

ELICITATION RECORD – Part 2 – Distribution

Tertile and Probability Method

Elicitation title	As in the Part 1 form
Session	As in the Part 1 form
Date	As in the Part 1 form
Quantity	The uncertain quantity whose distribution is to be elicited
Start time	Time when this part of the elicitation started

Definition	Repeat the definition of this quantity from Part 1. Give it a symbol to facilitate the recording of judgements about it. It will be called X in these notes.
Evidence	<p>Review the evidence specifically about X. (Refer to principal sources, but do not repeat lots of detail here.)</p> <p>[Although the evidence base has been set out in the Part 1 form, the facilitator should ask each expert to consider which items are of relevance to this quantity.]</p> <p><i>As in Part 1, this step is to avoid the ‘availability heuristic’, in which experts rely only on a subset of evidence that comes readily to mind.</i></p>
Plausible range	<p>Elicit from the experts the range of plausible values for X. In these notes, L is the lower bound of the range and U the upper bound. It may be useful to record absolute, logical bounds, but the objective here is to identify a range such that it is extremely unlikely (but not necessarily impossible) that X lies outside.</p> <p>Record how the range has been informed by the evidence base.</p> <p>[This range should not be unnecessarily wide, but it is important that it should not be too narrow. This should be a joint judgement of the experts, such that they all believe that X is extremely unlikely to be outside (U, L). After an initial specification of the range, the facilitator should probe by asking something like “Suppose an experiment produced a value [something below L or above U] for X; would this have to have been a flawed experiment, or might there be a way that X could have this value?”]</p> <p><i>There is substantial evidence that experts tend to be over-confident, in the sense that they do not allow enough probability for extreme values of X. Psychologists suggest that experts develop mental models that allow them to understand and predict the processes in their area of expertise. Experts themselves may recognise instances where the model does not</i></p>

	<i>apply, but part of their overconfidence in elicitation may be due to conditioning on inappropriate models. The suggested probing question invites them to think outside their models.</i>
Median	<p>The next steps should be done by each expert separately, without discussion. Each expert should specify their median value for X. This is a value such that they think 'X lies below the median' and 'X lies above the median' are equally likely propositions. Formally, if M is the median, then $P(X < M) = 0.5$.</p> <p>[The facilitator should instruct the experts to write down their own median values, but not to reveal them yet. Nothing should be written in this field until after the upper and lower tertiles have also been elicited.]</p> <p><i>The judgement of equal probability is generally found to be simple for experts, and is not subject to systematic biases.</i></p>
Upper and lower tertiles	<p>Each expert should now specify their upper and lower tertiles by considering the range from L to U and dividing it into three equally likely intervals. Formally, if T1 is the lower tertile and T2 is the upper tertile, then $P(L < X < T1) = P(T1 < X < T2) = P(T2 < X < U) = 0.33$.</p> <p>Before deciding definitely on these values, experts should be asked to check not only that they regard each of the three ranges (L to T1, T1 to T2 and T2 to U) as equally likely, but that the ranges T1 to M and M to T2 are also equally likely..</p> <p>[When asking for the tertiles, the facilitator should point out that generally experts would feel that values of X close to M are more probable than values close to L or U, and so the interval between T1 and T2 will typically be narrower than that between L and T1 or between T2 and U. In choosing T1 and T2, the experts are determining how much more likely values near M are. Obviously, M should lie between T1 and T2.</p> <p>When all the experts have written down their medians and tertiles, the facilitator should ask them to reveal their values, and should fill in this box and the one above in the record.]</p> <p><i>Psychologists suggest that people often make judgements by adjusting a previous judgement, and that when they do this they typically do not adjust far enough. The first judgement is called the anchor, and the process is called the 'anchoring and adjustment' heuristic. Anchoring and adjustment is suggested as one reason for over-confidence in the assessment of probability intervals, when experts are anchored by a previous central estimate (such as the median), and do not adjust far enough away from this. The SHELF protocols always elicit first the plausible range, and then ask experts to make such judgements relative to both the central value and the plausible range. In this way, they are anchored on both sides, which</i></p>

	<p><i>tends to cancel the effect.</i></p> <p><i>However, another common cause of poor judgement is the 'range-frequency compromise,' according to which when thinking about probabilities within a range people tend to want to share probability reasonably evenly across the range. The effect of this is that the two anchors will tend to cause experts to place their tertiles so as to make the three intervals more or less equally wide. This is why the facilitator should point out the essential unevenness of probability, and suggest a value closer to M. This needs to be done carefully, so as not to influence the experts too much.</i></p> <p><i>Although quartiles are generally somewhat easier for the experts to think about, there is some evidence that tertiles are elicited more accurately.</i></p>
Fitting	<p>The facilitator fits a distribution to each of the experts' assessments.</p> <p>(The distributions should be specified here, and if possible shown as density functions. If it is not straightforward to paste plots of the density functions into this record, they can be provided as an attachment, which should then be listed at the end.)</p> <p>[The facilitator should choose an appropriate family of distributions, and then fit the distribution by choosing parameters that give probabilities matching the elicited bounds, median and tertiles as closely as possible. Since L and U are not necessarily absolute bounds, there are 6 probabilities to match – two of 0.33, two of 0.17 and two of 0.0.</p> <p>The distributions should be shown to the experts, but at this stage we do not invite revision (unless the expert is insistent that the plotted distribution badly distorts his/her beliefs) or provide any other feedback.</p> <p>If feasible, the facilitator should compute the median and tertiles of an equally-weighted average of the density functions. Ideally, these should not be revealed to the experts immediately, but may be used at the facilitator's discretion in the next stage.]</p> <p><i>This stage of separate elicitations ensures that the initial divergence of opinion between the experts is recorded. The facilitator can refer to these if the group elicitation appears to be neglecting part of the original range of belief. There is evidence that group elicitation can itself lead to over-confidence, perhaps because the process of reaching consensus induces a false sense of decisiveness. So this step in the SHELF process allows the facilitator to see any narrowing of uncertainty, and to check that this is justified by the sharing of knowledge that has taken place. In particular, having the median and tertiles of the averaged densities gives the</i></p>

