MedCrave Step into the World of Research
Biometrics & Biostatistics International Journal

Review Article

Open Access CrossMark

Facilitated prior elicitation with the wolfram CDF

Abstract

One of the key advantages of the Bayesian paradigm is the ability to incorporate past experiences and expert opinion into statistical analyses. However, the principled, precise distillation of expert opinion into a probability distribution, a task known as prior elicitation, is challenging and involves considerations from psychology, computational science, software engineering, and other related fields. Moreover, for elicitation to be practical, applied statisticians need good computer tools to enable its use–to perform complex calculations and provide feedback. Such computer programs are called facilitators. Using the prior elicitation of a population proportion as a canonical example, in this article we contend that Wolfram Research Inc.'s Mathematica/CDF Player/CDF suite of technologies provides a new state-of-the-art platform for the development and dissemination of free facilitators. To illustrate its capabilities, we present a free facilitator that enables, in a real-time interactive environment, the computations and diagnostics required of an elicitation method known as the mode/percentile (MP) method.

Keywords: bayesian statistics, prior elicitation, interactive graphics, wolfram CDF

Volume 3 Issue 6 - 2016

David Kahle,[1] James Stamey,[1] Fanni Natanegara,[2] Karen Price,[2] Baoguang Han[3]
[1]Department of Statistical Science, Baylor University, USA
[2]Eli Lilly & Company, USA
[3]Biogen Inc., USA

Correspondence: David Kahle, Department of Statistical Science, Baylor University, Waco, TX, USA, Tel 254-710-6102, Email david_kahle@baylor.edu

Received: June 01, 2016 | **Published:** June 07, 2016

Introduction and background

One of the great strengths of the Bayesian paradigm is the ability to formally incorporate prior belief–typically expert opinion–into statistical analyses.[1] To do this, however, requires the representation of belief as a single probability distribution. The conversion of such belief, especially expert opinion, into a probability distribution is a task known as prior elicitation.

There have been several investigations into prior elicitation. Perhaps the most up-to-date concise resource is the definitive survey article by Garthwaite et al.,[2] which features nearly 150 references from the statistics and psychology literature.[2] Additionally, there are no less than three full-length texts dedicated to the subject.[3-5] Each of these texts explores the interaction between the statistician and the non-statistician and contains topics ranging from how questions should be asked to what questions should be asked to how to obtain good diagnostics to judge elicitation fidelity. The first text is particularly related to our investigation in this article, as it takes more time with some of the more technical statistical details.

While the body of prior elicitation works is fairly large, the subject is by no means considered settled, partly because of its wide-ranging complexity. Prior elicitation is by its very nature a multi-faceted problem, involving considerations from psychology, computational science, optimization theory, software engineering and others in a concerted effort to tackle a fundamentally statistical problem. For instance, much (but by no means all) of the psychological work done on the subject revolves around determining the extent to which humans can accurately, precisely, and consistently estimate statistical measures, particularly those of centrality, dispersion, and accumulation. To that end, generally speaking it appears that our ability to estimate quantities ranges from reasonable to not good, depending on the quantity: proportions appear to be easiest, and variances most difficult.[2,6,7] Such psychological investigations form what might now be considered the bedrock of prior elicitation, as they came chronologically early relative to statistical elicitation works and lend credibility to the ultimate purpose of the endeavor.

Apart from the psychological questions involved, various technical and practical statistical problems abound, although they are far scarcer in the literature. From the theoretical statistician's perspective, elicitation methods need to be proposed, formalized mathematically, and solved. In other words, typically expert opinion is elicited as numerical summaries of a distribution, and the summaries are converted into the usual distribution representations (i.e. with parameters); so the questions become: what information should be elicited?; how can the procedure be formalized into a general method?; and, once formalized, how might solutions be obtained? In this article these elements of the elicitation process are illustrated with the mode/percentile method for a population proportion. While this method has been described in various places, a proper mathematical formulation appears to be missing from the literature, so we present one in Section 2 before turning to our primary interest concerning practical implementations.[2,1]

A less well-studied aspect of prior elicitation comes from the applied statistician's perspective – for prior elicitation to be practical, there need to be tools that enable its use. Since even the simplest of prior elicitation procedures requires computers for calculations, what is really needed are computer programs dedicated to facilitating prior elicitation. We call such implementations facilitators; they are simply computer programs that aid the elicitation process in some way.

