and so involves no risk whatsoever.

Differing attitudes to risk can be illustrated in the following diagram:

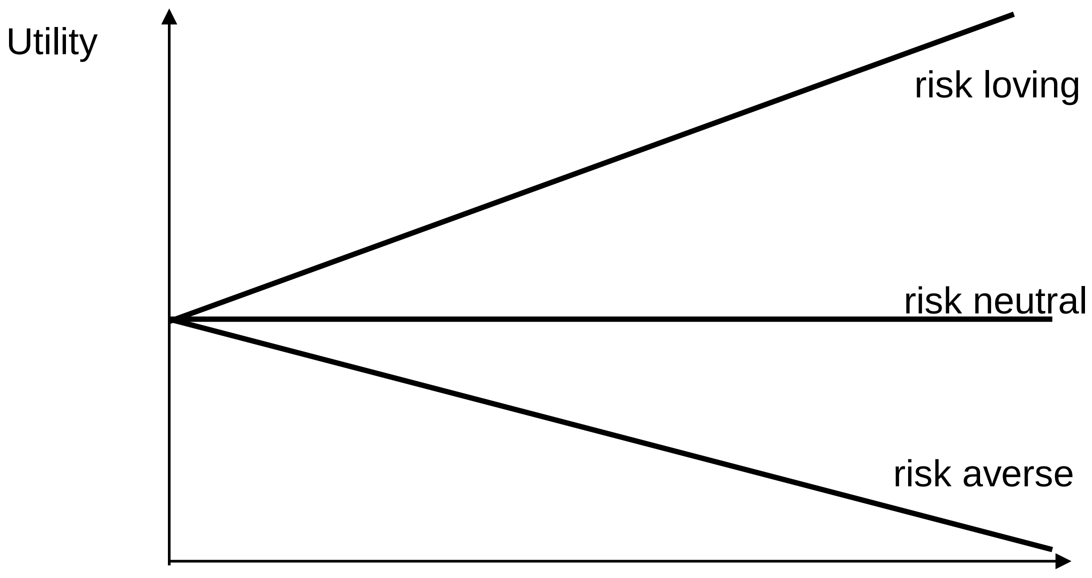

risk

The risk neutral agent is indifferent to risk. We might expect risk neutral behaviour from large institutions such as Insurance companies that can afford to average out wins and losses.

The risk averse agent wants to avoid risk. This might be expected of individuals and small companies that cannot

Risk loving behaviour is the attitude exhibited by gamblers. In practice we might expect that agents may be risk loving for 'small' gambles although risk averse for larger gambles.

If we could measure the risk attitude of the decision maker, then we could replace monetary payoffs by utility and use a decision criterion that maximises expected utility.

Assigning utilities to monetary outcomes

Von-Neumann and Morgenstern showed how revealed preference could be used to assign utility values to different gambles. Consider the gamble represented by decision $\mathbf{d_1}$ in the numerical example.

Decision	S_1	S_2	EMV
$\mathbf{d_1}$	500	-250	50
$\mathbf{p_i}$	0.4	0.6	

This gamble is a 40% chance of winning 500 and a 60% chance of losing 250. The expected (average) outcome is a win of 50.

We now ask the question, what is the value of the probability p at which an agent is indifferent between a certain outcome or the gamble

$$p * 500 + (1 - p) * (-250) \text{ ?}$$

Without loss of generality we can arbitrarily assign utility values of 1 and zero to the two payoffs 500 and -250 so that

$$U(500)=1 \quad U(-250)=0.$$

Then the expected utility of the gamble

$$p * U(500) + (1 - p) * U(-250) = p$$

is given by the probability p.

A risk neutral agent will be indifferent when the certain outcome is equal to the EMV of the gamble

$$p * 500 + (1 - p) *(-250)$$

.
e.g. for a certain outcome of 50 a risk neutral agent will be indifferent for p=0.4.

