ENGINEERING TRIPOS PART IIA 2022
CRIB WITH COMMENTS FROM ASSESSOR'S REPORT
MODULE 3E3: MODELLING RISK

Crib

Question 1

Part (a):

Answer is C: 180

Part (b):

Answer is H

Part (c)

Answer is F

Part (d)

Answer is B

Part (e)

Answer is B

Part (f)

Answer is E

Part (g)

Answer is B

Part (h)

Answer is F

Question 2:

Part (a)

The primary decision is whether to give out the loan or not.

Part (b)

Denote D=Default and ND=No Default and let P(ND)=x. Also denote all monetary values in thousands for simplicity (e.g., 1.5 instead of £1500, 100 instead of £100000).

Then, the payoff to the firm if a loan is given out is 111x and payoff without a loan is 100.

$111x>100$ when $x>100/111$.

The decision can be characterised as: Give loan when $x>100/111$ for an EMV of 111x and do not give out the loan otherwise for an EMV of 100.

Part (c)

Expected value of sample information (EVSI) is positive if we give a loan after a positive report and don't give a loan after a negative report. Otherwise, EVSI=0.

For the case when EVSI is positive, we can draw the tree of Figure A, with P(-) and P(+) denoting the probability of a negative and positive report.

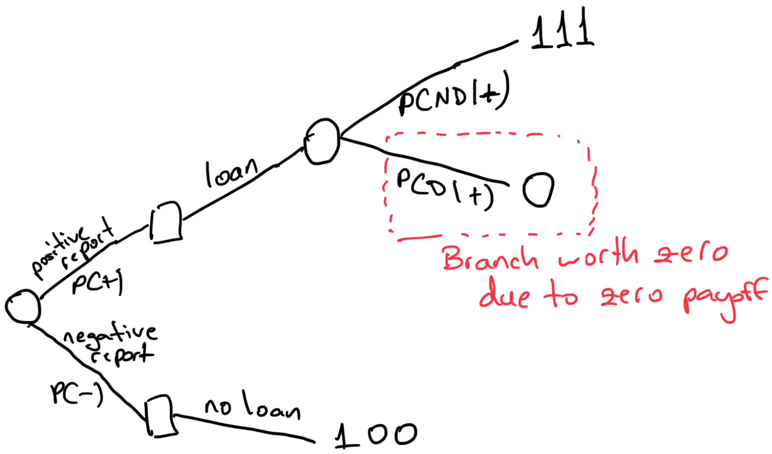

Fig A: Tree with Sample Information, when EVSI is Positive

Note:

$P(+)=P(+|ND)P(ND)+P(+|D)P(D)=(0.8)x + (0.3)(1-x)=0.3+0.5x$

And $P(-) = 1-P(+) = 0.7-0.5x$

Substituting P(+) and P(-) below gives the value of the tree of Figure A as

$$\begin{aligned} Tree\ Value\ (Fig.\ A) &= 100P(-) + (111)\frac{P(+|ND)P(ND)}{P(+)}P(+) \\ &= 100P(-) + (111)P(+|ND)P(ND) \\ &= (100)(0.7 - 0.5x) + (111)(0.8)(x) \\ &= 70 + 38.8x \end{aligned}$$

From part (b), EMV=111*x* if *x*>100/111 and 100 otherwise. Thus,

1) When *x*>111/100, the information is worth its cost of 1.5 if 70+38.8*x*-111*x*>1.5, which simplifies to *68.5-72.2x>0.*
2) When *x*<=*111/100*, the information is worth its cost of 1.5 if 70+38.8*x*-10>1.5, which simplifies to *38.8x-31.5>0.*

When conditions 1 or 2 are not satisfied, the information is not worth its cost.

Part (d)

With perfect information, the probability of a positive report is equal to the probability of no default (*x*). This gives us the tree in figure B.

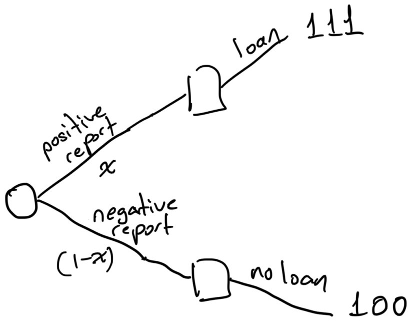

Fig B: Tree with Perfect Information.