To be truly practical, facilitators need to be *good* implementations: they should be stable (that is, they should not crash), free, readily accessible online or for download, cross-platform, fast, easy to learn and use, and so on. In principle, all of the routines required for elicitation can be performed in virtually any sufficiently flexible programming language: C, R, Python, Java, JavaScript, etc.; however, both the effort in development and quality (read: practicality) of the final product depend significantly on the language used. In this article we present a proof-of-concept facilitator for the prior elicitation of the population proportion π designed using Wolfram Research Inc.'s (WRI's) Computable Document Format (CDF) technology, written in Mathematica. After briefly advocating for the use of the CDF technology as a proper framework for the creation of such facilitators, we provide an overview of the facilitator available online at http://blogs.baylor.edu/baylorisms/beta-facilitator/.

The mode/percentile (MP) method

Discussions concerning elicitation are usually presented not as general principles that can be applied to any quantity, as in ``irrespective of the quantity elicited, the expert should specify two percentiles of the prior distribution”, but rather according to the quantity in question (see e.g.,[3,8]). In other words, sources typically present methods for eliciting the population proportion π in one section, methods for the population mean μ in another section, and so on, even if the principles motivating the methods overlap. This suggests that facilitators should be created with respect to the quantity being elicited, e.g. a π facilitator, a μ facilitator, and so on. In this section, we discuss one of the recommended procedures, the mode/percentile (MP) method, for the elicitation of a population proportion. Since the principle motivating the method is quite general, we formulate the MP method for the general case first and then simply apply it to the population proportion example. In Section 5, we present the corresponding facilitator. As in previous works on the prior elicitation of a population proportion, we assume a binomial sampling model with a conjugate beta prior. We further assume that $\alpha, \beta > 1$ to ensure a unique mode.

Previous work on the elicitation of the binomial proportion

Winkler RL[9] lists four methods for the prior elicitation of π : the hypothetical future sample (HFS) method, the equivalent prior sample information (EPS) method, the cumulative distribution function (CDF) method, and the probability density function (PDF) method. These are summarized in Table 1 below.

In cases where experts are inconsistent in their specifications, regularity is achieved through either statistician/expert dialogue or mathematical fitting such as least squares. Generally speaking, most methods for elicitation revolve around the specification of percentiles, means, medians, or modes, of which countless variations exist. For example, Hughes G[10]] list ten entirely separate methods. Eight of these are variations of those in Table 1; two more rely on the prior predictive beta-binomial distribution.

As our interest here is the implementation rather than the method, we now describe a variant of the CDF and PDF methods in Table 1 that we call the mode/percentile or ``MP” method; it can be found in various applied Bayesian texts.[1,10] In this method, the expert provides the ``best guess for the probability”, which is considered to be the mode, and the ``biggest value the probability could reasonably be”, which is considered to be the 95th percentile.[1]

Table 1 The four methods of prior elicitation for π as presented in[8]

Method	Quantities elicited
HFS	Two means: one from experience and one after begin given a hypothetical dataset
EPS	The mean and a corresponding sample size
CDF	Two or more percentiles
PDF	The mode and the two points half as likely

Abbreviations: HFS: The Hypothetical Future Sample Method; EPS: The Equivalent Prior Sample Information Method; CDF: The Cumulative Distribution Function Method; PDF: The Probability Density Function Method are each Methods of Prior Elicitation

A General formulation of the MP method

We now present a general formulation for the MP method and apply it to the population proportion problem with a beta prior. Let θ denote a (single) parameter to be elicited from a parametric family of distributions with probability density function $f(\theta|\eta)$, where $\eta \in H \subset R^k$, and k is the number of prior parameters to be elicited. Let p denote the u^{th} percentile, where $u \in (0,1)$. The MP method therefore corresponds to the solution of the system of equations

$$m = \arg\max_{\theta} f(\theta|\eta) \tag{1}$$

$$p = \int_{-\infty}^{u} f(\theta|\eta)\, d\theta \tag{2}$$

for m, where m, p, and u are considered known constants from the elicitation process.