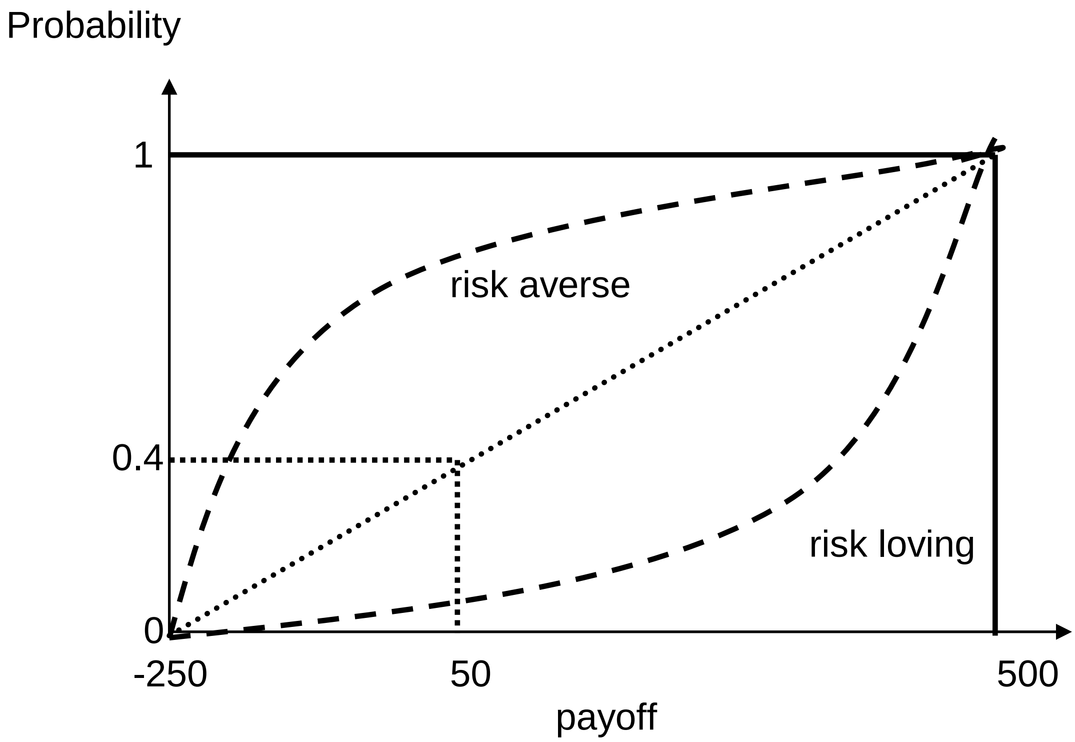

In the diagram, the indifference curve for a risk neutral agent lies on the straight line from -250, 0 to 500, 1.

A risk averse agent will have an indifference curve lying *above* this line since they require a probability of winning higher than the expected outcome to compensate them for bearing risk.

A risk loving agent will have an indifference curve lying below the straight line since they are prepared to take a gamble in which they will lose on average.

QUEUING MODELS AND MARKOV PROCESSES

Queues form when customer demand for a service cannot be met immediately. They occur because of fluctuations in demand levels so that models of queuing are stochastic.

Some definitions

- The number of servers is s

- The mean arrival rate (number of customers per unit of time) is λ

  - this is assumed to follow a Poisson distribution

  $$P_n = \frac{\lambda n e^{-\lambda}}{n!}$$

  where P_n is the probability of n arrivals in the time period.

- The mean service rate (number of customers served per unit time per server) is μ

  - this is assumed to follow an exponential distribution

  $$P(t) = 1 - e^{-\mu t}$$

  where P(t) is the probability of being served by time period t.

- We assume independence of the two processes λ and μ and require the condition that s * $\mu > \lambda$, otherwise the queue will grow indefinitely.

- The average service time is given by 1 / μ.

Costs of Queuing

Waiting Cost (WC)

This depends on the average time spent waiting in line

$$\mathbf{WC = C_w\, L} \qquad \text{where} \quad \partial \mathbf{L} / \partial \mathbf{s} < 0$$

and $\mathbf{C_w}$ is the waiting cost per customer per unit of time.

Service Cost (SC)

This is directly proportional to the number of servers s.

$$\mathbf{SC = C_s\, s}$$

where $\mathbf{C_s}$ is the cost per customer per unit of time. Increasing the number of servers s decreases **WC** but increases **SC**.

Total Cost (TQC)

$$\mathbf{TQC = WC + SC}$$

Total costs are nonlinear in **s** with a minimum at **s***, the optimal level of service.

