

INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

April 2014 examinations

Subject CT4 – Models Core Technical

Introduction

The Examiners' Report is written by the Principal Examiner with the aim of helping candidates, both those who are sitting the examination for the first time and using past papers as a revision aid and also those who have previously failed the subject.

The Examiners are charged by Council with examining the published syllabus. The Examiners have access to the Core Reading, which is designed to interpret the syllabus, and will generally base questions around it but are not required to examine the content of Core Reading specifically or exclusively.

For numerical questions the Examiners' preferred approach to the solution is reproduced in this report; other valid approaches are given appropriate credit. For essay-style questions, particularly the open-ended questions in the later subjects, the report may contain more points than the Examiners will expect from a solution that scores full marks.

The report is written based on the legislative and regulatory context pertaining to the date that the examination was set. Candidates should take into account the possibility that circumstances may have changed if using these reports for revision.

D C Bowie
Chairman of the Board of Examiners

June 2014

General comments on Subject CT4

Subject CT4 comprises five main sections: (1) a study of the properties of models in general, and their uses for actuaries, including advantages and disadvantages (and a comparison of alternative models of the same processes); (2) stochastic processes, especially Markov chains and Markov jump processes; (3) models of a random variable measuring future lifetime; (4) the calculation of exposed to risk and the application of the principle of correspondence; (5) the reasons why mortality (or other decremental) rates are graduated, and a range of statistical tests used both to compare a set of rates with a previous experience and to test the adherence of a graduated set of rates to the original data. Throughout the subject the emphasis is on estimation and the practical application of models. Theory is kept to the minimum required in order usefully to apply the models to real problems.

Different numerical answers may be obtained to those shown in these solutions depending on whether figures obtained from tables or from calculators are used in the calculations but candidates are not penalised for this. However, candidates may be penalised where excessive rounding has been used or where insufficient working is shown.

Comments on the April 2014 paper

The general performance was slightly better than that in April and September 2013. Well-prepared candidates scored highly across the whole paper, with an above average proportion of candidates scoring 70 per cent or more. The comments that follow the questions concentrate on areas where candidates could have improved their performance. Candidates approaching the subject for the first time are advised to include these areas in their revision.

- 1** Systems with long time frames, such as the operation of a pension fund, can be studied in compressed time.

Complex systems with stochastic elements, such as the operation of a life insurance company, can be studied by simulation modelling when mathematical or logical models cannot describe them in ways which are easy to interpret.

Different future policies of possible actions can be compared to see which best suits the requirements or constraints of a user OR to test the sensitivity of profits under different scenarios.

In a model of complex systems we can usually get control over experimental conditions so that we can reduce the variance of the results from the model without upsetting their mean values.

Alternative suggestions were also given credit, for example “to improve competitiveness by testing out new underwriting approaches based on postcodes”, or “to help understand the correlation between actions and decisions”, or “to allow the user to understand better the potential impact of changes over which he or she may have little control”. Most candidates scored fairly well on this question.

- 2** (i) All our models and analyses are based on the assumption that we can observe groups of identical lives (or at least, lives whose mortality characteristics are the same).

Although in practice, this is never possible.

We can at least subdivide our data according to characteristics known, from experience, to have a significant effect on mortality.

This ought to reduce the heterogeneity of each class so formed.

- (ii) Type of policy (which often reflects the reason for insuring)
Smoker/non-smoker status
Level of underwriting
Duration in force
Sales channel
Policy size
Occupation of policyholder OR socio-economic class
Known impairments
Postcode/geographical location
Marital status

Answers to part (i) were often unconvincing. The question asked specifically about the conduct of mortality investigations: many candidates wrote about the problems of selection in heterogeneous data, and this was not given credit. Most candidates scored full marks on part (ii).

- 3** A stochastic model is one which recognises the random nature of the input components, whereas a deterministic model does not contain any random components.

In a stochastic model the output of each run is one value from a distribution. By contrast, in a deterministic model, the output is determined once the set of fixed inputs and the relationships between them have been defined.

In a stochastic model, several independent runs are required for each set of inputs so that statistical theory can be used to help study the implications of a set of inputs. A deterministic model only requires one run.