In general, solving the simultaneous system (1) and (2) appears to be a very difficult task; however, in many practical situations it is readily solvable numerically. In particular, the case where $\theta = \pi$, the population proportion, and f is the beta density, so that $\eta = (\alpha, \beta)$, proves to be quite easy to solve. The assumption that $\alpha, \beta > 1$ implies that the unknown mode in the right hand side of (1) simplifies to $(\alpha-1)/(\alpha+\beta-2)$. Moreover, as[10] points out, the ubiquity of the integral on the right hand side of (2) when f is a beta PDF has lead to several decades of investigation by numerical analysts, resulting in very efficient routines for its evaluation for arbitrary u, α, and β. These two facts combine to change the apparently very difficult system into the much easier one.

$$m = \arg\max_{\theta} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1}(1-\theta)^{\beta-1} \tag{3}$$

$$p = \int_0^u \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1}(1-\theta)^{\beta-1} d\theta, \tag{4}$$

$$m = \frac{\alpha-1}{\alpha+\beta-2} \tag{5}$$

$$p = h(u, \alpha, \beta) \tag{6}$$

Where h is a function that, while lacking a closed form, can be evaluated at will. This last system can be sequentially solved by rearranging (5) to $\alpha = (1+(\beta-2)m)/(1-m)$ and substituting the result into (6) to obtain

$$p = h\left(u, \frac{1+m(\beta-2)}{1-m}, \beta\right) \tag{7}$$

a univariate root-finding problem in β that can be solved almost instantly with an off-the-shelf routine. The parameter α is then obtained by back-substitution.

Thus, while the MP method appears difficult in general, in the case of the population proportion with a beta prior the problem turns out to be quite simple. The same is true of the Poisson rate problem with a gamma prior and the normal mean problem with a normal prior. Care ought to be taken, however, as the problem appears to be somewhat more complex than its first impression. Intuitively, one might assume that the two pieces of information elicited (m and p) are simply converted to the two canonical parameters (α and β); however, after some investigation this is seen to be too naive. For example, no

Citation: Kahle D, Stamey J, Natanegara F, et al. *Facilitated prior elicitation with the wolfram CDF*. *Biom Biostat Int J*. 2016;3(6):193‒197.
DOI: 10.15406/bbij.2016.03.00081

solution (conversion) exists for the specification $m=.80$, $p=.92$, and $u=.90$, whereas $m=.35$, $p=.30$, and $u=.31$ has two solutions. Thus, the mapping from m and p to α and β is not simply 1–1. For more insight into these details see [11].

Current implementations

There are currently a few nice solutions for prior elicitation in some specific contexts (proportions, rates, means, etc.). In this section, we very briefly mention three such facilitators, two of which are effectively the same.

MATCH and SHELF

Known by the acronyms MATCH and SHELF, the first two facilitators are fairly general purpose tools for prior elicitation. Created by Ed Morris and Jeremy Oakley, MATCH stands for "Multidisciplinary Assessment of Technology Centre for Healthcare," and is a browser-based tool available at http://optics.eee.nottingham.ac.uk/match/uncertainty.php. The tool was created with a blend of PHP and JavaScript and features interactive elicitation over a user-specified interval support for the normal, Student-*t*, scaled beta, gamma, log-normal, and log-Student-*t* distributions using five different methods. Four of the methods are percentile-type variants (i.e. variants of the CDF method in Table 1) called the quartile, tertile, probability, and hybrid schemes. The last, called the roulette method, essentially allows users to draw their belief by stacking boxes ("chips") as in a histogram. The exact details of the conversions of the expert-specified quantities to a probability distribution (parameters) are not clear since they are not included; but once the information is specified, the information the expert provides is automatically converted to the specified probability distribution, the user is presented with the parameters of the elicited distribution and a nice interactive graphic linked to sliders that provides percentiles of the elicited distribution. A screenshot of MATCH can be seen in Figure 1,2.

