EXAMINATION

September 2007

Subject CT4 — Models
Core Technical

EXAMINERS' REPORT

Introduction

The attached subject report has been written by the Principal Examiner with the aim of helping candidates. The questions and comments are based around Core Reading as the interpretation of the syllabus to which the examiners are working. They have however given credit for any alternative approach or interpretation which they consider to be reasonable.

M A Stocker
Chairman of the Board of Examiners

December 2007

Comments

Comments on solutions presented to individual questions for this September 2007 paper are given below and further comments may be written within in the solutions that follow.

Q1 *This straightforward bookwork question was not especially well answered.*

Q2 *This was the most poorly answered question on the examination paper. Very few candidates recognised that the baker's problem could be modelled using the compound Poisson process described in Unit 2 section 3.4 of the Core Reading.*

Q3 *This was well answered, with many candidates scoring full marks.*

Q4 *Although most candidates performed the chi-squared test correctly, few realised that when using this to test a graduation some degrees of freedom are lost, a fact which is clearly stated in the Core Reading in Unit 12, section 7.3. In part (ii) comments tended not to be related to the data in the question; rather they focused rather mechanically on the shortcomings of the chi-squared test.*

Q5 *This straightforward bookwork question was well answered by many candidates.*

Q6 *Most candidates obtained the correct numerical answers in part (i) of this question, but answers to part (ii) were rather sketchy and vague.*

Q7 *This was more demanding than some previous questions on the Kaplan-Meier or Nelson-Aalen estimators, and the standard of the answers was lower than expected.*

Q8 *This exposed-to-risk question was easier than many questions on the same topic in previous papers. Most candidates scored well on parts (i) and (ii), although few explained that the method relied on the assumption of a uniform distribution of deaths. Answers to part (iii) were less impressive and tended to lack detail. Some candidates couched their answers to this part in aggregate terms, despite the question clearly referring to individual-level data.*

Q9 *Most candidates scored well on parts (i) and (ii). Common errors included the use of a three-state model (Deuce, Advantage and Game) which is inappropriate as the transition out of the state "Advantage" is ill-defined. Few candidates made attempts at parts (iii) and (iv) and several of these wrongly thought that part (iv) could be solved by finding the stationary distribution of the chain.*

Q10 *Parts (i) and (iv) of this question tested knowledge of Unit 7, sections 2, 3 and 5 of the Core Reading, which has not been tested in previous CT4 examination papers. Perhaps because of this, many candidates gave very sketchy and vague answers. In part (ii), while most candidates spotted that Type I censoring was present, only a small minority also registered the existence of random censoring. In part (iii) few candidates correctly interpreted the sex x smoking interaction. Part (v) was well answered by most candidates.*

Q11 *Many candidates only attempted parts (i) and (ii) of this question. The remainder was very poorly answered, with few candidates making serious attempts at part (vi), despite this being bookwork based on Core Reading, Unit 4, section 5.4.*

1 Factors to be considered include:

- the objectives of the modelling exercise,
- the validity of the model for the purpose to which it is to be put,
- the validity of the data to be used,
- the possible errors associated with the model or parameters used not being a perfect fit,
- representation of the real world situation being modelled,
- the impact of correlations between the random variables that drive the model,
- the extent of correlations between the various results produced from the model,
- the current relevance of models written and used in the past,
- the credibility of the data input,
- the credibility of the results output,
- the dangers of spurious accuracy,
- the ease with which the model and its results can be communicated.

Not all these factors needed to be mentioned for full marks to be awarded.

2 (a) Assume that, during each day, customers arrive at the shop according to a Poisson process.

Assume that the numbers of buns bought by each customer, the B_j, are independent and identically distributed random variables.

Then if X_t is the total number of buns sold between the beginning of the day and time t, (where t is measured in hours since the shop opens), X_t is a compound Poisson process defined by

$$X_t = \sum_{j=1}^{N_t} B_j ,$$

where the number of customers arriving between the shop opening and time t is N_t .

(b) The probability that the baker will run out of buns is

$$\Pr[K + ct - \sum_{j=1}^{N_t} B_j < 0]$$

for some t.