Running a stochastic model many times will produce a distribution of results for possible scenarios, whereas a deterministic model will produce results for a single scenario. Thus a deterministic model can be seen as a special case of a stochastic model.

For many stochastic models, it is necessary to use numerical approximations in order to integrate functions or solve differential equations. The results for a deterministic model can often be obtained by direct calculations.

Monte Carlo simulation is an example of a stochastic model: a collection of deterministic models each with an associated probability.

Most candidates made a reasonable attempt at this question. Some candidates wrote answers that repeated the same point using different words. Credit was only given once for each point. Comments that stochastic models are more costly or harder to interpret than deterministic models were not given any credit as they do not explain what a stochastic model is.

- 4** (i) A life alive at time t should be included in the exposed-to-risk at age x at time if and only if, were that life to die immediately, he or she would be included in the deaths data d_x at age x .
- (ii) The death rate at age 45 relates to the deaths reported aged 46 next.

Dealing with the exposure in country A first. Census data are at age last, so census data at age x correspond to the age next deaths data for age $x+1$.

$$E_x^c = \int_{1/12}^{1/13} P_{x,t} dt$$

Let P_n^A be the population aged 45 last in country A on the census date in year n .

Assuming the population varies linearly between census dates the central exposed to risk aged 45 last in the calendar year 2012 can be approximated by:

$$\begin{aligned} & 1/12 * 1/2 * (P_{12}^A + (P_{11}^A + 11/12 * (P_{12}^A - P_{11}^A))) + \\ & 11/12 * 1/2 * (P_{12}^A + (P_{12}^A + 11/12 * (P_{13}^A - P_{12}^A))) \\ & = 1/24 * (381,000 + 380,417) + 11/24 * (381,000 + 384,667) \\ & = 31,725.7 + 350,930.7 \\ & = 382,656 \end{aligned}$$

In country B we need to make the age definition of the exposure data match the deaths data.

Assuming that birthdays are spread uniformly over the calendar year, half of those aged 45 last will be aged 45 nearest and half will be aged 46 nearest.

Let P_n^B be the population of country B aged 45 last at the census date in year n then:

$$\begin{aligned} P_{11}^B &= 0.5 (374,000 + 354,000) = 364,000 \\ P_{12}^B &= 0.5 (381,000 + 372,000) = 376,500 \\ P_{13}^B &= 0.5 (385,000 + 375,000) = 380,000 \end{aligned}$$

Assuming the population varies linearly between census dates the central exposed to risk aged 45 last in the calendar year 2012 is:

$$\begin{aligned} & 7/12 * 1/2 * (P_{12}^B + (P_{11}^B + 5/12 * (P_{12}^B - P_{11}^B))) + \\ & 5/12 * 1/2 * (P_{12}^B + (P_{12}^B + 7/12 * (P_{13}^B - P_{12}^B))) \\ & = 7/24 * (369,208 + 376,500) + 5/24 * (376,500 + 377,958) \\ & = 217,498.2 + 157,178.8 \\ & = 374,677. \end{aligned}$$

Assuming the force of mortality is constant within each year of age

$$\mu_{45} = \frac{4,800}{382,656 + 374,677} = 0.006338 \text{ for the calendar year 2012.}$$

- (iii) The rate interval is the life year starting at age 45 exact.

The estimate relates to the age in the middle of the rate interval, which is 45.5 years.

In part (i) the phrase “if and only if” was required for the full mark. Part (ii) was demanding, and candidates struggled to obtain many marks. For full credit, the assumptions needed to be stated at the correct place in the argument, and not just listed at the end. In part (ii) full credit could be obtained for estimates of q_{45} , provided the initial exposed to risk was used.

A common error was to use the wrong age (i.e. 44 years last birthday). This was penalised in part (ii), but full credit could be scored for part (iii) if the answer to part (iii) was consistent with the age used in part (ii), and whether q_{45} or μ_{45} were estimated in part (ii).

