18th World IMACS / MODSIM Congress, Cairns, Australia 13-17 July 2009
http://mssanz.org.au/modsim09

Expert elicitation and its interface with technology: a review with a view to designing *Elicitator*

Low Choy, S. [1]**, A. James** [2] **and K. Mengersen** [1]

[1] *School of Mathematical Sciences, Queensland University of Technology, Brisbane, Queensland, Australia*
[2] *High Performance Computing unit, Queensland University of Technology, Brisbane, Queensland, Australia*
Email: s.lowchoy@crcbiosecurity.edu.au, elicitator@qut.edu.au

Abstract: Expert knowledge is valuable in many modelling endeavours, particularly where data is not extensive or sufficiently robust. In Bayesian statistics, expert opinion may be formulated as informative priors, to provide an honest reflection of the current state of knowledge, before updating this with new information. Technology is increasingly being exploited to help support the process of eliciting such information. This paper reviews the benefits that have been gained from utilizing technology in this way.

These benefits can be structured within a six-step elicitation design framework proposed recently (Low Choy et al., 2009). We assume that the purpose of elicitation is to formulate a Bayesian statistical prior, either to provide a standalone expert-defined model, or for updating new data within a Bayesian analysis. We also assume that the model has been pre-specified before selecting the software. In this case, technology has the most to offer to: targeting what experts know (E2), eliciting and encoding expert opinions (E4), whilst enhancing accuracy (E5), and providing an effective and efficient protocol (E6).

Benefits include:

- providing an environment with familiar nuances (to make the expert comfortable) where experts can explore their knowledge from various perspectives (E2);
- automating tedious or repetitive tasks, thereby minimizing calculation errors, as well as encouraging interaction between elicitors and experts (E5);
- cognitive gains by educating users, enabling instant feedback (E2, E4-E5), and providing alternative methods of communicating assessments and feedback information, since experts think and learn differently; and
- ensuring a repeatable and transparent protocol is used (E6).

Software has been successfully used for facilitating elicitation in other modelling frameworks, such as Bayesian networks (Uusitalo 2007). However there are few general purpose packages available for supporting elicitation in a Bayesian setting. The most accessible, *Elicitor* (Kynn, 2006), implements an indirect approach to elicitation for regression, and is implemented as a module within one of the more widely used freely accessible Bayesian packages WinBUGS. In this paper we illustrate how many of these potential technology-based benefits have been incorporated into a new tool *Elicitator*, designed to support elicitation using a new method (Low Choy et al., in press).

The features of a modern computing platform can be reported according to the elicitation services provided: communication, cognitive benefits including thinking and learning, automation and accuracy. Some features provide the full range of services, such as feedback and education. Other features provide particular services, namely exploration, interaction and providing several modes of communication. Several benefits arise from dynamic and interactive graphical interfaces, use of freely available and open source libraries, and incorporating a relational database. *Elicitator* provides these benefits, and is the first such tool to also support elicitation from multiple experts; made possible by managing projects and underlying data integrity. The use of a modern computing environment shows potential to simplify and streamline the elicitation process in Bayesian regressions with several covariate sets as well as in Bayesian networks with complex relationships.

Keywords: *Expert elicitation; Prior information; Informative Bayesian analysis; Software tool; Design; Human-Computer Interface*

1. INTRODUCTION

Quantifying expert opinions underpins many modern modelling efforts in a world where preliminary decisions are often required before extensive and robust data are available. This equally applies whether expert knowledge is used to provide standalone models (e.g. Low Choy et al., 2009, Case A) or for later inclusion with observed data within a Bayesian framework (e.g. Low Choy et al., 2009, Case B).

A desire to achieve elicitation in a transparent, repeatable and robust manner suggests judicious use of technology (Kynn, 2006; Low Choy et al., 2009). Despite this, few general purpose software packages have been developed to support elicitation, as lamented recently by Leal et al. (2007). Most of these are not freely available, instead applied to particular case studies (Chaloner and Duncan 1983; Du Mouchel 1988; Goldman et al. 1988; Chaloner et al., 1993; O'Hagan, 1997; Kadane and Wolfson, 1998; Al-Awadhi and Garthwaite, 2006; Denham and Mengersen, 2007; O'Leary et al., 2008). Few elicitation tools have been designed for more general application on standard computer platforms. These rare efforts have been implemented in TROLL and CADA (Kadane et al., 1980), Excel (Kuhnert et al., 2005; Martin et al., 2005; Leal et al., 2007; O'Leary et al., 2008), or Blackbox Pascal as an OpenBUGS module (Kynn, 2006).

