

Excerpts from a typical dialogue during an oral exam in 732A66 Decision theory

Note! You should not try to learn the answers below by rote. The questions put will vary a lot in their formulations. The dialogue just gives you the structure and the scope of what you are supposed to be able to discuss and explain.

Examiner: Which are the components of a statistical model for decision making?

Student 1: There are different states of the world, a set of possible actions and for each combination of state of the world and action there is a known consequence.

Examiner: What do we mean by consequence? Can you give examples?

Student 2: Say that the actions are about investments. Then, the consequences can be about net gains of this investment, whether we win or lose or remain unaffected by the action taken.

Examiner: How does the model become “statistical”?

Student(s): It depends on whether we have non-probabilistic or probabilistic criteria in choosing actions. With non-probabilistic criteria we use a payoff or loss function (table), which may have randomness in them. In that case we would expect payoff/expected loss. With probabilistic criteria we use a probability distribution over the states of the world (prior probabilities if no data is collected for the decision and posterior probabilities if data is collected).

Examiner: Say that you are seated for a coach trip to Stockholm (from Linköping). Would you use the seat belt?

Student(s): If it is compulsory to use it, I would.

Examiner: ...but if it is not compulsory, but recommended?

Student(s): R1: I think I would... R2: I'm not sure I would, depends maybe where I sit...

Examiner: Could you elaborate on that...how should we describe this decision problem?

Student(s): *will write things on the whiteboard, discuss openly, there may be issues that makes the examiner to interfere, it is probable that non-probabilistic criteria will be discussed here.*

...

Examiner: Is there a difference between payoff and utility? If so, in what way?

Student(s): If we measure the consequences in terms of money, a certain amount of money may have different importance between different individuals. Some are happy with small amounts of payoff while others require more. Hence, the utility of say 1000 SEK can be different for two persons.

Examiner: Could we say that payoff is non-personal while utility is personal?

Student(s): Yes!

Examiner: That must imply difficulties in obtaining utility functions then? How can we do that?

Student(s): It is difficult, but a logical strategy is to find out the worst consequence and the best consequence. Then the utility function is normalized so that the worst consequence gets utility 0 and the best utility 1. Then for a certain consequence the utility of that consequence is equal to the probability p that makes you indifferent between obtaining that consequence for certain and obtaining the worst consequence with probability $1 - p$ and the best with probability p .

Examiner: What are your own utilities when it comes to the decision of using a seat belt or not?

Student(s): *will elaborate, but probably not reasoning in terms of normalized utilities*

Examiner: Your utility function can characterize you as an actor in this decision problem. Which kind of actors are there?

Student(s): risk avoider (risk adverse), risk neutral and risk taker

Examiner: Describe these and try to tell which one you are in this problem

...

Examiner: We have talked about the value of information. What do we mean by that?

Student(s): *give an account of value of perfect information and (perhaps upon subsequent questions from the examiner) value of sample information, EVPI, EVSI, ENGSI, single-stage and sequential plans*

Examiner: A lot of problems discussed during the course has been with a few states of the world and few actions to choose between. What kind of loss functions (or payoff functions) are used in these problems?

Student(s): These are 0-k functions but with different k:s for different combinations

Examiner: Could you give examples where the loss function is not a 0-k function?

Student(s): We can do point estimation with a quadratic loss function. Then the action is the point estimator that minimizes posterior loss (which is the posterior mean).

Examiner: Is the state of the world always a finite set of possible states?

Student(s): No, we can model the state of the world with a continuous probability distribution. This may be the case when the state of the world is a quantity of interest, e.g. the demand of something that is measured in kilograms, litres, etc.

Examiner: Could you give arguments to why in this case a Bayesian approach to decision making is beneficial compared to a frequentist approach?

Student(s): *there may be several arguments, but Bayesian inference takes care of the uncertainty about the state of the world by modeling it. If it is not modeled but it is evident that it is a continuously varying quantity it can be very hard to use non-probabilistic criteria*

...

Examiner: We have taken up Bayesian networks in the course. Could you elaborate on what they are?

Student(s): *give an account of Bayesian networks*

Examiner: Could we model the entire decision problem graphically?

Student(s): Yes, by using an influence diagram.

Examiner: And what is that?

...

Examiner: *Turning to Student 1:* Assume you are the examiner of this course. What question would you ask your colleague here (about things we haven't discussed so far)?

Student 1: *asks Student 2 a question*

Student 2: *(tries to) answer(s) the question*

Examiner: *Turning again to Student 1:* Are you happy with that answer? If, not how would you like to correct?

Student 1: ...

...and this is repeated switching order between student asking and student answering