A (slightly simpler) alternative answer to the exposed to risk for country A was:

$$\begin{aligned} & 1/12 * 1/2 * (P^A_{11} + P^A_{12}) + 11/12 * 1/2 * (P^A_{12} + P^A_{13}) \\ &= 1/24 (374,000 + 381,000) + 11/24 * (381,000 + 385,000) \\ &= 31,458.3 + 351,083.3 \\ &= 382,542; \end{aligned}$$

and for country B the corresponding was:

$$\begin{aligned} & 7/12 * 1/2 * (P^A_{11} + P^A_{12}) + 5/12 * 1/2 * (P^A_{12} + P^A_{13}) \\ &= 7/24 (364,000 + 376,500) + 5/24 * (376,500 + 380,000) \\ &= 215,979.2 + 157,604.2 \\ &= 373,583, \end{aligned}$$

giving a final answer of 0.00635. Because it is slightly simpler than the correct solution it did not score full credit.

- 5** (i) A person who is aged 50 years at the start of the investigation, is not a heavy drinker,
and has not lived for 12 months or more in a tropical country.

- (ii) From the third result we have:

$$\exp(\beta_T) = 3,$$

so that

$$\beta_T = \log_e 3 = 1.099.$$

From the second result we have:

$$\exp(\beta_C + \beta_T) = \exp(\beta_C) \exp(\beta_T) = 4,$$

hence

$$\exp(\beta_C) = \frac{4}{3},$$

and

$$\beta_C = \log_e \frac{4}{3} = 0.2877.$$

Finally, using the first result we have:

$$\exp(\beta_C + 10\beta_A) = \exp(\beta_C) \exp(10\beta_A) = 2,$$

$$\text{hence } \exp(10\beta_A) = \frac{6}{4},$$

$$\text{and } \beta_A = \frac{1}{10} \log_e \frac{6}{4} = 0.0405.$$

- (iii) The chance of a 50-year old non-heavy-drinking person who has always lived in the UK remaining free of the disease for 10 years is:

$$S_{T=0}(10) = \exp\left(-\int_0^{10} h_0(t) dt\right) = 0.8.$$

The chance of a person of the same age and drinking habits who has lived for more than 12 months in a tropical country remaining free of the disease for 10 years is therefore:

$$S_{T=1}(10) = \exp\left(-\int_0^{10} h_0(t) e^{\beta_T} dt\right) = 0.8 e^{\beta_T} = 0.8^3 = 0.512.$$

Many candidates scored highly on this question, particularly on part (ii). In part (i) very few candidates wrote that the baseline hazard referred to a person aged 50 at the start of the investigation. Candidates who missed this important point were penalised. In part (iii) full credit was given to answers that evaluated $-\int_0^{10} h_0(t) dt$, though this was not necessary.

However, candidates who assumed that the baseline hazard was constant were penalised, as this assumption is neither necessary nor correct for the Cox regression model.

- 6** (i) Suppose we observe n individuals ($i = 1, \dots, n$), for a period E and the number of events observed to happen to individual i is d_i .

Then the Poisson likelihood is:

$$\prod_{i=1}^n \frac{(\mu E)^{d_i} e^{(-\mu E)}}{d_i!}.$$

Taking logarithms of the likelihood we have:

$$\log_e L = \sum_{i=1}^n d_i \log \mu + \sum_{i=1}^n d_i \log E - \sum_{i=1}^n E\mu - \sum_{i=1}^n \log_e (d_i!).$$

Differentiating with respect to μ we obtain:

$$\frac{\partial(\log_e L)}{\partial \mu} = \frac{\sum_{i=1}^n d_i}{\mu} - nE.$$

Setting this to zero and solving gives $\hat{\mu} = \frac{\sum_{i=1}^n d_i}{nE}$.

Since the second derivative $-\frac{\sum_{i=1}^n d_i}{\mu^2}$ is negative, we have a maximum.

- (ii) The number of minutes exposed to risk for each student is given in the table below.

<i>Student</i>	<i>Exposed to risk (minutes)</i>
1	5
2	25
3	10
4	5
5	10
6	5
7	10
8	20
9	30
10	60

The total length of time exposed to risk is thus 180 minutes, or 3 hours.

During this time, 4 buses arrived.

The maximum likelihood estimate is thus $4/3 = 1.33$ buses per hour.