Tree Value (Fig. B)=111*x* + (100) (1-*x*) = 100+11*x*

We need to subtract the appropriate EMV (111x or 100) from part (b) to get the expected value of perfect information (EVPI):

If *x*>100/111, EVPI=100+11*x*-111*x*= (1-*x*)100

If *x*<=100/111, EVPI=100+11*x*-100=11*x*

Question 3:

Part (a)

Winston should recommend pooled security, as pooled queues always dominate dedicated queues in terms of waiting times and the number in the queue. An exception could be the loss in efficiency if people need to travel a long distance from the head of the pooled queue to the respective security lane.

Part (b)

To satisfy the target, we start with 2 lanes and increase the number of lanes until the target is met.

λ_p = 0.9/min
Denote service time by τ. $(1/\mu) = \tau = 2$ mins . then $\rho = \lambda_p \tau/s = 0.9$.
From table C, $L_q = 7.6737<8$.
Hence, just two security lanes are enough.

Part (c)

$\lambda = 0.45$/min $\tau = 2$ mins $\rho = \lambda \tau = 0.9$.
$W_q = \tau \rho/(1-\rho) (CV_A^2+CV_S^2)/2 = 18$ minutes
$Lq = \lambda W_q = 0.45X18 = 8.1$ people per queue. Since we have 2 queues, we have 16.2 people waiting on average.

Part (d)

Customer Type	Frequent Flyer	Regular
Arrival rate λ	0.9*0.65=0.585	0.9*0.65=0.315
Service rate μ	1/1.5=0.667	1/2/93=0.3413
s	1	1
ρ=λ/(μ*s)	0.8775	0.9230
Lq (from table C)	6.45	10.58
Wq=Lq/λ (By Little's Law)	11.02564	33.5873
W= Wq + Service Time	12.52564	36.5173

Performance Metrics:

Average time spent in the queue by a passenger: (0.65)(11.02564)+(0.35)(33.5873) =18.9222

Average time total time spent by a passenger: (0.65)(12.52564)+(0.35)(36.5173) =20.9227

Average number of passengers waiting in the queue: 6.45+10.58=17.03

Overall, this system is slightly worse off than part C. However, in this system, the frequent travellers are significantly better off.

Further Suggestion: The faster travel lane can be used as a way to provide value added service. The airlines can be charged for this service provided by the airport. The extra revenue can be used to improve service for regular passengers and thereby promote equity.

Part (e)

Resource pool	Calculation	Capacity (trays/hour)
Positioning	Capacity = 3 × (1/8 min) × (60 min/hour)	22.5 trays/hour
Washing and Drying	Capacity = 10 × (1/60 min) × (60 min/hour)	10 trays/hour
Inspecting and Packing	Service time = 80%(9 min) + 20% (22 min) = 7.2 + 4.4 = 11.6 min Capacity =2 × (1/11.6min) × (60 min/hour)	10.3 trays/hour
Sterilizing	Capacity = (24 trays/batch)/(2.5 hours/ batch)	9.6 trays/hour

Sterilizing is the bottleneck. Therefore, the capacity of the system is 9.6 trays/hour.

Part (f)

To achieve a system capacity of 12 trays/hour, we need to increase the capacity of the second, third, and fourth steps because their current capacities are lower than 12 trays/hour.

- For the second step, we need to purchase at least two more washing and drying machines to reach a capacity of 12 trays/hour.
- For the third step, we need a third person, which will bring the capacity to 3(1/11.6 min) × (60 min/hour) = 15.5 trays/hour.
- For the fourth step, we will need to buy another sterilizing oven which would double the rate to 19.2 trays/hour.

Some possible improvements over this plan are possible: a) Only two people are needed in the first step so moving one person from step 1 to step 3 may be possible depending on the skill requirements. b) We may need a third additional washer/dryer since the rate for this resource pool is exactly the required rate. c) It may be possible to buy a smaller sterilizer as a capacity of 19.2 trays/hour is more than necessary.

Part (g)

Reducing the damage rate from 20% to 5% would change the service time to 95%(9 min) + 5% (22 min) = 9.65 min in step 3 and this would increase the capacity to 2 (1/9.65 min) × (60 min/hour) = 12.4 trays/hour with the current two people. Thus, the current two people could keep up with a demand rate of 12 trays/hour.

Therefore, if the damage rate could be reduced, the savings would be the wages of an entire person (remember that we had to hire an additional worker to keep up with the demand rate of 12 trays/hour if the damage rate had stayed the same, see part (b)).

Dean has merely calculated the time saved and not recognized that there will be idle time for the third person required to handle the 12/hour rate in the previous setting. Thus, the savings would be 8 hours of one person's time multiplied by that person's wage rate (recall that union contracts prohibit us from hiring part-time employees for this task).