This paper was motivated by the need to construct a software tool to support a new, indirect approach to elicitation for regression (Section 2). We review the ways in which technology has been used previously to support elicitation (Section 3). These concepts informed design of a new software tool *Elicitator* (Section 4). Finally we summarize how technology may support elicitation, now and in the future (Section 5).

2. TECHNOLOGY AND ELICITATION: A REVIEW

A statistical approach to expert elicitation treats it as a data collection exercise. A well-designed approach requires several main steps (**Table 1**): divine the purpose (E1), formulate the statistical model (E3), appropriately target (E2) and encode (E4) expert knowledge, and design an accurate (E5) and repeatable elicitation protocol (E6). Suppose that the purpose (E1) and statistical model (E3) are specified before choosing elicitation software. This section focuses on contributions that technology may make within steps (E2,E4-E6) of the elicitation framework.

Table 1. Main steps in designing an elicitation process; (for details see Low Choy et al., 2009).

Step	Description
E1	Divine the purpose and type of information required, e.g. standalone model or prior for input to Bayesian analysis or other model.
E2	Appropriately target expert knowledge to ensure they can accurately conceptualize and communicate it, yet can be translated mathematically.
E3	Formulate the model. A Bayesian setting requires specification of the data model (likelihood) and the expert model (prior).
E4	Design method for encoding elicited info Determine overall strategy (direct or indirect), the summary statistics to be elicited, communication methods, how to capture expert uncertainty, and statistical inference procedure for calculating priors.
E5	Manage uncertainty Minimize major sources of bias, including cognitive and linguistic. Consider eliciting from multiple experts. Verify and validate elicited information.
E6	Specify the elicitation protocol and logistics, from motivating and selecting experts to preparing them, as well as assessing how expert opinion is to be combined, then determining delivery of elicitation (interview, questionnaire, software & other tools).

2.1. Targeting Expert Knowledge (E2)

Typically expert elicitation is only required when experts have not previously had to refine and quantify their knowledge in the desired way. Otherwise this information is published, for instance as an expert 'model', for example as bioregional boundaries (Low Choy et al. 2009; Case B). An important role for technology is thus to assist the expert in *exploring* their knowledge.

Ideally this exploration is achieved within a *familiar* environment suited to their knowledge and communication preferences (Spetzler and Stäel von Holstein, 1975; O'Hagan et al., 2006; Kynn, 2008). The tool *Elicitor* (Kynn, 2006) permits users to tailor graphs, using an appropriate title and units for the response variable (e.g. number of occupied sites). If run on a laptop computer, an elicitation tool can be highly portable, allowing consultation with experts in their customary environment, indoors or out.

Familiarity may also be a direct consequence of the visual interface. For instance experts may be comfortable with exploring, both conceptually and graphically, the change in response (y-axis) with respect to a single covariate (x-axis), holding all other covariates constant. Using graph paper Willems et al. (2005) asked experts to express their uncertainty in navigable channel width (y-axis), conditional on the previous year's observation (x-axis). Similarly, two software packages (Kynn, 2006; Al-Awadhi and Garthwaite, 2006) have asked experts to manipulate an electronic graph.

These two software packages were used to elicit from landscape ecologists with well-developed conceptual understanding. In contrast a map-based tool (Denham and Mengersen, 2007) targeted ecologists with strong field-based knowledge. Using environmental factors mapped in a GIS, experts were asked to describe ecological response at particular sites. Both of these indirect approaches capitalized on an expert's ability to express their opinions about *observable* quantities (Kadane et al., 1980) but required interviews of an hour or more to capture such rich information. An alternative more rapid elicitation approach simply asked whether each covariate was thought to increase, decrease or have no effect on response. This has been implemented in a spreadsheet (O'Leary et al. (2008), with similar aim but different mathematical formulation to previous work (Kuhnert et al., 2005; Martin et al., 2005). *Targeting expert knowledge* in these three different ways provided very different results (O'Leary et al., 2008).

2.2. Encoding Expert Knowledge (E4)

By *automating tedious* or *repetitive* tasks, technology ensures consistency and reduces error—by elicitors or experts (Kynn 2008, Recom. 5). In *Elicitor* (Kynn 2006), the priors are recalculated and plots adjusted each time the expert updates their elicitations. Similarly the tool of Denham and Mengersen (2007) re-computes the prior distributions whenever an expert updates their elicitations. To support the multivariate mixture model priors encoded in Accad et al. (2005), several tools had to be developed—both to transfer information between GIS and statistical packages, as well as to compute encoded priors—for prior models based on different variable sets. Instantaneous computation ensures consistency, accuracy, and responsiveness to small incremental changes, and also obviates the need for calculation 'downtime'.