- (iii) The Poisson model normally assumes a fixed exposed-to-risk for each person, but in this investigation the waiting times vary with the student.

This is not a problem provided we can regard the students as identical, and we replace each student who catches a bus with an identical student at the moment the bus leaves.

But in this study the students were not replaced.

In the investigation above, there are gaps when no student was at the bus stop (e.g. between 4.50 and 4.55 p.m.). Buses may have arrived during these gaps.

Had we observed these buses, our estimate of the rate may have been different.

The assumption that the arrival of buses follows a Poisson process may not be valid as arrival times may not be independent due to traffic conditions, and/or they may not be random due to timetabling.


Answers to this question were very disappointing, with a substantial minority of candidates making only token attempts. In part (i) many candidates wrote answers in terms of a general Poisson parameter λ , rather than μE . This attracted a modest penalty. Full credit could be obtained in part (i) for candidates who used the total number of deaths, $\sum_{i=1}^n d_i$, and noted that

the sum of n independent Poisson variables is Poisson, so the likelihood $L = e^{-n\mu E} (n\mu E)^{\sum_{i=1}^n d_i}$.

Many candidates obtained credit for proceeding to derive a maximum likelihood estimator where the expression for the likelihood was incorrect.

A very common error in part (ii) was to assume that six buses had arrived rather than four, two buses having arrived at 4.35 p.m. and two at 4.50 p.m. This was penalised, as the question explicitly stated that "only one bus arrived at any given time". Less common, but still present in too many scripts, was the serious error of only including exposure for students who caught the bus, and omitting the exposure for those who were censored. In part (iii) most candidates' comments were restricted to the appropriateness of the assumption that the arrival of buses followed a Poisson process.

7 (i)



[2]

(ii) With the state numbers in the diagram above we can write:

EITHER

Using the Markov assumption

OR

The Chapman Kolmogorov equation is

$${}_{dt+t}p_x^{24} = {}_t p_x^{21} {}_{dt}p_{x+t}^{14} + {}_t p_x^{22} {}_{dt}p_{x+t}^{24} + {}_t p_x^{23} {}_{dt}p_{x+t}^{34} + {}_t p_x^{24} {}_{dt}p_{x+t}^{44}.$$

$$\text{But } {}_{dt}p_{x+t}^{34} = 0$$

$$\text{and } {}_{dt}p_{x+t}^{44} = 1.$$

So:

$${}_{dt+t}p_x^{24} = {}_t p_x^{21} {}_{dt}p_{x+t}^{14} + {}_t p_x^{22} {}_{dt}p_{x+t}^{24} + {}_t p_x^{24}.$$

Assuming that, for small dt

$${}_{dt}p_{x+t}^{ij} = \mu_{x+t}^{ij} dt + o(dt)$$

$$\text{where } \lim_{dt \rightarrow 0} \frac{o(dt)}{dt} = 0,$$

then substituting, we have

$${}_{dt+t}p_x^{24} = {}_t p_x^{21} \mu_{x+t}^{14} dt + {}_t p_x^{22} \mu_{x+t}^{24} dt + {}_t p_x^{24} + o(dt)$$

$$\text{so that } {}_{dt+t}p_x^{24} - {}_t p_x^{24} = {}_t p_x^{21} \mu_{x+t}^{14} dt + {}_t p_x^{22} \mu_{x+t}^{24} dt + o(dt)$$

$$\text{and hence } \frac{d}{dt}({}_t p_x^{24}) = \lim_{dt \rightarrow 0} \frac{{}_{t+dt}p_x^{24} - {}_t p_x^{24}}{dt} = {}_t p_x^{21} \mu_{x+t}^{14} + {}_t p_x^{22} \mu_{x+t}^{24}.$$

- (iii) The maximum likelihood estimate (MLE) of the death rate from heart disease for persons with the condition is $\frac{25}{1,139} = 0.02195$.

The MLE of the death rate from heart disease for persons without the condition is $\frac{10}{2,046} = 0.00489$.