Question 4

Part (a)

Y=19,998+23*2022+.02972*1800000=120,000. The expected annual demand in region 1 is 120,000 tons. Standard error for the model is given as 40,000. Then the 95% confidence interval is given by 120,000 +/- 2*40,000 tons, i.e. [40000,200000].

Part (b)

We can see that the confidence intervals for the variables Year and Advertising include zero. This means that neither variable is significant. Hence, the variables are not that good at explaining annual demand.

We can also see that the r-square is .2209. In other words, the model explains 22.09% of the variability in demand. This is relatively low relative to forecasting models we built in class. A mitigating factor could be how volatile demand inherently is for the product at hand.

Part (c)

Note: Elaborating on two problems is sufficient. Three are provided below. Other reasonable answers can be accepted as correct.

Problem 1:
It seems from the positive coefficient for Year that the sales are increasing every year. As the firm grows, the amount it spends on advertising may also grow. That would make the two variables (Year and Advertising) highly correlated. That, in turn, leads to a multicollinearity problem. A multicollinearity problem could explain the insignificance of the two explanatory variables.

One would need to see the correlation table to have a better understanding of the extent of this problem. If they are indeed highly correlated I would run regressions with each of the explanatory variables separately to select the best model.

Problem 2:
By definition, the data are from previous years and we are predicting demand for the year 2022. That means that we are extrapolating. Extrapolation is dangerous.

To address the issue, I would want to see a regression with a hold-out sample of one year and compare the residuals of that model to those of the current model.

Problem 3:
Note that we also do not know if the firm has previously spent more than the current year's advertising budget. Again, we may be extrapolating. To address this issue, I would need to see a table of previous years' advertising budgets.

Part (d)

We need to take into account the "reactive" capacity that costs 600 + 150 + 50 = 800 Rand per ton. The Overall cost structure is Co = c – s = 600 – 300 = 300 Rand. Cu is 1000 – (1000 – 800) – 600 = 200 Rand, where p=1000 is diminished by (1000-800) = 200 Rand to account for the guaranteed revenue to the distribution centre by using reactive capacity, if needed.

So F* = Cu/(Cu+Co) = 200/500 = 0.4 is the critical fractile. From Table A, z* = -0.25. Expected monthly demand (given assumptions listed in the question) is 120,000/12 and standard deviation is 40,000/sqrt(12). The optimal inventory level is therefore

$$Q^* = \mu + z^* \sigma = \frac{120000}{12} - 0.25\frac{40000}{\sqrt{12}} = 7113 \text{ tons.}$$

Each distribution centre should thus order 7113 tons per month.

Part (e)

Let's call the orders fulfilled by the initial order "primary sales" and note that additional demand can be met with reactive orders while any excess inventory is sold at salvage value. L_n values are given by Table B.

$$\text{E[reactive orders]} = \text{E[lost sales from primary]} = \sigma L_n(z^*) = \frac{40000}{\sqrt{12}} \times L_n(-0.25)$$

= 11547 x 0.53634 = 6193

E[primary sales] = E[demand] – E[reactive orders] = 10000 – 6193 = 3807

E[Salvage] = amount ordered – E[primary sales] = 7113 – 3807 = 3306

E[profit] = p x E[Demand] + s x E[salvage] – c x Q* - 800 x E[reactive orders]
= 1000 x 10 000 + 300 x 3306 – 600 x 7113 – 800 x 6193 Rand
= 1,769,600 Rand

So each of the 5 distribution centres make 1.7696 Million Rand for a total of 8,848,000 (8.848 Million) Rand per month.

But the main centre makes 50 x 6193 = 309,650 Rand per distribution centre during the transfer associated with reactive orders, so the total profit for the whole group is:

8,848,000 + 5*309,650 = <u>10,396,250 Rand per month</u>

Part (f)

Qualitatively, the 50 Rand/ton transfer for reactive capacity leads to waste. The artificially high reactive ordering costs cause the distribution centres to order more initially and rely less on reactive capacity. Lowering the reactive capacity cost will reduce the amount ordered initially, and eliminating this waste should increase overall profits.

Co stays the same: Co = c − s = 600 − 300 = 300 Rand. Removing the 50 Rand reactive capacity cost means Cu = 750-600 = 150 Rand, instead of 200 Rand.

Then, F* = Cu/Cu+Co = 150 / 450 = 0.33, so the z* = -0.43. The new inventory order level is

$Q^* = \mu + z^* \sigma$ = 10000 – 0.43 x 11547 = 5035 tons.

