Eliciting Probability Distributions with `shelf`

Jeremy Oakley

School of Mathematics and Statistics

6/9/16

The University Of Sheffield.

Background

Aim: to represent an expert's subjective uncertainty about a fixed quantity using a probability distribution

Background

Aim: to represent an expert's subjective uncertainty about a fixed quantity using a probability distribution

Toy example

Define X: shortest distance (in miles) by road from Birmingham to Newcastle.

Background

Aim: to represent an expert's subjective uncertainty about a fixed quantity using a probability distribution

Toy example

Define X: shortest distance (in miles) by road from Birmingham to Newcastle.

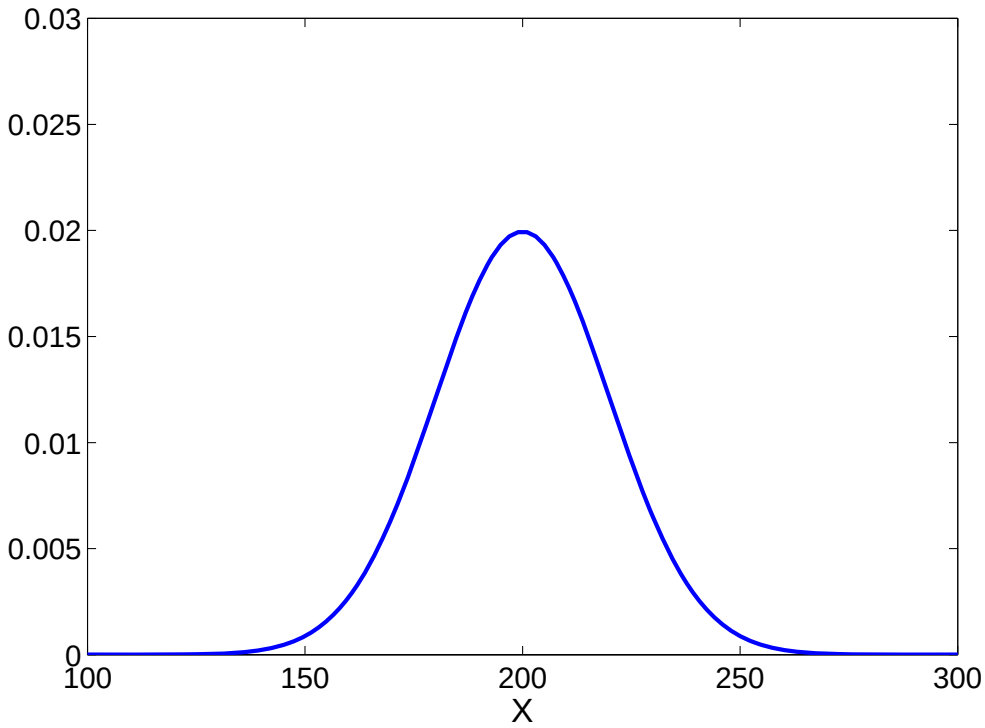

Applications

Applications

- Decision-making with little/no data; quantifying uncertainty about decision-model inputs

Applications

- Decision-making with little/no data; quantifying uncertainty about decision-model inputs
- Experimental design; clinical trial planning

Applications

- Decision-making with little/no data; quantifying uncertainty about decision-model inputs
- Experimental design; clinical trial planning
- Bayesian inference (but very rarely!)

The task

The task

- Experts usually won't provide distributions directly

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:
  - Direct probability: "I think $P(X < 150)$ is..."

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:
  - Direct probability: "I think $P(X < 150)$ is..."
  - Quantile: "My median value is 200; I think $P(X < 200) = 0.5$."

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:
  - Direct probability: "I think $P(X < 150)$ is..."
  - Quantile: "My median value is 200; I think $P(X < 200) = 0.5$."
- Analyst ("Facilitator") fits a probability distribution to these judgements and provides feedback (additional quantiles)

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:
  - Direct probability: "I think $P(X < 150)$ is..."
  - Quantile: "My median value is 200; I think $P(X < 200) = 0.5$."
- Analyst ("Facilitator") fits a probability distribution to these judgements and provides feedback (additional quantiles)
  - "Based on the fit, we're assuming you are 99% certain $X < 250$."

The task

- Experts usually won't provide distributions directly
- Expert(s) instead make a small number of probability judgements:
  - Direct probability: "I think $P(X < 150)$ is..."
  - Quantile: "My median value is 200; I think $P(X < 200) = 0.5$."
- Analyst ("Facilitator") fits a probability distribution to these judgements and provides feedback (additional quantiles)
  - "Based on the fit, we're assuming you are 99% certain $X < 250$."
- Software needed for fitting a range of distributions, displaying results, providing feedback.

