R functions for fitting distributions
Instructions

0) Quick summary

0.1) Install the *rpanel* library, then `source` the `shelf2.R` function.

0.2) For eliciting from a single expert, use one of the commands

`quartile.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

`hybrid.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

`roulette.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

`tertile.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

`probability.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

Add the argument `sig.fig=` *no. of significant figures*

if you want to change the number of significant figures used in displays.

0.3) For eliciting from a group, use the command

`elicit.group.values(N.experts=` *no. of experts* `Lo=`*lower limit*`,Up=`*upper limit*`,method =` *elicitation method*`)`

with *elicitation method*

elicitation method	Individual judgements	Group elicitation	SHELF template
`“q”`	quartile	quartile	Q
`“r”`	roulette	roulette	R
`“t”`	tertile	tertile	T
`“qp”`	quartile	probability	QP
`“rp”`	roulette	probability	RP
`“tp”`	tertile	probability	TP

Add the argument `sig.fig` as before, if desired. Enter the appropriate values for each group member in the data window, then quit the data window. Quit the 1st control panel to move onto eliciting a consensus distribution.

If you are re-running the `elicit.group.values` command in a single session, and wish to use your last defined group data, add the argument `revise=T`

0.4) Output from the functions is stored in the global variables `elicited.consensus.data` and `elicited.group.data`.

1) Obtaining the R package

R is a free software environment for statistical computing and graphics, and can be downloaded from the R homepage at http://www.r-project.org

2) Installing rpanel

These functions use the rpanel library (Bowman and Crawford, 2008). To install the rpanel library start R, and go to Packages→Install package(s)…, then select a CRAN mirror. Then select rpanel and click OK.

3) Running the SHELF function

1. Save the function `shelf2.R` into a folder on your hard drive
2. Start R, and go to *File→Change dir…*, then select your chosen folder from step 1
3. Type `source("shelf2.R")` in the commands window.

Note that the various control panels used in these routines will appear in separate windows. The graphics will appear within the R window, and so you will need to re-size the R window so that both the graphics and control panels are visible.

4) Eliciting distributions from a single expert

There are five available methods in SHELF. Results from each function are stored in the variable `elicited.consensus.data`, which you can view in the commands window. Go to *File→Save workspace…* if you wish to save the data.

4.1) The quartile method

1. Start the quartile method elicitation by typing

   `quartile.single(Lo=` *lower limit*`,Up=` *upper limit*`)`

   in the commands window. The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the expert must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.
2. Elicit the median value using the slider on the control panel
3. Check the “Elicit lower quartile box” and elicit the lower quartile, using the slider on the control panel.
4. Check the “Elicit upper quartile box” and elicit the upper quartile, using the slider on the control panel.

5. Change the feedback to “Four intervals” in the “Feedback” box. The expert should judge the four coloured sections to equally probable, given the specification of the quartiles. The expert may wish to adjust the quartiles before proceeding.
6. Change the feedback to “Density” in the “Feedback” box. This will show the fitted density function, together with the parameters of the distribution, for the distribution selected in the “Distribution” box. “Best fitting” will show the distribution that most closely fits the expert’s quartiles.
7. Quantiles of the fitted distribution will be reported in the density plot window, and the choice of quantiles can be changed using the sliders in the control panel.
8. The elicited quartiles can be changed at any time, with the distributions re-fitted automatically.
9. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.
10. If the interval (L,U) is set at anything other than (0,1), the distribution option “Beta” will change to “Scaled Beta”. This fits a Beta distribution to the transformation (X – L) / (U – L)

4.2) The hybrid method

1. Start the hybrid method by typing

   `hybrid.single(Lo=` *lower limit*`,Up=` *upper limit*`)`

   in the commands window. The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the expert must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.
2. Elicit the median value M using the slider on the control panel
3. Check the “Elicit probability (a)” box and elicit the probability for the interval specified in the graphics window, using the slider on the control panel. (The interval is calculated as [L,(2M+L)/3]).
4. Check the “Elicit probability (b)” box and elicit the probability for the interval specified in the graphics window, using the slider on the control panel. (The interval is calculated as [(2M+U)/3,U]).
5. Change the feedback to “Show histogram” in the “Feedback” box. This shows a density which is piecewise uniform across the intervals

[L,(2M+L)/3], [(2M+L)/3,M], [M,(2M+U)/3], [(2M+U/3],U]. The expert may wish to adjust the median or probabilities before proceeding.