3 The transition matrix for the chain is:

$$\begin{pmatrix} 1-\alpha & \alpha & \\ 1-\alpha & & \alpha \\ & 1-\alpha & \alpha \end{pmatrix}.$$

To determine the long-run probability, we need to solve the equation $\pi P = \pi$, which reads:

$$\begin{aligned}
&\text{(I)} && \pi_1 = (1-\alpha)\pi_1 + (1-\alpha)\pi_2 \\
&\text{(II)} && \pi_2 = \alpha\pi_1 + (1-\alpha)\pi_3 \\
&\text{(III)} && \pi_3 = \alpha\pi_2 + \alpha\pi_3 .
\end{aligned}$$

The probabilities must also satisfy:

(IV) $\pi_1 + \pi_2 + \pi_3 = 1$.

(III) gives $\pi_2 = \left(\frac{1-\alpha}{\alpha}\right)\pi_3$.

Substituting in (I) gives $\pi_1 = \left(\frac{1-\alpha}{\alpha}\right)^2 \pi_3$,

and so (IV) leads to $\left(\left(\frac{1-\alpha}{\alpha}\right)^2 + \left(\frac{1-\alpha}{\alpha}\right) + 1\right)\pi_3 = 1$.

We know that $\pi_3 = 0.75$, which leads to:

$$\left(\frac{(1-\alpha)^2 + \alpha(1-\alpha) + \alpha^2}{\alpha^2}\right) \times 0.75 = 1,$$

$$\Rightarrow 0.75\left(\left(1-2\alpha+\alpha^2\right) + \left(\alpha+\alpha^2\right) + \alpha^2\right) = \alpha^2,$$

$$\Rightarrow 0.25\alpha^2 + 0.75\alpha - 0.75 = 0 \ .$$

Using the quadratic equation formula, this leads to

$$\alpha = \frac{-0.75 \pm \sqrt{0.75^2 + 4 \times 0.25 \times 0.75}}{2 \times 0.25}.$$

As $\alpha > 0$, we must have $\alpha = 0.7913$.

4 (i) The null hypothesis is that graduated rates are the same as the true underlying rates in the population.

To test overall goodness-of-fit we use the chi-squared test.

$\sum_x z_x^2 \sim \chi^2_m$, where m is the number of degrees of freedom.

In this case, we have 10 ages.

The graduation was carried out by reference to a standard table, so we lose a number of degrees of freedom because of the choice of standard table.

So, $m < 10$, and let us say $m = 8$.

The observed value of the test statistic is $\sum_x z_x^2 = 15.8623$

The critical value of the chi-squared distribution with 8 degrees of freedom at the 5 per cent level is 15.51.

Since 15.8623 > 15.51,

we reject the null hypothesis and conclude that the graduated rates do not adhere to the data.

[Credit was given for using other values of m, say m = 7 or m = 9, provided candidates recognized that some degrees of freedom should be lost for the choice of standard table. Note that if m = 9, the null hypothesis will not be rejected.]

(ii) From the data we can see that the actual deaths are lower than those expected at all ages.

The graduated rates are too high; the graduation should be revisited.

At these ages the force of mortality increases with age,
so a suitable adjustment may be to reduce the age shift relative to the standard table from 2 years.

The standardised deviations also appear to show a systematic increase with age, showing that departure of the graduated rates from the actual rates increases with age.

There appear to be no outliers (all the z_xs have absolute values below 1.96).

5 (i) We assume that mortality rates progress smoothly with age.

Therefore a crude estimate at age x carries information about the rates at adjacent ages, and graduation allows us to use this fact to "improve" the estimate at age x by smoothing.

This reduces the sampling errors at each age.

It is desirable that financial quantities progress smoothly with age, as irregularities are hard to justify to clients.

(ii) Any two of the following three methods are acceptable:

By parametric formula:

Should be used for large experiences, especially if the aim is to produce a standard table;

Depends on a suitable formula being found which fits the data well.

Provided the number of parameters is small, the resulting curve should be smooth.

With reference to a standard table

Should be used if a standard table for a class of lives similar to the experience is available, and the experience we are interested in does not provide much data.

The standard table will be smooth,

and provided the function linking the graduated rates to the rates in the standard table is simple, this smoothness will be "transferred to the graduated rates".

Graphical

if a quick check is needed, or data are very scanty.

The graduation should be tested for smoothness using the third differences of the graduated rates, which should be small in magnitude and progress regularly with age.

If the smoothness is unsatisfactory, the curve can be adjusted ("hand-polishing") and the smoothness tested again.

6 (i) (a) Assuming a uniform distribution of deaths between ages 58 and 62 implies that half of those who die between those ages die between ages 58 and 60.