MATCH is essentially a nice web-interface that runs routines from SHELF, the SHeffield ELicitation Framework, created by Tony O'Hagan and Jeremy Oakley. SHELF is a set of R functions, documents, and templates that facilitate the elicitation of probability distributions from either a single expert or from a group of experts.[12] Instead of an R package, SHELF is a giant R script file that can be sourced into R. It is written entirely in R and uses interactive panels via the rpanel package.[13] Once loaded, it functions almost identically to MATCH but with a rougher looking and less responsive interface. Originally released in 2008, SHELF was updated to version 2.0 in late 2010. It is available online at http://www.tonyohagan.co.uk/shelf/.

Betabuster

Created by Chun-Lung Su at UC Davis, BetaBuster is a facilitator for eliciting beta distributions. Freely available online at http://www.epi.ucdavis.edu/diagnostictests/betabuster.html, BetaBuster is a Windows-only executable written in Java. It is less interactive than MATCH and SHELF, having only input fields and clickable increase/decrease tabs where MATCH and SHELF have sliders (and fewer input fields). The chief difference between the two, however, is the method of elicitation. Where MATCH and SHELF use predominantly CDF variants, BetaBuster implements the MP method described in Section 2.2; it is thus in many ways a predecessor of the facilitator presented in this work. While BetaBuster is less capable than MATCH and SHELF, it is notable because it was perhaps the first facilitator. Nevertheless, while BetaBuster is not state-of-the art, it is still common to see it used in real analyses (e.g.[14]), probably because the MP method and BetaBuster's implementation of it are more inviting than MATCH, whose methods have unfamiliar names and a more flexible but less intuitive interface. While MATCH's interface is actually very simple, for first timers it can be surprisingly difficult to use.

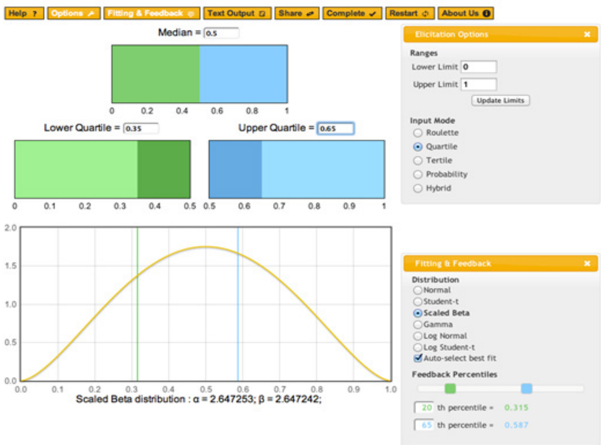

Figure 1 The Multidisciplinary Assessment of Technology Centre for Healthcare (MATCH) tool.

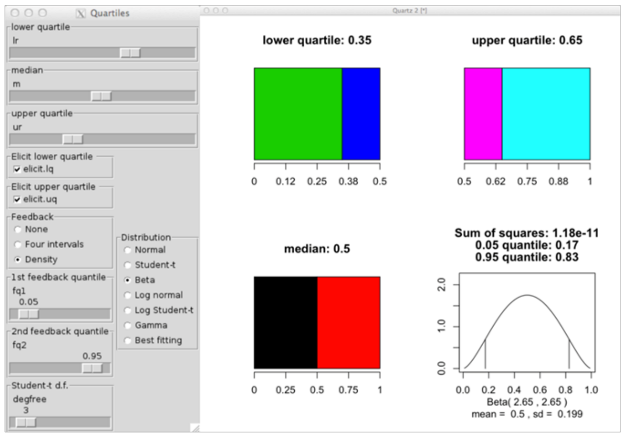

Figure 2 The Sheffield Elicitation Framework (SHELF) tool.

Wolfram research inc.'s (WRI) computable document format (CDF) technology

In Summer 2011, Wolfram Research Inc. (WRI), the developers of the popular computer algebra system Mathematica and online computational knowledge engine Wolfram|Alpha, released a new technology called the Computable Document Format (CDF). The innovation came in the form of (1) a new kind of file format with extension .cdf and (2) a new application called the CDF Player. When run, the files themselves (CDFs) work just like interactive computer programs, where one. cdf file is one program, and they can be run as stand-alone programs or in the browser embedded in an HTML webpage.