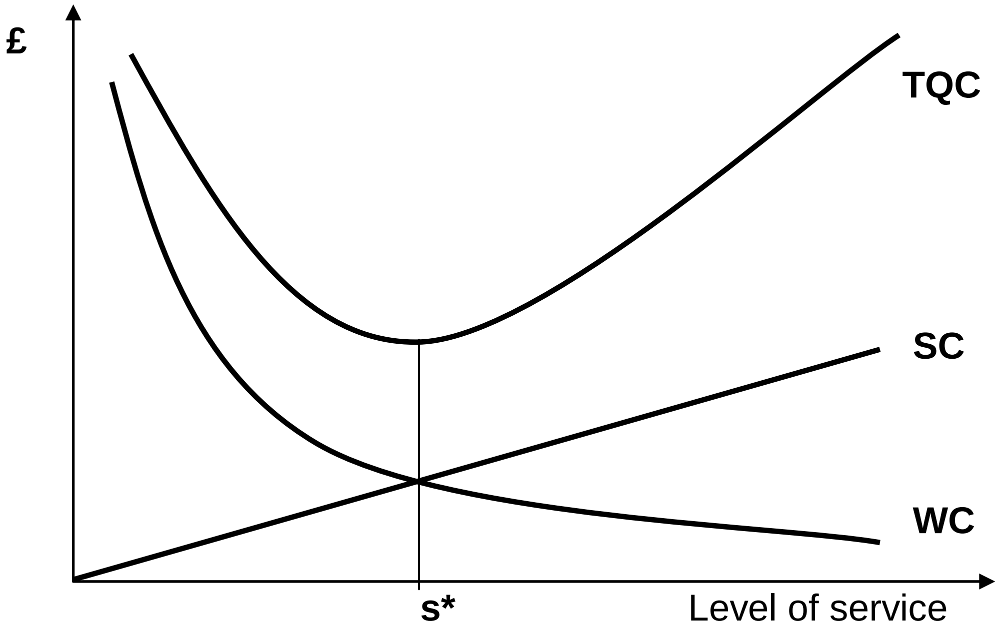

A Simple Model with One Server

Consider the simplest model with s=1. Then the following results hold:

- the probability of the system being busy: $\rho = \lambda / \mu$.
- the probability of n customers: $P_n = (1-\rho) \rho^n$.
- the average number of customers: $L = \lambda / (\mu - \lambda)$.
- the average time spent in the system: $W = L / \lambda$
- the average time spent queuing: $W_q = \lambda / [\mu (\mu - \lambda)] = W - 1/\mu$.

A general result in a steady state queuing process is that $L = \lambda W$. This relationship is known as Little's law.

A Model with s Servers

Consider a model with s servers. Then the following results hold:

- the probability of the system being busy: $\rho = \lambda /(s \mu)$.
- the probability of 0 customers P_0 is given by the formula

$$P_0 = \frac{1}{\sum_{n=0}^{s-1} \frac{(\lambda/\mu)^n}{n!} + \left(\frac{(\lambda/\mu)^s}{(s-1)!}\right)\left(\frac{\mu}{s\mu - \lambda}\right)}$$

- the probability of n customers: P_n is given by

$$P_n = \frac{(\lambda/\mu)^n}{s! s^{n-s}} P_0 \quad \text{for} \quad n > s \quad \text{and}$$
$$P_n = \frac{(\lambda/\mu)^n}{n!} P_0 \quad \text{for} \quad 0 < n \leq s.$$

- the average number of customers: L is given by

$$L = \frac{(\lambda/\mu)^s \lambda\mu}{(s-1)!(s\mu - \lambda)^2} P_0 + \frac{\lambda}{\mu}$$

- the average time spent in the system: $W = L / \lambda$

- the average time spent queuing: $W_q = W - 1/\mu$.

A general result in a steady state queuing process is that $L = \lambda W$. This relationship is known as Little's law.

MARKOV PROCESSES

Markov processes are used to describe a system moving over time between different states.

Suppose that there are n states: $S_1, ..., S_n$
and the probability of being in a state at time t: $p_1(t), ..., p_n(t)$

The probabilities are assumed to change over time following a simple stochastic process.

The probabilities $p_1(t), ..., p_n(t)$ can be represented in a row vector

$$\mathbf{p}(t+1) = [\ p_1(t)\ .\ .\ .\ p_n(t)\ ].$$

Transitional Probability

The transitional probability $P_{ij}(t)$ is the probability of moving from state i to state j at time t. This is the conditional probability

$$P_{ij}(t) = p_j(t+1) \mid p_i(t)$$

A crucial assumption of the Markov model is that transitional probabilities are independent of time so that

$$P_{ij} = p_j(t+1) \mid p_i(t)$$

The transitional probabilities can be represented in a transition matrix of dimension n x n:

$$\mathbf{P} = \begin{bmatrix} P_{11} & \dots & P_{1n} \\ & P_{ij} & \\ P_{n1} & \dots & P_{nn} \end{bmatrix}$$

The elements in each row of the matrix **P** must sum to unity since, whatever state you are in at time t, you must end up in one of the n states in period t+1.