Illustration

- Expert has provided a median and tertiles (33rd and 66th percentiles)

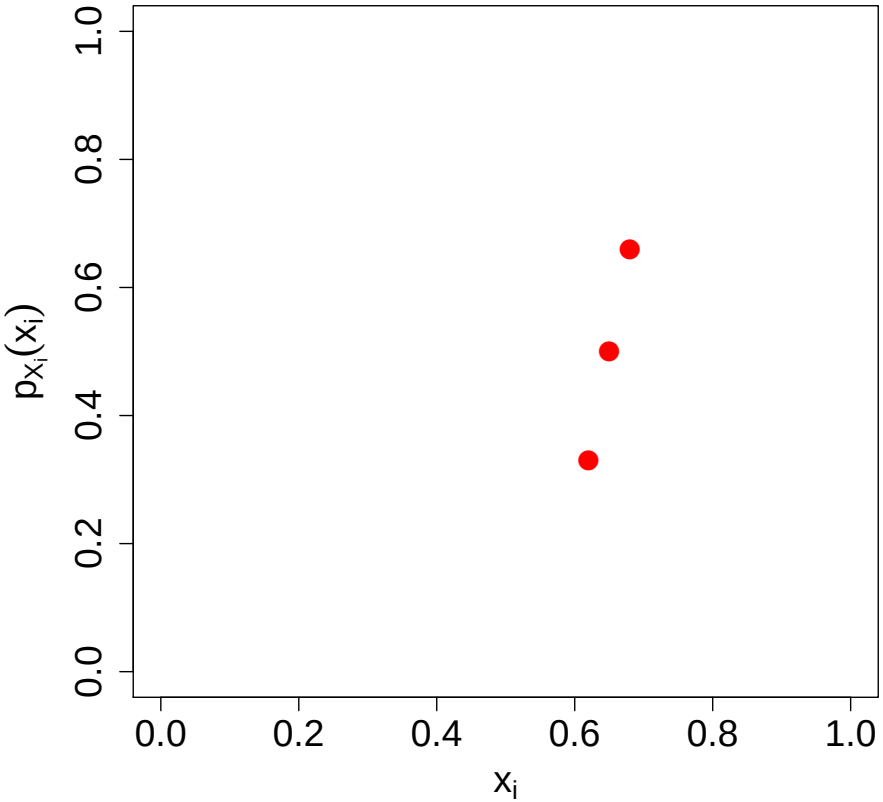

Illustration

- Expert has provided a median and tertiles (33rd and 66th percentiles)
- Fit a distribution to these points on the CDF (the Beta(30.6, 16.6) distribution here)