Therefore

$$l_{60} = l_{58} - 0.5(l_{58} - l_{62})$$

$$= 88{,}792 - 0.5(88{,}792 - 84{,}173)$$

$$= 86{,}482.5.$$

(b) ALTERNATIVE 1

Let the constant force of mortality be μ.

Then we have ${}_4p_{58} = \exp\left(-\int_0^4 \mu dx\right) = e^{-4\mu}$.

But ${}_4p_{58} = \frac{l_{62}}{l_{58}} = \frac{84{,}173}{88{,}792} = 0.94798$.

Therefore $e^{-4\mu} = 0.94798$,

so that $-4\mu = \log_e(0.94798) = -0.05342$,

whence $\mu = 0.01336$.

Therefore with a constant force of mortality,

$l_{60} = l_{58}\exp[-2(0.01336)] = 88{,}792(0.97363)$

so $l_{60} = 86{,}452$.

ALTERNATIVE 2

Let the constant force of mortality be μ.

Then we have ${}_4p_{58} = \exp\left(-\int_0^4 \mu dx\right) = e^{-4\mu}$.

But ${}_4p_{58} = \frac{l_{62}}{l_{58}}$.

Now $l_{60} = l_{58} \cdot {}_2p_{58}$.

and, since ${}_2p_{58} = e^{-2\mu} = \sqrt{e^{-4\mu}} = \sqrt{\dfrac{l_{62}}{l_{58}}}$,

$$l_{60} = l_{58}\sqrt{\frac{l_{62}}{l_{58}}} = \sqrt{l_{58}l_{62}} = \sqrt{(88,792)(84,173)}$$

so l_{60} = 86,452

(ii) The actual value of l_{60} from the tables is 86,714.

This shows that neither assumption is very accurate, but that the uniform distribution of deaths (UDD) is closer than the constant force of mortality.

The UDD assumption is better than the constant force of mortality assumption because UDD implies an increasing force of mortality over this age range, which is biologically more plausible than the assumption of a constant force.

The fact that the actual value of l_{60} is considerably greater than that implied by the UDD assumption suggests that the true rate of increase of the force of mortality over this age range in English Life Table 15 (males) is even greater than that implied by UDD.

7 (i) (a) If, for player i, T_i is the number of games played before he is dismissed, and C_i is the total number of games played before 1 December, and d_i = 1 if the player had been dismissed before 1 December and 0 otherwise.

then

EITHER

from the data given we can create the two variables

min(T_i,C_i)
and d_i,
e.g. for player 1, min(T_i,C_i) = 12 and d_i = 0

OR

The required data for the Kaplan-Meier estimator are therefore

Player	*min(T_i,C_i)*	d_i
1	12	0
2	12	0
3	5	1
4	12	0
5	7	1
6	12	0
7	10	0
8	0	1
9	5	1
10	8	0
11	2	1
12	5	0
13	5	0
14	0	1
15	4	0

(b) Censoring in these data arises because not all players have been dismissed before 1 December. Those players who have yet to be dismissed on that data are right-censored.

This censoring is random [NOT Type I], because the metric of "duration" is the number of games played since the start of the season, and this may vary from player to player.

(ii) ALTERNATIVE 1 (where censorings are assumed to occur immediately before events)

t_j	N_j	D_j	C_j	$\frac{D_j}{N_j}$	$1-\frac{D_j}{N_j}$
0	15	2	0	2/15	13/15
2	13	1	3	1/13	12/13
5	9	2	0	2/9	7/9
7	7	1	6	1/7	6/7

Then the Kaplan-Meier estimate of the survival function is

t	$\hat{S}(t)$
$0 \le t < 2$	0.8667
$2 \le t < 5$	0.8000
$5 \le t < 7$	0.6222
$7 \le t < 12$	0.5333

Therefore the value of the chosen statistic, $\hat{S}(10)$ is 0.5333.

ALTERNATIVE 2 (where censorings are assumed to occur immediately after events)

t_j	N_j	D_j	C_j	$\dfrac{D_j}{N_j}$	$1-\dfrac{D_j}{N_j}$
0	15	2	0	2/15	13/15
2	13	1	1	1/13	12/13
5	11	2	2	2/11	9/11
7	7	1	6	1/7	6/7

Then the Kaplan-Meier estimate of the survival function is

t	$\hat{S}(t)$
$0 \le t < 2$	0.8667
$2 \le t < 5$	0.8000
$5 \le t < 7$	0.6545
$7 \le t < 12$	0.5610

Therefore the value of the chosen statistic, $\hat{S}(10)$ is 0.5610.