CDFs are written using Mathematica. The relationship between Mathematica, the CDF Player, and a CDF file can be well understood by analogy with the Adobe Acrobat, Adobe Reader, and portable document format suite created by Adobe Systems Inc. (Adobe).

Adobe is the maker of the ubiquitous portable document format (with extension .pdf). Documents in this format, called PDFs, are often authored and edited in Adobe Acrobat, a paid application, but can be viewed by anyone with Adobe Reader, a freeware application. Analogously, CDFs are authored in Mathematica, but can be viewed using the CDF Player which, like Adobe Reader, is free and can operate through the browser. CDFs therefore meet a basic requirement for any practical facilitator: they are easily disseminated and can be freely used by anyone.

CDFs can draw from nearly the full range of Mathematica capabilities, which makes the CDF technology an incredibly powerful platform for developing facilitators. In particular, CDF-based facilitators enjoy huge design advantages over other programming languages and frameworks due to their ability to draw on the power of Mathematica. Here are a few examples:

1. Mathematica has a vast array of symbolic and numerical algorithms that authors can draw upon with simple commands. Numerical integration, optimization, linear algebra and much more are all built-in. Moreover, they are customizable. For example, when a numerical optimization problem is presented, not only are standard numerical algorithms such as Newton's method available, but so too are other schemes: differential evolution, Nelder-Mead, simulated annealing and so on, all by simply changing a word, not re-implementing the whole scheme. Symbolic schemes and/or simplifications and rearrangements can also be used. To top it off, the implementations are efficient, stable, and robust.

2. Mathematica boasts the largest array of probability distributions of any current computing platform. This includes key functionality not currently available in other platforms: (1) symbolic representation and manipulation of probability distributions, (2) symbolic handling and numerical evaluation of distribution-related functions (PDFs/CDFs, survival, hazard, and characteristic functions) and functionals (moments/expected values, percentiles), and (3) the analytical and numerical evaluation of probability statements. This is on top of standard functionality such as sampling from distributions and fitting linear and nonlinear models.

3. Mathematica includes a tremendous variety of graphics options, from 2D and 3D plotting of functions and data (e.g. scatterplots) to 2D and 3D histograms and contour plots.

4. Perhaps most importantly, Mathematica has very simple mechanisms for making interactive content. For example, a user can create an interactive graphic of a normal density with sliders to change its mean and standard deviation in real-time using a mere few lines of code. All of the interactive programmatic overhead and graphical user interface (GUI) details are handled by Mathematica and are highly customizable.

While the Mathematica-inherited capabilities alone are nice, they are by themselves insufficient for a platform for developing facilitators. In particular, to be truly practical facilitators need to be freely available to be used. This is precisely the beauty of the WRI's CDF Player technology. Not only can the routines of facilitators be easily designed and written in Mathematica, but as CDFs they can be created and distributed as either stand-alone applications or seamlessly embedded in webpages, a particularly convenient medium.

The combination of each of the above makes Mathematica a potent development platform for facilitators. Compared to other platforms (e.g. Java or Javascript), the development cycles of facilitators can be much shorter with Mathematica due (1) its wide array of built-in functionality, and (2) its automated handling of so many of the interactive GUI tasks. Taken together, these allow the creator to focus on the statistical details of the facilitator rather than the programming details of the software engineering. It should further be noted that Mathematica and the CDF Player can be run on a wide array of operating systems (Windows, OS-X, Linux), and the CDFs can be viewed through any of several standard browsers (Internet Explorer, Safari, Firefox, Chrome, etc.). Lastly, CDFs also function on multi-touch devices with much the same development simplicity. This kind of interface allows for a significantly enhanced prior elicitation experience.

The beta prior facilitator

The beta prior facilitator is a proof-of-concept CDF facilitator available through any standard web browser at http://blogs.baylor.edu/baylorisms/beta-facilitator/. It is also available for download from the same site, thereby allowing it to be used offline in exactly the same way and with the same functionality as online. A screenshot of the facilitator is included in Figure 3,4.