We can now recompute expected profits with this new order level:

E[reactive orders] = E[lost sales from primary] = $\sigma L_n(z^*)$ = 11 547 x $L_n(-0.43)$

= 11 547 x 0.65026 = 7 509 tons

E[primary sales] = E[demand] – E[reactive orders] = 10 000 – 7 509 = 2 491 tons

E[salvage] = amount ordered – E[primary sales] = 5 035 – 2 491 = 2 544 tons

E[profit] = p x E[Demand] + s x E[salvage] – c x Q* - 750 x E[reactive orders]

= 1000 x 10 000 + 300 x 2 544 – 600 x 5 035 – 750 x 7 509 Rand

= 2 110 450 Rand

So each of the 5 distribution centres make 2 110 450 for a total of

<u>10,552,250 Rand per month.</u>

The main centre no longer contributes further revenues from the receipt of money from the distribution centres.

So the improvement from eliminating this waste increases expected profits by 10,552,250 − 10,396,250 = 156,000 Rand per month, or 1,872,000 Rand per year.

Comments from Assessor's Report

Question 1:

No. candidates attempted: 82 (with IB), 5 (without IB). Attempted by all candidates as Questions 1 and 2 were in the section of the exam meant for all students.

Part (a) assessed students' understanding of Little's Law. Part (b) assessed students' understanding of queueing theory. Part (c) assessed students' ability to identify different contributors to the quality of a forecasting process such as behavioral biases and incentives. Part (d) assessed students' understanding of process flow analysis and the impact of bottlenecks on the output of a process with a quality risks, and the impact of technology improvements on throughput. Part (e) tested students understanding of inventory models and the impact of a change in business models on the order quantity of an organization facing uncertainty. Part (f) tested students' ability to make tradeoffs between risk and return across project portfolios and to identify portfolios on the efficient frontier. Part (g) asked students to calculate the expected value of potential choices for a risk-neutral individual and identify that they should be indifferent between a subset of the choices. Part (f) asked about the same set of choices (as the previous part), but this time for a risk-averse individual.

Of the multiple-choice questions (i.e., Question 1), a vast majority of students selected the correct choice for parts (a), (c), (e), (f), and (g). The parts that students found more difficult were (b), (d), and (h).

Part (b) had 8 choice options, which is more than most other parts. This created more opportunities to make mistakes. In the question, three statements were made about the likely impact of a change to a queueing system on the utilization of servers (physicians), interarrival-times, and the average number of patients in service. The students who answered correctly realized that all three statements were false. Those who made mistakes typically answered either that the second statement (on inter-arrival times) or the third statement (on average number of patients in service) were false.

Part (d) asked about a process improvement that would combine two steps of a process which took 25 seconds and 15 seconds into a single step, but performed the combined step in 30 seconds. The students that answered correctly realized that the 10 second “improvement” would actually hurt throughput because the slowest step of the process would now take 30 seconds instead of 25 seconds and throughput is determined by the slowest step. Therefore, the correct answer about whether to adopt this new technology was “No”. Of those who got the answer wrong, some students saw the improvement in total time taken for all steps and simply answered “Yes”. Others may have considered that costs could also matter and selected that the answer could not be determined without further information. However, the question asked only about the impact on throughput, without consideration for costs.

Part (h) had 6 choice options, which is more than most other parts. This created more opportunities to make mistakes. Most students answered the related question of part (g) correctly and selected that a risk-neutral individual should be indifferent between 4 possible decisions, all of which gave the same positive expected (and equally risky) payoff. In part (h) those who answered wrongly selected the option D, which had a certain payoff of zero. This assumes that the risk-averse individual would prefer a payoff of zero instead of a positive expected (risky) payoff. This is not necessarily true as the question does not state how risk-averse the individual is. Therefore, the correct answer was that the answer could not be determined without additional information.

Question 2:

No. candidates attempted: 82 (with IB), 5 (without IB). Attempted by all candidates as Questions 1 and 2 were in the section of the exam meant for all students.

Part (a) was answered correctly by a vast majority of students who identified that the core decision was whether to give out a loan or not. The students who answered it incorrectly talked about carrying out credit assessments without mentioning what the core decision (to give a loan or not) should be.

Parts (b) (c) and (d) asked students to evaluate the core decision without any further information, with imperfect information and perfect information, respectively. Part (b) could have been solved with or without drawing a decision tree, a decision tree was necessary for parts (c) and (d). The students who answered correctly characterized the full set of decisions, drew the appropriate decision tree, and entered the correct set of probabilities at each of the branches. Some chose to do this with an arbitrary probability assigned to the probability of no default, expressed as a variable *x* or assumed an appropriate probability for "no default" [for example, 0.85] and correctly calculated all other probabilities based on that assumption and the information given in the question.