An estimate of the variance of the maximum likelihood estimator of the death rate from heart disease for persons with the condition is

$$\frac{0.02195}{1,139} = 0.000019271.$$

An estimate of the variance of the maximum likelihood estimator of the death rate from heart disease for persons without the condition is

$$\frac{0.00489}{2,046} = 0.00000239.$$

The null hypothesis H_0 is that there is no difference between the means.

The variance of the difference between the two estimates is therefore:
 $0.000019271 + 0.00000239 = 0.000021659.$

THEN EITHER

A 95% confidence interval around the difference is therefore:

$$\begin{aligned} & (0.02195 - 0.00489) \pm 1.96\sqrt{0.000021659} \\ & = 0.01706 \pm 1.96*0.004654 \\ & = 0.01706 \pm 0.009122 \\ & = (0.007938, 0.026182) \end{aligned}$$

which does not include zero

OR

Under the null hypothesis the difference \sim Normal (0, 0.000021659).

Our observed value of the difference is $0.021949 - 0.004888 = 0.01706$
A z-score for the actual difference of 0.01706 is therefore
 $(0.01706/\sqrt{0.000021659}) = 3.67$

and since this is greater than 1.96 we reject the null hypothesis at the 95% level

THEN

so the difference is statistically significantly different from zero

Most candidates scored highly on part (i), though attempts at part (ii) were more variable. For part (iii) most candidates correctly computed the estimated rates and a smaller (though substantial) number computed the correct variances. Few candidates, however, attempted a formal test of the statistical significance of the difference.

In part (iii) some candidates calculated 95% confidence intervals around each estimate and argued that since these do not overlap, the difference between the two estimates is statistically significant. This approach will not always produce the same conclusion as testing the difference directly (because $\sqrt{X} + \sqrt{Y} \neq \sqrt{X + Y}$). It was given partial credit.

8 (i) EITHER

Censoring is the mechanism which prevents us from knowing when an individual entered the investigation or the exact date of death.

OR

We do not know the exact duration for an individual, only that it lies within some range.

- (ii) Right-censoring cuts short the investigation in progress so we do not know exactly when the event of interest happened, we only know it happened after a certain date.

An example of this might be in a mortality investigation conducted over a period of one year, all those still alive at the end of the year will die some time after the end of the investigation, but we do not know when.

Left-censoring prevents us from knowing when entry into the state which we wish to observe took place.

An example arises in medical studies in which patients are subject to regular examinations. Discovery of a condition tells us only that the onset fell in the period since the previous examination, the time elapsed since onset has been left censored.

Interval-censoring happens if we can only say that an event of interest fell within some interval of time, rather than exactly when it happened.

For example in a mortality investigation when we only know the calendar year of death rather than the precise date of death.

- (iii) Right-censoring is present as the observation was cut short while in progress for those toys which were unplugged, taken and which remained working at the end of the trial.

Type I censoring is present as the trial ended at a predetermined time, so all those toys still working were Type I censored.

The censoring is likely to be non-informative censoring. The toys which were unplugged and taken are unlikely to have any special features such as working for longer or shorter overall than the rest of the toys in the trial.

Random censoring is present as the action of the cleaner censored the toys at times which were random.

(iv) Rearranging the data:

Hour	0	4	10	11	13	31
Toys in trial	500	500	488	471	446	443
No. of exits	0	12	17	25	3	8
Reason for exit		fail	unplugged	fail	taken	fail


The Nelson-Aalen estimate for Λ is $\Lambda_t = \sum_{x_j \leq x} \frac{d_j}{n_j}$.

t_j	n_j	d_j	c_j	d_j/n_j	Λ_t
0	500	0	0		
4	500	12	0	12/500	0.024000
10	488	0	17		
11	471	25	0	25/471	0.077079
13	446	0	3		
31	443	8	0	8/443	.095137

Since $S(t) = \exp(-\Lambda_t)$ we have:

t	$S(t)$
$0 \leq t < 4$	1
$4 \leq t < 11$	0.976286
$11 \leq t < 31$	0.925817
$31 \leq t < 32$	0.909248

(v)



(vi) We do not know the length of time for which a new toy has a 60% chance of surviving, only that it is some time in excess of 32 hours.