8 (i) The central exposed to risk at age x, E_x^c, is the observed waiting time in a multiple-state or a Poisson model. It is the sum of the times spent under observation by each life at age x.

In aggregate data, the central exposed to risk is an estimate of the number of lives exposed to risk at the mid-point of the rate interval.

The initial exposed to risk requires adjustments for those lives who die, whom we continue observing until the end of the rate interval.

It may be approximated as $E_x^c + 0.5d_x$, where d_x is the number of deaths to persons aged x.

(ii) The age definition used for both deaths and exposed to risk is the same, so no adjustment is necessary.

Using the census formula, and assuming that the population aged 22 and 23 years changes linearly over the year, we have, for the central exposed to risk:

$$E_x^c = \int_0^1 P_{x,t}\,dt,$$

so that

$$E_x^c = \frac{1}{2}(P_{x,0} + P_{x,1}).$$

The initial exposed to risk, E_x , is then obtained using the approximation $E_x^c + 0.5d_x$.

This assumes that deaths are uniformly distributed across each year of age.

Therefore, at age 22 we have

$$E_{22} = \frac{1}{2}(150+160) + \frac{20}{2} = 165,$$

and

$$E_{23} = \frac{1}{2}(160+155) + \frac{25}{2} = 170.$$

Hence $q_{22} = \dfrac{20}{165} = 0.1212$ and $q_{23} = \dfrac{25}{170} = 0.1471$.

[The complete derivation was not required for full marks.]

(iii) ALTERNATIVE 1

The central exposed to risk is calculated as $\sum_i (b_i - a_i)$, for all lives i for whom $b_i - a_i > 0$,

where a_i and b_i are measured in years since the person's 22nd birthday, and

where b_i is the earliest of

the date of person i's death
the date of person i's 23rd birthday
the end of the calendar year 2005
the date of person i's exit from observation for reasons other than death

and a_i is the latest of

the date of person i's 22nd birthday
the start of the calendar year 2005
the date of person i's entry into observation.

The initial exposed to risk is then calculated by adding on to the central exposed to risk a quantity equal to $1-b_i$ for all lives who died aged 22 last birthday during the calendar year 2005.

ALTERNATIVE 2

The initial exposed to risk is calculated as $\sum_i (b_i - a_i)$,

where a_i and b_i are measured in years since the person's 22nd birthday, and

where b_i is the earliest of

the date of person i's 23rd birthday
the date of person i's exit from observation for reasons other than death

and a_i is the latest of

the date of person i's 22nd birthday
the start of the calendar year 2005
the date of person i's entry into observation.

for all lives i for whom $b_i - a_i > 0$.

9 (i) State space:

{Deuce, Advantage A(ndrew), Advantage B(en),
Game A(ndrew), Game B(en)}.

Transition matrix:

	Deuce	*Adv A*	*Adv B*	*Game A*	*Game B*
Deuce	0	0.6	0.4	0	0
Adv A	0.4	0	0	0.6	0
Adv B	0.6	0	0	0	0.4
Game A	0	0	0	1	0
Game B	0	0	0	0	1

The chain is Markov because the probability of moving to the next state does not depend on history prior to entering that state (because the probability of each player winning a point is constant)

(ii) The chain is reducible because it has two absorbing states Game A and Game B.

States Game A and Game B are absorbing so have no period. The other three states each have a period of 2 so the chain is not aperiodic.

(iii) The game either ends after 2 points or it returns to Deuce.

The probability of it returning to Deuce after two points is:

Prob A wins 1st point × Prob B wins 2nd point
+ Prob B wins 1st point × Prob A wins 2nd point

$= 0.6 \times 0.4 + 0.4 \times 0.6 = 0.48$.

[This can also be obtained by calculating the square of the transition matrix.]

Need to find number of such cycles N such that:

$0.48^N < 1 - 0.9$,

so that

$$N > \frac{\ln 0.1}{\ln(0.48)} > 3.14.$$

But the game can only finish every two points so we require 4 cycles, that is 8 points.

(iv) (a) Define A_X to be the probability that A ultimately wins the game when the current state is X.

We require A_{Deuce}.