The First Order Markov Process

The simplest Markov process is the first order process defined by the matrix equation

$$\mathbf{p}(t+1) = \mathbf{p}(t)\ \mathbf{P}$$

or, equivalently,

$$p_j(t+1) = \Sigma_i\ p_i(t)\ P_{ij}.$$

The first order Markov process is called a *zero memory process* because the state next period depends only on the current state and not on any past states.

Higher Order Markov Processes

More generally, a k^{th} order Markov process can be defined by

$$\mathbf{p}(t+1) = \mathbf{p}(t)\ \mathbf{P}_1 + \mathbf{p}(t-1)\ \mathbf{P}_2 + \ldots + \mathbf{p}(t-k+1)\ \mathbf{P}_k$$

where $\mathbf{P}_1$, $\mathbf{P}_2$, . . . , $\mathbf{P}_k$ are n x n transition matrices. The k^{th} order Markov process depends on the state as far as k-1 periods back.

Predicting Future States

Starting from an initial state p(0) at time t=0 and assuming a first order Markov process, the state at time t=1 is determined by the matrix equation

$$\mathbf{p}(1) = \mathbf{p}(0)\ \mathbf{P}\ .$$

For period t=2 we have

$$\mathbf{p}(2) = \mathbf{p}(1)\ \mathbf{P} = \mathbf{p}(0)\ \mathbf{P}^2$$

where $\mathbf{P}^2 = \mathbf{P} * \mathbf{P}$ is the matrix product of $\mathbf{P}$ with itself.

In general, by substitution, it can be seen that

$$\mathbf{p}(t) = \mathbf{p}(0)\ \mathbf{P}^{t}.$$

The Steady State of a Markov Process

Most Markov processes eventually converge to a steady state.
This is because in general

$$\text{Limit}_{t\to\infty} \mathbf{P}^t \text{ converges to a fixed matrix.}$$

When this is true, after a certain time **p**(t) will not change anymore so that **p**(t) = **p**(t-1) = **p*** and **p*** will satisfy the equation

$$\mathbf{p^*} = \mathbf{p^*}\,\mathbf{P}\ .$$

The steady state, if it exists, can be computed analytically by solving this system of n equations subject to the adding-up restriction that the elements of **p*** must sum to one, or formally,

$$\Sigma_i\ p^*_i = 1.$$

This is a system of n+1 equations in n unknowns but it can be solved by dropping one of the equations.

Special case: Absorbing states

An absorbing state is a state from which, once the state is entered, there is no possibility of exit. An absorbing state is analogous to a black hole in astrophysics. In the transition matrix, an absorbing state will have a row with a one on the diagonal and elsewhere zeroes.

When a transition matrix has absorbing states, then in the limit, everything will end up in these states.

An Example

Suppose the transitional matrix is given by

$$\mathbf{P} = \begin{bmatrix} 0.6 & .0.2 & 0.2 \\ 0.2 & 0.7 & 0.1 \\ 0.2 & 0.3 & 0.5 \end{bmatrix}$$

with initial state

$$\mathbf{p}(0) = [\ 0.324\ 0.441\ 0.235] \ .$$

Then

$$\mathbf{p}(1) = \mathbf{p}(0)\ \mathbf{P}$$

$$= [\ 0.330\ 0.444\ 0.226] \ .$$

and

$$\mathbf{p}(2) = \mathbf{p}(1)\ \mathbf{P} = \mathbf{p}(0)\ \mathbf{P}^2$$

$$= [\ 0.332\ 0.445\ 0.224] \ .$$

Continuing to project forwards

$$\mathbf{p}(11) = \mathbf{p}(0)\ \mathbf{P}^{11}$$

$$= [\ 0.333\ 0.444\ 0.222]$$

and

$$\mathbf{p}(12) = \mathbf{p}(0)\ \mathbf{P}^{12}$$

$$= [\ 0.333\ 0.444\ 0.222]$$

at which point the process has converged to a steady state (to 3 decimal places) so that

$$\mathbf{p}(t > 12) = \mathbf{p^*} = [\ 0.333\ 0.444\ 0.222] \ .$$