6. Change the feedback to “Density” in the “Feedback” box. This will show the fitted density function, together with the parameters of the distribution, for the distribution selected in the “Distribution” box. “Best fitting” will show the distribution that most closely fits the expert’s probabilities.
7. Quantiles of the fitted distribution will be reported in the density plot window, and the choice of quantiles can be changed using the sliders in the control panel.
8. The elicited median and probabilities can be changed at any time, with the distributions re-fitted automatically.
9. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.
10. If the interval (L,U) is set at anything other than (0,1), the distribution option “Beta” will change to “Scaled Beta”. This fits a Beta distribution to the transformation (X – L) / (U – L)

4.3) The roulette method

The roulette method is a fixed interval method, in which the expert allocates chips to a number of bins, with the probability assigned to a particular bin given by the proportion of chips allocated to that bin.

1. Start the trial roulette method by typing

   `roulette.single(Lo=` *lower limit*`,Up=` *upper limit*`)`

   in the commands window. The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the expert must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.

   Ten equally sized bins will be drawn between L and U. The argument `nchips` is only used for plotting purposes, setting the maximum height on the y-axis. You can use any number of chips in total when making the probability judgements

2. Place chips in the marked bins by clicking on the buttons in the control panel.
3. When finished, check the “Show fitted distribution box”. This will show the fitted density function, together with the parameters of the distribution, for the distribution selected in the “Distribution” box. “Best fitting” will show the distribution that most closely fits the expert’s probabilities.

4. Quantiles of the fitted distribution will be reported in the density plot window, and the choice of quantiles can be changed using the sliders in the control panel.
5. The allocation of chips to bins can be changed at any time, with the distributions re-fitted automatically.
6. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.
7. If the interval (L,U) is set at anything other than (0,1), the distribution option "Beta" will change to "Scaled Beta". This fits a Beta distribution to the transformation (X – L) / (U – L)

4.4) The tertile method

1. Start the tertile method by typing

   `tertile.single(Lo=` *lower limit*`,Up=` *upper limit*`)`

   in the commands window. The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the expert must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.
2. Elicit the median value using the slider on the control panel
3. Elicit the 1st tertile t_1, so that $P(X < t_1) = 1/3$, using the slider on the control panel.
4. Elicit the 2nd tertile t_2, so that $P(X < t_2) = 2/3$, using the slider on the control panel. Note that in the display, the red, white and blue sections should all represent regions of equal probability
5. Change the feedback to "Show histogram" in the "Feedback" box. This shows a density which is piecewise uniform across the intervals $[0,t_1]$, $[t_1, t_2]$, $[t_2,U]$. The expert may wish to adjust the median or tertiles before proceeding.
6. Change the feedback to "Density" in the "Feedback" box. This will show the fitted density function, together with the parameters of the distribution, for the distribution selected in the "Distribution" box. "Best fitting" will show the distribution that most closely fits the expert's median and tertiles.
7. Quantiles of the fitted distribution will be reported in the density plot window, and the choice of quantiles can be changed using the sliders in the control panel.
8. The elicited median and tertiles can be changed at any time, with the distributions re-fitted automatically.

9. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.
10. If the interval (L,U) is set at anything other than (0,1), the distribution option “Beta” will change to “Scaled Beta”. This fits a Beta distribution to the transformation (X – L) / (U – L).

4.5) The probability method

1. Start the probability method by typing

   `probability.single(Lo=`*lower limit*`,Up=`*upper limit*`)`

   in the commands window. The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the expert must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.
2. Specify values of X_0, X_1 and X_2 in the boxes (you will need to press return after typing in each value).
3. Use the sliders to specify the three probabilities
   a. P_1: P(L<X<X_1)
   b. P_2: P(X_2<X<U)
   c. P_0: P(L<X<X_0)

   (We recommend that you consider the three probabilities in this order).
4. Change the feedback to “Show histogram” in the “Feedback” box. This shows a density which is piecewise uniform across the intervals [L,X_1], [X_1,X_0], [X_0,X_2], [X_2,U]. The histogram will be displayed in red if the three elicited probabilities are incoherent. The expert may wish to adjust the probabilities before proceeding.
5. Change the feedback to “Density” in the “Feedback” box. This will show the fitted density function, together with the parameters of the distribution, for the distribution selected in the “Distribution” box. “Best fitting” will show the distribution that most closely fits the expert’s probabilities.
6. Quantiles of the fitted distribution will be reported in the density plot window, and the choice of quantiles can be changed using the sliders in the control panel.
7. The elicited median and probabilities can be changed at any time, with the distributions re-fitted automatically.

8. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.
9. If the interval (L,U) is set at anything other than (0,1), the distribution option “Beta” will change to “Scaled Beta”. This fits a Beta distribution to the transformation (X – L) / (U – L)

5) Eliciting a distribution from a group of experts

1. Start the group tool with the command

`elicit.group.values(N.experts=` *no. of experts*`,` `Lo=`*lower limit*`,` `Up=`*upper limit*`,` `method =` *elicitation method*`)`

elicitation method can be any of the following

elicitation method	Individual judgements	Group elicitation	SHELF template
`"q"`	quartile	quartile	Q
`"r"`	roulette	roulette	R
`"t"`	tertile	tertile	T
`"qp"`	quartile	probability	QP
`"rp"`	roulette	probability	RP
`"tp"`	tertile	probability	TP

The arguments `Lo` and `Up` correspond to the lower and upper limits L and U, and so you should specify your own values for *lower limit* and *upper limit*. It will help visually to make the range (L,U) as small as possible, but the experts must be happy to give values outside this range zero probability; L and U cannot simply be extreme quantiles. The default number of significant figures used in the output is 2, but this can be changed by including the argument `sig.fig=` and specifying your own value.

no. of experts should be any value greater than 1.

If you are using the trial roulette method, ten equally sized bins will be drawn between L and U. You can include an optional argument

`Nchips =` *maximum number of chips* (default value 10).

This is used for plotting purposes only, setting the maximum height on the y-axis. You can use any number of chips in total when making the probability judgements

Specify the number of experts when prompted.

1. A data editor window will appear.

   For the quartile and tertile method, the first column (`expert.prob)` gives the cumulative probabilities. The remaining columns (`expert.value)` correspond to each expert, and you should enter the quartile or tertile values for each expert as appropriate.

   For the roulette method, the first two columns give the lower and upper limits of each bin. In the remaining columns you should specify the number of chips in each bin for each expert.

   Columns with headings `var3, var4` etc. can be ignored

   Close the window when finished.

2. A graphics window will appear, with a separate fitted distribution for each expert. You can change the type of fitted distribution fitted using the control panel. If the interval (L,U) is set at anything other than (0,1), the distribution option "Beta" will change to "Scaled Beta". This fits a Beta distribution to the transformation (X – L) / (U – L). The "best fitting option" finds the closest fitting distribution for each expert. This is reported in the plot legend.

3. If fitting a Student-t (or log Student-t) distribution, the R function will not automatically set the degrees of freedom. This can be chosen manually using the slider in the control panel.

4. Check the "Show linear pool" box to see an equal weighted linear pool of the expert's distribution, using the current option for "Distribution". (We recommend that you do not show the experts this before eliciting a consensus distribution).

5. You can modify the original judgements at any time by clicking the "change individual judgements" button, which re-opens the data window. Closing the data window does not automatically update the graph, so you will need to click on any other option within the control panel again once you have closed the data window.

6. Close the control panel to start eliciting a consensus distribution, using the same method. Follow the appropriate instructions in section 4.

7. If you have used the roulette method at the individual stage, and are using the probability method at the group stage, you can see the consensus allocation of chips to bins by selecting the option "Chips & bins" in the feedback box. This is calculated for an allocation of 20 chips by rounding the fitted probabilities in each bin.

8. In the control panel for the consensus distribution, there is the extra option of showing the linear pool. This is based on the "best fitting" option for each expert.

9. If you wish to report feedback from the linear pool, type the command

   ```
   linear.pool.feedback()
   ```

in the commands window. This uses the linear pool constructed from the best fitting distribution for each expert. You can specify your own set of quantiles $q_1, q_2, \ldots, q_n$ for feedback by adding the argument

`quantiles=c(`$q_1, q_2, \ldots, q_n$`)`

and you can see the feedback using each type of distribution by adding the argument

```
output="all"
```

10. Results from group elicitation are stored in the variable `elicited.group.data`, which you can view in the commands window. Go to *File→Save workspace…* if you wish to save the data. If you want to restart the group tool without re-entering all the individual data, add the argument `revise=T` to the `elicit.group.values` command.

5.1) The probability method versus the hybrid method

These two methods are fairly similar, in that both cases experts are asked to consider probabilities of the uncertain quantities lying in specified intervals. In the probability method, the facilitator chooses in the intervals, whereas the intervals are pre-determined in the hybrid method. Choosing suitable intervals can be hard, without first having some idea of what the experts believe, hence the recommendation to use a different method at the individual stage. Once individual beliefs have been reported, we think the probability method is superior to the hybrid method.

The hybrid method is available for those who wish to experiment with it, but it is not recommended as a formal SHELF procedure for group elicitation.

Reference

Bowman, A. W. and Crawford, E. (2008). R package 'rpanel': simple

control panels (version 1.0-5) URL http://www.stats.gla.ac.uk/~adrian/rpanel