By definition $A_{\text{Game A}} = 1$ and $A_{\text{Game B}} = 0$.

Conditioning on the first move out of state Adv A:

$$A_{\text{Adv A}} = 0.6 \times A_{\text{Game A}} + 0.4 \times A_{\text{Deuce}} = 0.6 + 0.4 \times A_{\text{Deuce}}.$$

Similarly:

$$A_{\text{Adv B}} = 0.6 \times A_{\text{Deuce}},$$

and

$$A_{\text{Deuce}} = 0.6 \times A_{\text{Adv A}} + 0.4 \times A_{\text{Adv B}} = 0.6 \times A_{\text{Adv A}} + 0.24 \times A_{\text{Deuce}}.$$

So,

$$A_{\text{Deuce}} = \frac{0.6}{0.76} A_{\text{Adv A}},$$

$$A_{\text{Adv A}} = 0.6 + 0.4 \times \frac{0.6}{0.76} A_{\text{Adv A}},$$

and

$$A_{\text{Adv A}} = 0.8769,$$

and

$$A_{\text{Deuce}} = 0.6923.$$

ALTERNATIVELY

Probability A wins after 2 points = 0.6*0.6 =0.36

Probability that A wins from Deuce
$= \sum_{i=1}^{\infty}$ Probability A wins after i points have been played

= Probability A wins after 2 points
+ Probability A wins after 4 points +…..
(as period 2)

= 0.36 + 0.48 * 0.36 + 0.48^2 * 0.36 +…….

= 0.36/(1-0.48) as a geometric progression

= 0.6923

(b) This is higher than 0.6 because Ben has to win at least two points in a row to win the game.

10 (i) Fully parametric models are good for comparing homogenous groups, as confidence intervals for the fitted parameters give a test of difference between the groups which should be better than non-parametric procedures, or semi-parametric procedures such as the Cox model.

But parametric methods need foreknowledge of the form of the hazard function, which might be the object of the study.
The Cox model is semi-parametric so such knowledge is not required.

The Cox model is a standard feature of many statistical packages for estimating survival model, but many parametric distributions are not, and numerical methods may be required, entailing additional programming.

(ii) Type I censoring, since the investigation ends after a period which is fixed in advance.

Random censoring, since death from a cause other than a heart attack is a random variable and may occur at any time.

(iii) The likelihood ratio statistic is a common criterion.

Suppose we fit a model with p covariates and another model with $p+q$ covariates which include all the p covariates of the first model.

Then if the maximised log-likelihoods of the two models are L_p and L_{p+q}, then the statistic

$$-2(L_p - L_{p+q})$$

has a chi-squared distribution with q degrees of freedom, under the hypothesis that the extra q covariates have no effect in the presence of the original p covariates.

This statistic can be used either will full likelihoods or with partial likelihoods in the Cox model

This statistic can be used to test the statistical significance of any set of q covariates in the presence of any other disjoint set of p covariates.

(iv) Holding other factors constant,

females have a lower risk of heart attack than males,

and smokers have a higher risk than non-smokers,

but the effect of smoking varies for men and women.

The relative risks, compared with the baseline category of male non-smokers are as follows.

female non-smokers	$\exp(-0.4)$	$= 0.67$
male smokers	$\exp(0.5)$	$= 1.65$
female smokers	$\exp(-0.4+0.5-0.25)$	$= 0.86$

(or any other numerical example to illustrate the previous points)

(v) Let the required age for the woman smoker be $50+x$.

The hazard for this woman is

$h(t,x) = h_0(t) \exp(0.01x - 0.4 + 0.5 - 0.25)$,

The hazard for a male non-smoker aged 50 at the initial interview is simply $h_0(x)$, since this is the baseline category.

Thus we have

$h_0(t) \exp(0.01x - 0.4 + 0.5 - 0.25) = h_0(t)$

so that

$\exp(0.01x - 0.4 + 0.5 - 0.25) = 1$

or

$\exp(0.01x - 0.15) = 1$

so that

$0.01x = 0.15$

Therefore $x = 15$, and the woman's age at interview must be 65 years.

11 (i) (a) The parameters are:

- the rate of leaving state i, λ_i, for each i,
- the jump-chain transition probabilities, r_{ij}, for $j \neq i$, where r_{ij} is the conditional probability that the next transition is to state j given the current state is i.

