Theory of Statistics

(I) Introduction

- what is the purpose of a theory of statistics?

- there is an object ψ whose value is unknown.

- eg ψ = half-life of a neutron in seconds

  ψ = price of a stock at a particular future time.

  ψ = party of U.S. president after 2020 election.

- there is data x which presumably contains evidence concerning the following problems

  (1) specify an estimate $\psi(x)$ together with an assessment of how accurate it is.

  (2) indicate whether there is evidence in favor of or against a specified value ψ_0 for ψ and also how strong/weak this evidence is

- (1) is the <u>estimation problem</u> and (2) the <u>hypothesis assessment problem</u>

- so a theory of statistical reasoning indicates how we go from x to answer (1) and/or (2)

- is data enough?

- it seems that we need additional ingredients beyond the data to which the rules of statistical inference are applied to give answers

- like a problem of logical reasoning which is comprised of the premises and the rules of inference.

- ingredients

- this varies but common to all is the notion of a statistical model $\{f_\theta : \theta \in \Theta\}$ = a collection of probability distributions on a set $\mathcal{X}$ one of which it is assumed produced $x \in \mathcal{X}$

- here $\theta \in \Theta$ indexes the possible distributions and so there is a true value $\theta_{true} \in \Theta$ corresponding to the distribution which produced x.

- also α is such that $\alpha = \overline{\alpha}(\Theta) \in \mathcal{A}$ (overload notation)
  so that $\alpha_{true} = \overline{\alpha}(\Theta_{true})$

- sometimes a prior prob. dist π on Θ is added where π reflects beliefs about $\Theta_{true} \in \Theta$

- additionally it is common to add a loss function $L : \Theta \times \mathcal{A} \to [0,\infty)$ where $L(\Theta, \alpha) = 0$ iff $\alpha = \overline{\alpha}(\Theta)$

- other ingredients?

- what characteristics should the ingredients satisfy?

- note - in general all ingredients are subjective in the sense that they are not dictated by the application.

- is that bad/good?

- good as their judgement can be employed.

- bad because the choices made could be badly wrong or even biased in the sense that inferences are pre-determined.

- for scientific contexts we require each ingredient must be falsifiable, namely, is the ingredient in question somehow contradicted by the (objective) data

- so far the model there is model checking: is x surprising for each f_θ?

- for the prior there is checking for prior-data conflict: given the model is there an indication that θ_{true} lies in the tails of π?

- there are consistent methods for model checking and checking for prior-data conflict.

- what about checking the loss L? - I don't know how

- even if the model and prior pass their checks it is possible that biased choices have been made.

- with methodology for checking the ingredients and assessing and controlling bias, criticisms concerning subjectivity are largely voided.

- caution - don't obsess about the ingredients as they are always wrong

  - the real question is whether or not the inferences drawn for estimation or hypothesis assessment are seriously distorted because of the choices made.

  - the ingredients are just devices so that we can reason concerning the questions of interest

- <u>the rules of statistical inference</u>

- what are they?

- they are independent of the correctness/incorrectness of the ingredients, i.e., for inference we assume the ingredients are correct

- the rules should be based on as clear a specification of statistical evidence as possible and thus place some requirements on what the ingredients need to be

- hint: a model and a prior suffice

- theory of statistical reasoning for an ideal problem

ingredients

- specify model, prior, loss, ...
- measure bias
- control bias by design
- collect data
- check ingredients and modify if necessary

rules of statistical inference

- apply rules to get answers to estimation and hypothesis assessment problems.