Using software to automate mundane tasks can free the elicitor to focus more on *encouraging interaction* with the expert. With 'lag' time virtually eliminated, this can improve the flow in the elicitation conversation, allowing the elicitor to provide feedback or follow-on questions in 'real time'. Hence more time can be devoted to developing the interpersonal relationship required to establish rapport and trust necessary for elicitation (O'Hagan et al., 2006).

Whenever expert models (prior distributions) are recalculated *feedback* can be provided, helping experts to: refine their understanding of definitions and requirements, explore their knowledge, maintain self-consistency and therefore greatly reduce cognitive biases (Kadane and Wolfson 1998; Low Choy et al., 2009, point vi). Unusual or influential elicitations can be identified using feedback diagnostic plots (Denham and Mengersen, 2007), density plots or boxplot-style summaries (Leal et al., 2007).

2.3. Eliciting Accurate Expert Knowledge (E5)

In addition to acknowledging that tools should target expert knowledge appropriately (E2), elicitation tools can be used to tailor elicitation by recognizing that people *learn and think in different ways*. As noted by Chaloner and Duncan (1983) "Experts are very different both in terms of opinions and how they parse the problem. What was easy for one expert was hard for another. Developing a single method for elicitation that is good for everybody will be difficult, if not impossible." Statistical learning and thinking can be undertaken in aural, oral, visual or kinetic modes. Some people respond better when information is imparted personally (aural), when they have the opportunity to discuss it (oral), or through graphs or maps emphasizing geometric or spatial aspects and patterns (visual). Others require an activity (kinetic) such as interacting with a computer or a field trip. In addition some people relate better to numeric information when presented concretely in context with a few preferring the more abstract presentation in terms of equations.

Many tools have taken advantage of *combining modes of communication*. When information is presented in several modes, experts can learn about their other less utilized mode(s), and in addition elicitors can reach a wider range of experts. Some tools supplement aural and oral modes with the visual mode via graphs (Chaloner and Duncan, 1983; Du Mouchel, 1988; Kynn, 2006; Leal et al., 2007 or maps in GIS or Geographic Information Systems (Lehmann et al., 2002; Denham and Mengersen, 2007). In many of these cases, modern graphical user interfaces (GUIs) provide a kinetic mode via dynamic interaction with graphs.

Presenting the same information in more than one way may also encourage the expert to *explore* their knowledge from different perspectives, encouraging an *active* rather than passive process of distilling their knowledge. In turn this releases experts from the pressure to understand all representations of statistical information, in fact encouraging *questioning* and therefore learning. For example users can edit summary statistics (O'Hagan, 1998; Kynn, 2006) or manipulate a plot (Denham and Mengersen 2007) and obtain feedback on other summary statistics. Self-consistency is encouraged by allowing comparisons between individual elicitations provided by an expert (Du Mouchel, 1988; Low Choy et al., 2009, Case A; Denham and Mengersen, 2007).

2.4. Elicitation Protocol (E6)

Solid preparation reduces linguistic uncertainty by ensuring that modellers understand the expert's language and context, and experts understand what information is sought. Multiple modes of presentation also provide an opportunity to *educate* users about the model and other statistical concepts. Software may also contain tutorials: such as tools to help visualize probabilities, e.g. a probability wheel and random samples of proportion of area (Kynn, 2006) or a graphical introduction to multiple comparisons (Du Mouchel, 1988).

A key benefit of automation can be the *repeatability* and *documentation*, and therefore *transparency*, of both intermediate and final results. For instance a tool (Kynn, 2006) may report final estimated prior distributions, together with elicited information in textual, database and graphical form. Thus the elicitation is documented in a way that can be revised at another time. To this end, recent tools (Denham and Mengersen 2007; Kynn 2006) exploit graphical user interfaces to provide windows that can be independently manipulated and saved.

A practical advantage of elicitation software is that several aspects of elicitation methodology and logistics can be set to *defaults*. This may be helpful for novice elicitors. For example, Kynn (2006) enforces the order of elicitation to start with elicitation of an intercept, followed by covariates one at a time, in the same order as the regression equation. The tool of Du Mouchel (1988) requires credible intervals on individual effects to be specified first before advancing to multiple comparisons.

3. CASE STUDY: ELICITATOR

Here we outline how a new tool supports a novel approach to elicitation of priors for logistic regression; details in Low Choy et al. (in press). For logistic regression, coefficients β relate binary observations $Y_i \sim \text{Bern}(\mu_i)$, to covariates X_{ij}, $j=1,\ldots,J$ for cases $i=1,\ldots I$ using link $\text{logit}(\mu_i)= \beta X_i$ to probability of success μ_i.