[Alternatively the parameters may be expressed as σ_{ij}, where $\sigma_{ii} = -\lambda_i$ and (for $j \neq i$), $\sigma_{ij} = \lambda_i r_{ij}$.]

(b) The assumptions are as follows.

- The holding time in each state is exponentially distributed. The parameter of this distribution varies only by state i. The distribution is independent of anything that happened prior to the current arrival in state i.

- The destination of the jump on leaving state i is independent of holding time, and of anything that happened prior to the current arrival in state i.

ALTERNATIVELY

The holding time in each state is exponentially distributed and the destination of the jump on leaving state i is independent of holding time

Both holding time distribution and destination of jump on leaving state i are independent of anything that happened prior to arrival in state i

(ii) (a) The estimator *[it is the MLE but this need not be stated]* of λ_i, $\hat{\lambda}$, is the inverse of the average duration of each visit to state i.

so $\hat{\lambda}_A = 4$ per hour, $\hat{\lambda}_B = 5$ per hour, $\hat{\lambda}_C = 1.5$ per hour

The estimator *[it is the MLE but this need not be stated]* of r_{ij}, $\hat{r}_{ij}$, is the proportion of observed jumps out of state i to state j.

$\hat{r}_{AB} = 11/20$
$\hat{r}_{AC} = 9/20$
$\hat{r}_{BA} = 80/125 = 16/25$
$\hat{r}_{BC} = 9/25$
$\hat{r}_{CA} = 24/27 = 8/9$
$\hat{r}_{CB} = 1/9$

(b) The estimated generator matrix (in hr^{-1}) is:

$$\begin{pmatrix} -4 & 11/5 & 9/5 \\ 16/5 & -5 & 9/5 \\ 4/3 & 1/6 & -3/2 \end{pmatrix}$$

(iii) Distribution is binomial with mean $n.r_{ij}$ and variance $n.r_{ij}(1 - r_{ij})$, where n is the given number of transitions.

(iv) Null hypothesis is that the Markov property applies to successive transitions, or that the observed triplets are from a Binomial distribution with the estimated parameters (given the number of transitions to the middle state).

Using test statistic given in the hint, we can draw up the table below.

Triplet	n_{ijk}	$E = n_{ij}\hat{r}_{jk}$	$\dfrac{(n_{ijk} - E)^2}{E}$
ABC	42	39.6	0.1455
ABA	68	70.4	0.08182
ACA	85	80	0.3125
ACB	4	10	3.6
BAB	50	44	0.8182
BAC	30	36	1
BCA	38	40	0.1
BCB	7	5	0.8
CAB	64	66	0.0606
CAC	56	54	0.07407
CBA	8	9.6	0.2667
CBC	7	5.4	0.4741

Test statistic 7.7335

Under the null hypothesis, the test statistic follows a χ^2 distribution with the following number of degrees of freedom:

	Number of triplets	12
Minus	Number of pairs	6
Plus	Number of states	3
Minus	One	1
		8 degrees of freedom

The critical value of χ_8^2 at the 5% significance level is 15.51

As 7.7335 < 15.51 there is no evidence to reject the null hypothesis.

[Alternative approaches could be taken which resulted in a slightly different result for the test statistic. These were given full credit where appropriate.]

(v) *[Refer back to part (i) — the test in (iv) has only tested that there is no evidence that the destination that the next jump depends on the previous state occupied. Need to test the other assumptions].*

Holding times — are these exponentially distributed?

A chi-squared goodness of fit test would be appropriate

Is destination of jump independent of the holding time?

There is no obvious test statistic for doing this. A suitable test would be to classify jumps as being from short, medium and long holding times and investigating these graphically.

(vi) APPROXIMATE METHOD

Divide time into very short intervals, h, such that $\sigma_{ij}h$ is much less than 1.

Simulate a discrete-time Markov chain $\{Y_n : n \geq 0\}$, with transition probabilities $p^*_{ij}(h) = \delta_{ij} + h\sigma_{ij}$.

The jump process, X_t is given by $X_t = Y_{[t/h]}$.

EXACT METHOD

Simulate the jump chain as a Markov chain, with transition probabilities $p_{ij} = \sigma_{ij} / \lambda_i$.

Once the path $\{\hat{X}_n : n = 0,1,...\}$ has been generated, the holding times $\{T_n : n = 0,1,...\}$ are a sequence of independent exponential random variables, having parameter $\lambda_{\hat{X}_n}$.

END OF EXAMINERS' REPORT