	<p><i>facilitator a quick check on the degree to which subsequent discussion leads to a narrowing of uncertainty.</i></p> <p><i>The process of averaging the density functions is known as the linear opinion pool (with equal weights). It is one of the formulae which proponents of eliciting separately from experts use to combine the resulting distributions. We use it in SHELF simply as a benchmark.</i></p>
Group elicitation	<p>After discussion of the different distributions, and sharing of knowledge and reasoning about the differences, record group consensus values for probabilities $P1 = P(L < X < X1)$ and $P2 = P(X2 < X < U)$, and finally for $P0 = P(L < X < X0)$, where $X1$, $X2$ and $X0$ are selected by the facilitator.</p> <p>[This discussion can take a substantial amount of time. The facilitator should only cut it short if he/she feels that the experts are no longer exchanging information and arguments but are just repeating opinions. The agreed median will inevitably be some sort of compromise. Before discussion, there are two components of uncertainty in the group – uncertainty that each expert has and is expressed in that expert's quartiles, as well as variability between the experts' judgements. The agreed probabilities should reflect the group's overall uncertainty that remains after the discussion.</p> <p>The facilitator needs to select three values, $X1$, $X0$ and $X2$ (with $X1 < X0 < X2$), for the group elicitation. The choice of $X1$ and $X2$ is again at the facilitator's discretion, with the aim of eliciting group $P1$ and $P2$ values that are not too close to zero or 0.5. $X0$ is selected between these values</p> <p>Tertiles of the equally-weighted opinion pool distribution may work well for this choice, assuming they have not previously been shown to the experts. $X0$ might be set to the median of the equally-weighted opinion pool distribution, or to the average of $X1$ and $X2$. As before, the chosen values could be rounded.</p> <p>The facilitator might need to adjust the above recommendations in the light of the group discussion. If, for instance, the group seems to be discounting the opinion of one expert whose original elicitation favoured higher values of X than the other experts, then the facilitator might expect to see the consensus distribution concentrating on lower values than indicated by the equally-weighted opinion pool, and so might choose lower values of $X1$, $X2$ and $X0$.</p> <p>The discussion and sharing of opinions will often mean that $P1$ and $P2$ are less than the corresponding probabilities derived from the equally-weighted opinion pool distribution. $P1$ and $P2$ should be elicited before $P0$.</p> <p>The facilitator needs to manage the discussion so that divergent views are properly considered, and to ensure that strong</p>

	<p>personalities and/or groups of people with overlapping experiences do not inappropriately dominate. In general, we would expect the group debate and sharing of opinions to result in a distribution that is narrower than the linear opinion pool, but it should not be much narrower or markedly different in other ways, unless the discussion has clearly altered individual experts' opinions.</p> <p>In some situations, experts may not be able to reach consensus. It may then be necessary for the facilitator to elicit two (or more) distributions, representing consensus views within the opposing factions.]</p> <p><i>The individual elicitations and subsequent discussion provide the facilitator with enough information to choose meaningful probabilities to request, so that it is not necessary to make use of quantiles. Psychology research has shown that in general people assess probabilities more accurately than quantiles.</i></p> <p><i>P1 and P2 are elicited before P0 to minimise the effect of the anchoring and adjustment heuristic – P0 is effectively anchored on both sides.</i></p>
Fitting and feedback	<p>Record here the (potentially iterative) process of fitting, feedback and revision of the group judgements.</p> <p>[The facilitator first fits a distribution to the group's probability values. This should be shown to the experts, and the fitted probabilities compared with the elicited probabilities. The experts are invited to consider whether the fit is close enough, or whether some values might be varied in order to fit others (that are believed to be more pivotal) better. The facilitator then feeds back to the experts some implied probabilities in the fitted distribution, such as the 10th and 90th percentiles, or the median and tertiles. The experts are invited to consider whether these are reasonable reflections of the group's knowledge.</p> <p>If revision is needed, this may be followed by further rounds of fitting and feedback until the experts are comfortable with the fitted distribution and its implications.]</p>
Chosen distribution	<p>Record and show here the finally fitted distribution.</p>
Discussion	<p>The facilitator should record here any difficulties that arose during the elicitation of this distribution. Also the experts' reactions to the process and to the finally fitted distribution.</p> <p><i>The elicitation record should be open about any concerns with the finished distribution. The SHELF protocols are designed to avoid many of the pitfalls of elicitation, but no process is perfect. It is important to be critical and realistic about the result. Nevertheless, it is important also to remember that, despite whatever deficiencies it might have, the elicited distribution is our best attempt. It has been developed using a</i></p>

	<i>robust protocol, and since expert knowledge is needed in the wider enterprise there is no alternative!</i>
--	---

End time	Time when elicitation of this distribution was completed.
Attachments	List any attachments, e.g. plots of distributions.