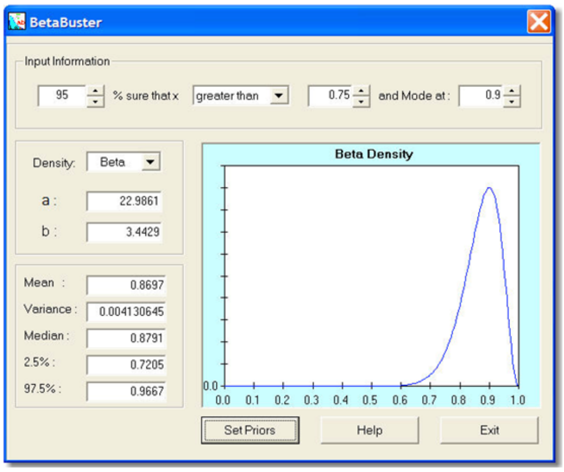

Figure 3 The Beta Buster tool.

At root, the CDF beta facilitator contains many of the features of BetaBuster in a greatly enhanced form. Input fields to specify the mode (m), percentile (p), and corresponding percentage (u) are present, but they can now be provided either on the cumulative distribution or survival scales (i.e. "less than p " or "greater than p "). The confidence region is illustrated in the graphic as a shaded region that changes with the user's specifications. Instead of input fields for α and β , the CDF facilitator provides sliders.

However, the facilitator has several elements that BetaBuster does not have. In the control panel on the bottom, a ``Scale ESS" slider allows the user to interactively change how informative the elicited prior is *after elicitation*. In Bayesian data analysis, one of the measures of the strength of a prior -- how influential it will be an analysis -- is its equivalent sample size (ESS,[1,15]). In the context of a binomial sampling scheme, the beta prior on π can be understood in terms of a hypothetical dataset. In particular, if a $Beta(\alpha,\beta)$ prior is elicited, it can be thought of as observing α successes out of $\alpha+\beta$ trials. For example, a prior with $\alpha=3$ and $\beta=7$ can be thought of as having observed an experiment where a success occurred in three out of ten trials.

Citation: Kahle D, Stamey J, Natanegara F, et al. *Facilitated prior elicitation with the wolfram CDF*. *Biom Biostat Int J.* 2016;3(6):193‒197.
DOI: 10.15406/bbij.2016.03.00081

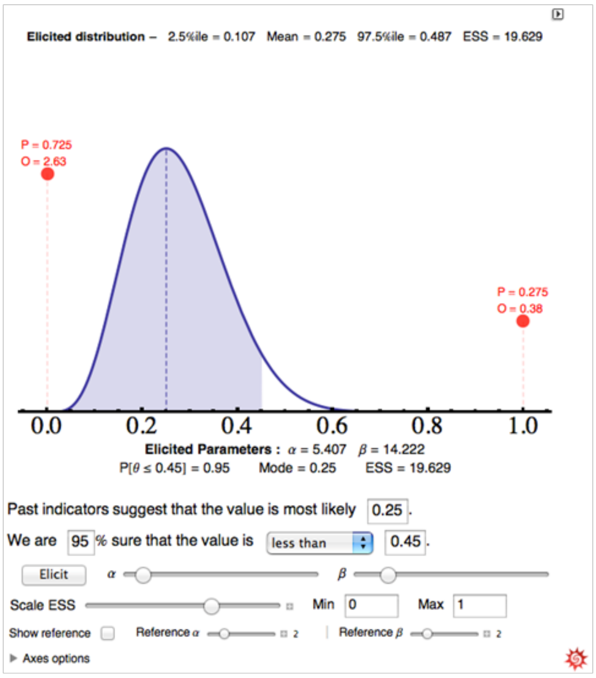

Figure 4 The Beta Facilitator – An Elicitation Tool for the Population Proportion Using a Beta Prior.

The purpose of the Scale ESS slider is to allow the user to increase or decrease the strength of the prior after it is elicited. When a prior is elicited, the parameters are provided directly below the graphic along with the ESS. If, after elicitation, the statistician or expert wishes to decrease the strength of the prior, they can use the ESS slider to "discount" the ESS of the prior. This works by scaling the elicited α and β; in effect trying to reduce the ESS while minimally affecting the elicited quantities. The graphic and all quantities change in real-time as the slider is adjusted, providing the user with instant feedback.