Answers to part (i) were very disappointing. Many candidates said that censoring was when individuals were removed from the investigation for reasons other than death. This was

given no credit as it is not a definition. Most candidates scored better on part (ii), though fewer could define and explain left and interval censoring than were able to define right censoring, and many candidates could have scored more highly by giving more precise definitions. Most candidates scored highly on part (iii) and the calculation in part (iv). A common error in part (iv) was to suggest that $S(t)$ remained constant for an indeterminate time after 31 hours. As we have no information after 32 hours, the upper limit of the range for which $S(t)$ is estimated should be 32 hours. Only a minority of candidates answered part (vi) correctly.

9 (i) Smoothness,

EITHER because we are likely to want to use the data for financial calculations, and clients expect these to progress smoothly.

OR because we believe the underlying quantities vary smoothly with age.

Adherence to data

because we want the graduated rates to reflect as closely as possible the experience on which they are based.

Suitability for the purpose to hand

EITHER In life insurance work, losses result from premature deaths (benefits are paid sooner than expected) so we must not underestimate mortality,

OR In annuity work, losses result from delayed deaths (benefits are paid for longer than expected) so we must not overestimate mortality.

(ii) (a) To test for overall goodness of fit we use the χ^2 test.

The null hypothesis is that the graduated rates are not significantly different from the underlying rates in the new experiences.

The test statistic $\sum_x z_x^2 \approx \chi_m^2$ where m is the degrees of freedom.

The calculations are below:

<i>Age deviation</i>	<i>Standardised deviation</i>	<i>Standardised deviation (squared)</i>
60	2.40	5.760
61	0.08	0.006
62	0.80	0.640
63	0.76	0.578
64	1.04	1.082
65	0.77	0.593
66	1.30	1.690
67	1.76	3.098
68	0.28	0.078
69	0.68	0.462
70	0.93	0.865
Sum	14.852	

The observed value of the test statistic is 14.852.

We compare this with the critical value of the χ^2 distribution with degrees of freedom equal to the number of ages minus at least 1 for the choice of standard table.

With 8 degrees of freedom the critical value at the 95% significance level is 15.51 [with 9 d.f. it is 16.92, and with 10 d.f. 18.31].

Since $14.852 < 15.51$ [or 16.92]

we do not reject the null hypothesis.

- (b) Overall, the graduated rates seem to represent the underlying rates in the new experience.
- (iii) There may be individual ages at which the graduated rates and the observed rates differ substantially, but if these ages are a small proportion of the whole the chi-squared test may not detect them.

Use Individual Standardised Deviations Test.

There may be a consistent but small bias in one direction, but the deviations are not large enough to allow detection by the chi-squared test.

Use Signs Test or Cumulative Deviations Test.

There may be bias over sections of the age range.

Use Grouping of Signs Test or Serial Correlations test, or Cumulative.

Deviations Test over sections of the age range (chosen independently of the pattern of deviations).

The graduation may not be sufficiently smooth if the linking function is complex.

Use Third Differences Test.

- (iv) We can carry out a modified version of the Individual Standardised Deviations Test.

Under the null hypothesis we expect the individual standardised deviations to have a Normal (0,1) distribution.

Only 1 in 20 of the z_x s should lie above 1.96 in absolute value

OR

None should lie above 3 in absolute value.

OR

Range	0,1	1,2	2,3
Expected	7.5	3.1	0.4
Actual	7	3	1

We have one deviation with an absolute value of 2.4.

OR

We have no deviations above 3 in absolute value.

OR

The distribution of the deviations is close to that we might expect under a Normal (0,1) distribution

therefore we have no strong reason to reject the null hypothesis

In part (i), most candidates mentioned smoothness and adherence to data, but fewer mentioned suitability for the purpose to hand. Most candidates scored well on part (ii), the most widespread error being a failure to deduct degrees of freedom for the choice of standard table. Answers to part (iii) were sometimes vague. Few candidates realised that the chi-squared test only fails to detect small bias (it will detect large bias), and few mentioned smoothness. Some candidates used the terms “overgraduation” and “undergraduation” incorrectly to refer to situations where the graduated rates are systematically biased above and below the true rates respectively.