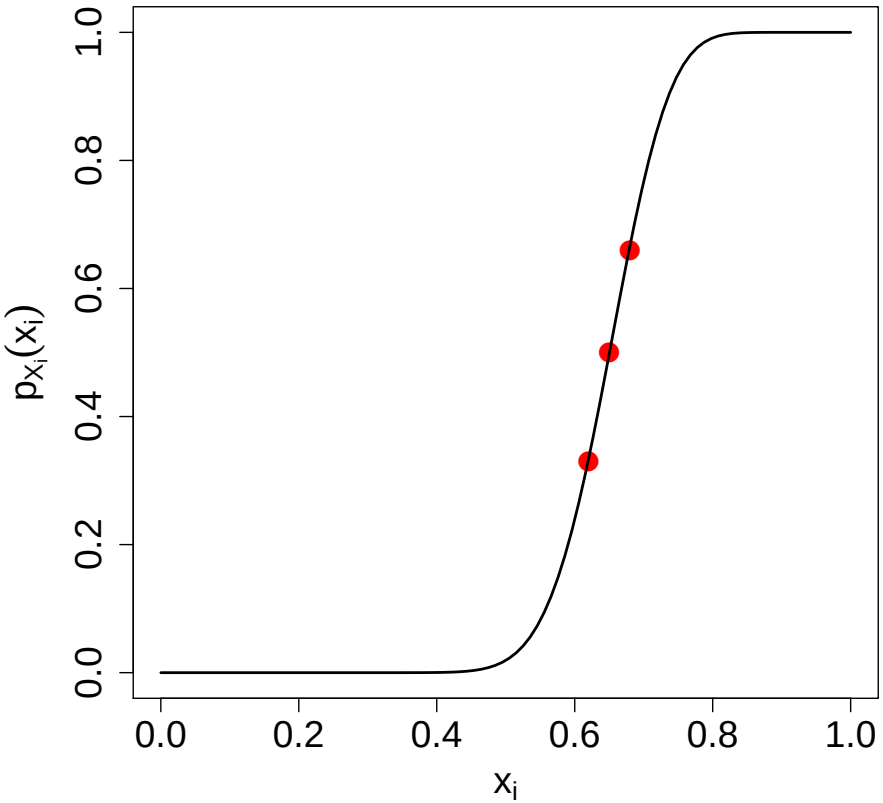

Illustration

- Expert has provided a median and tertiles (33rd and 66th percentiles)
- Fit a distribution to these points on the CDF (the Beta(30.6, 16.6) distribution here)
- Consider further percentiles from the fitted Beta(30.6, 16.6) distribution (eg 5th and 95th). Check these seem appropriate.

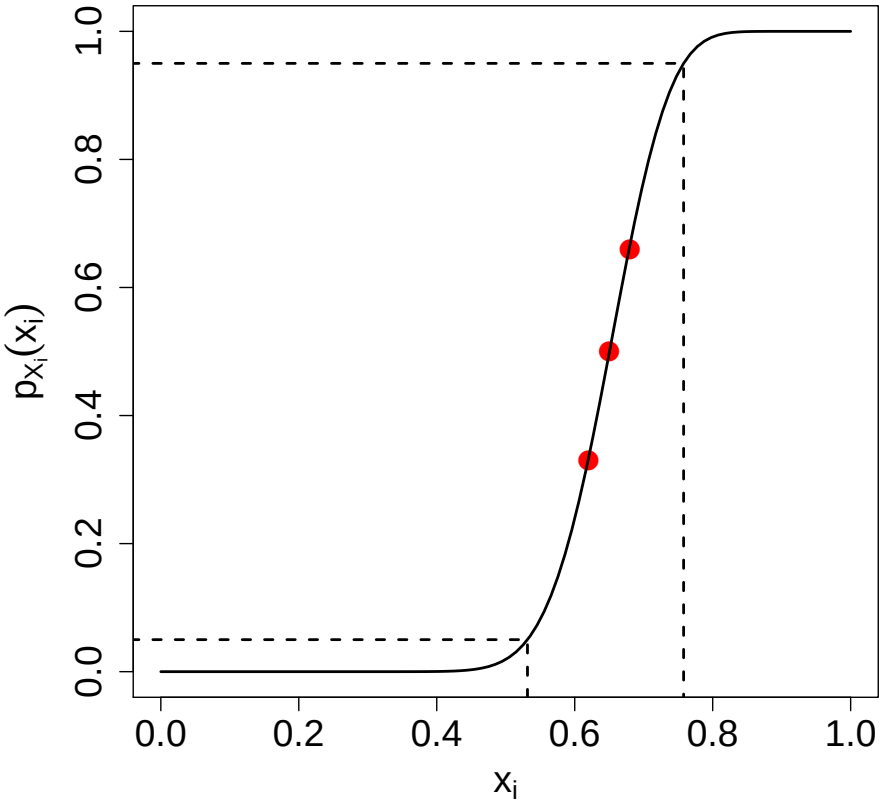

SHELF and shelf

SHELF and `shelf`

- SHELF: the SHeffield Elicitation Framework, developed by Tony O'Hagan and myself
  - Protocols and templates for eliciting a single distribution from a group of experts

SHELF and `shelf`

- SHELF: the SHeffield Elicitation Framework, developed by Tony O'Hagan and myself
  - Protocols and templates for eliciting a single distribution from a group of experts
- `shelf`: an R package for implementing the methods

SHELF and `shelf`

- SHELF: the SHeffield Elicitation Framework, developed by Tony O'Hagan and myself
  - Protocols and templates for eliciting a single distribution from a group of experts
- `shelf`: an R package for implementing the methods
- History:
  - Initially just functions for download

SHELF and `shelf`

- SHELF: the SHeffield Elicitation Framework, developed by Tony O'Hagan and myself
  - Protocols and templates for eliciting a single distribution from a group of experts
- `shelf`: an R package for implementing the methods
- History:
  - Initially just functions for download
  - MATCH tool front-end, developed by Ed Morris at Nottingham

SHELF and `shelf`

- SHELF: the SHeffield Elicitation Framework, developed by Tony O'Hagan and myself
  - Protocols and templates for eliciting a single distribution from a group of experts
- `shelf`: an R package for implementing the methods
- History:
  - Initially just functions for download
  - MATCH tool front-end, developed by Ed Morris at Nottingham
  - Now a package on CRAN (and github). Used tutorial on "R Function of the Day", and Wickham (2015)