In parts (b), (c) and (d), the majority of the students drew the decision tree correctly. Those who made mistakes typically did not calculate the probabilities or expected values correctly, which led to mistakes. A minority of students based their calculations on a wrong or incomplete decision tree, which led to incorrect conclusions.

In part (c), an additional source of error was that some students failed to identify that some of the decision branches would never be chosen. Those branches did not need to be considered. Taking those branches into account made the calculations unnecessarily lengthy and led to mistakes.

It is also worth noting that while part (d) was simpler to solve than part (c), it was attempted by fewer students. This could be due to students running out of time towards the end of the exam if they chose to leave Question 2 to the end.

Question 3:

No. candidates attempted: 45 (with IB), 3 (without IB). Attempted by slightly over half of candidates as Questions 3 and 4 were in the section of the exam where students opted to answer one question or the other.

Most of the students answered part (a) correctly by explaining the logic of pooling security, making the necessary recommendation, and listing an appropriate exception to the superior performance of the pooled queue. Those who answered wrongly either did not remember that pooled queues always outperform dedicated queues on the core performance metrics discussed in the course or failed to identify an exception (e.g. the need to travel a longer distance from the head of the pooled queue to the respective security lane could be a downside that is not taken into account by the performance metrics).

In part (b), most of the students calculated the correct number of security lanes that were needed (i.e. 2). Those who answered wrongly either did not round to a whole number of security lanes or calculated a higher number of security lanes.

For part (c), most of the students calculated the number of people in both the queues (i.e. 16.2). A handful of them only calculated the number of people in one queue (8.1) and ignored that there were 2 lanes.

For part (d), most of the students had made the correct suggestions for frequent and non-frequent flyer systems described within the question. They also compared the situation presented in this part with that of the previous part (3c) to identify which system was better. Some students also made suggestions on the possibility of providing a faster travel lane as a value-added service. While most of the students had these recommendations explained qualitatively, the performance metrics calculated were sometimes different from the correct answer.

For part (e), most of the students calculated the capacity of the system correctly by systematically calculating the capacity across the resource pool. They were also able to identify that ‘sterilizing’ is the bottleneck within the system. The students who gave incorrect suggestions calculated the service time incorrectly. The difficulty there was the probabilistic nature of the service times for different types of inflow.

In part (f), for the most part, all the students answered these questions correctly and gave logical explanations for the changes/additions to be made to meet the updated demand. They correctly identified that washing and drying need two more people and that, inspecting and parking would need a third person to keep up with the demand. Some students gave the direction in which the number of resources in each step should increase or decrease to but did not calculate the exact quantity. A minor error was to only mention what was necessary for meeting the output requirements (i.e. increases in capacity at some steps) without acknowledging excess capacity at another, which would lead to unnecessary costs.

In part (g), some of the students answered the question correctly by explaining why they thought that Dean's calculations were incorrect/incomplete. Others assumed that the calculations provided by Dean (a character in the question) were correct. An effort to justify those wrong calculations led to an incorrect response by the students.

Question 4:

No. candidates attempted: 37 (with IB), 2 (without IB). Attempted by slightly under half of candidates as Questions 3 and 4 were in the section of the exam where students opted to answer one question or the other.

In part (a), nearly all the students answered this question correctly by calculating the 95% confidence interval. A handful of students made some arithmetic errors in their calculations.

In part (b), nearly all students mentioned that the r-square value was low. Many also mentioned that the explanatory variables themselves were not significant as could be observed by the confidence intervals. The typical wrong answer omitted the latter factor and only focused on r-square.

Part (c) could have a variety of correct answers. Those who did well (a vast majority) mentioned at least two correct factors such as the lack of quality data, the need for more accurate information, the dangers of extrapolation, multicollinearity and noted how the problems identified could be addressed. Those who did not do well (a minority) only listed one factor or did not mention solutions to the problems they identified.

In part (d) many of the students calculated the monthly orders wrongly. However, their logic in approaching the problem was typically correct and they got partial credit for that.

In part (e), many students calculated the expected profit wrongly. Those who they followed the correct logic in approaching the question got partial credit for that.

In part (f), the question asked for both a quantitative solution and a qualitative explanation. Some of the students only provided one or the other and thus lost marks for the aspects of the question they omitted. Some of the students did not attempt the question. This could be because the question was at the end and they ran out of time.