Most candidates correctly identified the individual standardised deviations test as the only test which could be carried out in part (iv).

10 (i) Markov chain.

(ii) (a) It is not irreducible

because a heating element cannot move to a state of being in better condition.

(b) It is not periodic

because it can remain in each state (or any other suitable reason).

(iii) EITHER

The second order transition matrix is:

0.25	0.2	0.26	0.29
0	0.25	0.3	0.45
0	0	0.25	0.75
0	0	0	1

Hence probability in Poor condition at the second inspection is 0.26.

OR

The required probability is equal to

Prob [Excellent to Excellent to Poor] +
 Prob [Excellent to Good to Poor] +
 Prob [Excellent to Poor to Poor]

which is $(0.5 \times 0.2) + (0.2 \times 0.3) + (0.2 \times 0.5) = 0.26$.

(iv)

	<i>Excellent</i>	<i>Good</i>	<i>Poor</i>
Excellent	0.6	0.2	0.2
Good	0.2	0.5	0.3
Poor	0.5	0	0.5

(v) Long-term probabilities satisfy $\pi = \pi P$.

$$0.6\pi_E + 0.2\pi_G + 0.5\pi_P = \pi_E \quad (1)$$

$$0.2\pi_E + 0.5\pi_G = \pi_G \quad (2)$$

$$0.2\pi_E + 0.3\pi_G + 0.5\pi_P = \pi_P \quad (3)$$

Also $\pi_E + \pi_G + \pi_P = 1$.

(2)–(3) gives:

$$\pi_G = \frac{5}{8}\pi_P.$$

$$\text{So } \pi_E = \frac{25}{16}\pi_P.$$

$$\text{Hence } \left(\frac{25}{16} + \frac{5}{8} + 1\right)\pi_P = 1.$$

$$\text{Stationary distribution } \pi_E = \frac{25}{51}, \pi_G = \frac{10}{51}, \pi_P = \frac{16}{51}.$$

(vi) The expected number of failures of heating elements is:

$$(0.1\pi_E + 0.2\pi_G + 0.5\pi_P) * 100 = 24.51.$$

The cost of each failure is £1,050 so the expected cost over a year is £25,735.

(vii) The transition matrix for the condition of the element at the start of each cycle will now be:

	<i>Excellent</i>	<i>Good</i>
<i>Excellent</i>	0.8	0.2
<i>Good</i>	0.5	0.5

The revised stationary distribution satisfies $\rho = \rho P$

$$0.8\rho_E + 0.5\rho_G = \rho_E \quad (1)$$

$$0.2\rho_E + 0.5\rho_G = \rho_G \quad (2)$$

$$0.2(1 - \rho_G) + 0.5\rho_G = \rho_G$$

$$\rho_E = \frac{5}{7}, \rho_G = \frac{2}{7}$$

Expected cost of failures is now:

$$(0.1\rho_E + 0.2\rho_G) * 100 * 1050 = £13,500.$$

But we also now have extra heating element replacement costs of:

$$(0.2p_E + 0.3p_G) * 100 * 50 = £1,143.$$

So overall profits have improved by:

$$£25,735 - £13,500 - £1,143 = £11,092.$$

In part (i) "Markov jump chain" is not correct as the question makes reference to time, and the Markov jump chain loses the information about the timing of the transitions. Answers to part (ii) were very good. In part (iii) the full matrix was not required but some indication of where the numbers have come from or what they are was needed. In parts (iv) onwards many candidates offered a 4×4 matrix as a solution as follows:

	<i>Excellent</i>	<i>Good</i>	<i>Poor</i>	<i>Failed</i>
<i>Excellent</i>	0.5	0.2	0.2	0.1
<i>Good</i>	0	0.5	0.3	0.2
<i>Poor</i>	0	0	0.5	0.5
<i>Failed</i>	1	0	0	0

this is incorrect because it implicitly assumes that the kiln is run for a complete cycle with a failed element. It was penalised in part (iv) but full credit could be scored in part (v) for the correctly followed-through stationary distribution.

The better prepared candidates scored highly on this question.

END OF EXAMINERS' REPORT