It has been suggested that elicitation is best performed in an environment with instant feedback.[4] In particular, if done properly studies suggest that feedback might enable the expert to calibrate their information to produce priors more faithful to their actual belief: "To improve experts' calibration, the feedback from the known procedures must be immediate, frequent, and specific to the task..." ([4], p.18, summarizing[16]). To provide such feedback, the beta facilitator provides the marginal probability and odds of success and failure in the graphic, labeled ``P" and ``O", assuming the elicited distribution. These are calculated based on the prior predictive beta-binomial distribution, which is built-in in Mathematica. As sliders change, every element of the graphic changes dynamically in real-time, so that the statistician and the expert know right-away the ramifications of the expert's belief on a meaningful scale.

A few other functional elements are included in the facilitator as well: the ability to change the axes/viewing window and the inclusion of a reference beta distribution with slider-manipulable parameters. These can be helpful when trying to manually select a beta distribution that in some way matches the elicited beta.

Discussion and future directions

As advances are made in computer technology, specifically software technology and especially web-based technologies, there are increasingly many opportunities for the computer-assisted conversion of expert knowledge into probability distributions to be incorporated directly into statistical analyses. Currently, the arguments in this article suggest that WRI's CDF technology ushers in a new state-of-the-art platform for quickly developing top-notch facilitators with minimal effort.

Acknowledgement

None.

Conflict of interest

None.

References

1. Christensen R, Johnson WO, Branscum AJ, et al. *Bayesian ideas and data analysis: An introduction for scientists and statisticians*. USA: CRC Press; 2011.
2. Garthwaite PH, Kadane JB, O Hagan A. Statistical Methods for Eliciting Probability Distributions. *Journal of the American Statistical Association*. 2005;100(470):680–701.
3. O Hagan A, Buck CE, Daneshkhah A, et al. *Uncertain Judgements: Eliciting Experts' Probabilities*. USA: Wiley; 2006.
4. Meyer MA, Booker JM. Eliciting and Analyzing Expert Judgment: a Practical Guide. *Society for Industrial and Applied Mathematics*. 1991;7:29.
5. Cooke RM. Experts in Uncertainty: Opinion and Subjective Probability in Science. *Ethics*. 1991;103(3):599–601.
6. Shuford EH. Percentage Estimation of Proportion as a Function of Element Type, Exposure Time, and Task. *Journal of Experimental Psychology*. 1961;61(5):430–436.
7. Lathrop RG. Perceived Variability. *Journal of Experimental Psychology*. 1967;73(4 pt 1):498–502.
8. Lunn D, Jackson C, Spiegelhalter DJ, et al. The BUGS book: A Practical Introduction to Bayesian analysis. CRC Press. 2012;52:91.
9. Winkler RL. The Assessment of Prior Distributions in Bayesian Analysis. Journal of the American Statistical association. 1967;62(319):776–800.
10. Hughes G, Madden L. Some Methods for Eliciting Expert Knowledge of Plant Disease Epidemics and their Application in Cluster Sampling for Disease Incidence. Crop Protection. 2002;21(3):203–215.
11. Wu Y, Shih WJ, Moore DF. Elicitation of a Beta Prior for Bayesian Inference in Clinical Trials. Biom J. 2008;50(2):212–223.
12. http://www.tonyohagan.co.uk/shelf/
13. http://www.jstatsoft.org/v17/i09/.
14. Okura H, Nielsen SS, Toft N. Prevalence of Mycobacterium Avium Subsp. Paratuberculosis Infection in Adult Danish non-Dairy Cattle Sampled at Slaughter. *Preventive Veterinary Medicine*. 2010;94(3):185–190.
15. Morita S, Thall PF, Müller P. Determining the Effective Sample Size of a Parametric Prior. *Biometrics*. 2008;64(2):595–602.
16. Lichtenstein S, Fischhoff B, Phillips LD. Calibration of Probabilities: The State of the Art to 1980. *Springer*. 1977;306–334.

Citation: Kahle D, Stamey J, Natanegara F, et al. *Facilitated prior elicitation with the wolfram CDF. Biom Biostat Int J.* 2016;3(6):193–197. DOI: 10.15406/bbij.2016.03.00081