3.1. Elicitation method

This indirect elicitation method targets the expert's assessment of probability of presence (E2) and extends the conditional mean approach of Bedrick et al. (1997). The stages are (Figure 2): [S1] setup project, [S2] elicit and [S3] encode then [S4] verify expert knowledge, and finally [S5] output an expert model.

Step [S2] follows an indirect strategy, and asks experts to estimate the probability of success Z_k for several cases k with known covariates $X_{1k}, X_{2k}, \ldots, X_{Jk}$. They are asked [S2a] to describe the range of values with varying likelihood (percentiles) as well as their best estimate (mode). This information is used [S2b] to numerically estimate μ_k and γ_k in $p(Z_k|X_k)$. Feedback is provided and the expert given an opportunity to modify [S2c]. This is repeated for several cases $k=1,\ldots,K$. In step [S1] the elicitor imports covariates X_k for these elicitation cases.

At step [S3], the information provided by the expert across all the cases can be combined to form the expert model. A Beta regression is used to relate expert data Z_k to the covariates

$$Z_k \sim \text{Beta}(\mu_k, \gamma_k), \ \text{logit}(\mu_k)=X_k\beta \qquad \text{(Eqn 1)}$$

linking shape and scale parameters a_k and b_k to the expected probability of presence $\mu_k= a_k/\gamma_k$ and effective "expert" sample size $\gamma_k = a_k+b_k$. Since this is a regression, the usual diagnostics can be supplied and interpreted by the elicitor. This feedback may help clarify whether the overall fitted relationships between expected probability of presence and covariates is consistent with the expert's case-by-case assessments [S4]. If not, the expert may choose to revise elicitations [S2]; otherwise they may accept the final model [S5].

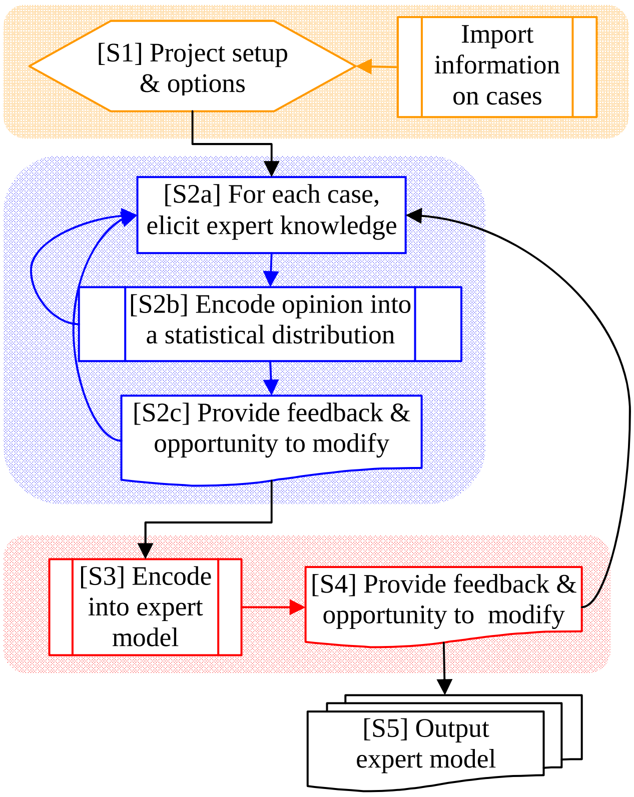

Figure 2: Main stages in new elicitation method.

3.2. Elicitation services

This section focuses on how *Elicitator* harnesses technology to provide the elicitation services within the elicitation framework (E1-E6) and previous literature of Section 2.

Communication services. *Elicitator* is implemented using the Java platform and links to the R package for statistical computation, MySQL for relational database management and JFreeChart for displaying graphs. This reliance on freely available open source libraries and software exploits existing resources and ensures that the tool can be maintained and extended over a long period of time, a key *logistical* issue for potential users (E6). The software is also flexible enough to import text data from a variety of database or GIS packages; the alternative would embed elicitation within a statistical package (Kynn 2006) or GIS (Denham and Mengersen 2007). Hence although *Elicitator* does not depend on GIS software, it can be run alongside a GIS which may be used to *explore* (e.g. zoom in and out, pan) and interrogate covariates spatially (E2). This loose coupling to GIS avoids potential version compatibility issues, particularly with commercial GIS packages, and helps provide a *familiar* environment for experts accustomed to maps (E2). To aid *familiarity* plots can be tailored to the application by importing covariates and category labels (E2) at step [S1].

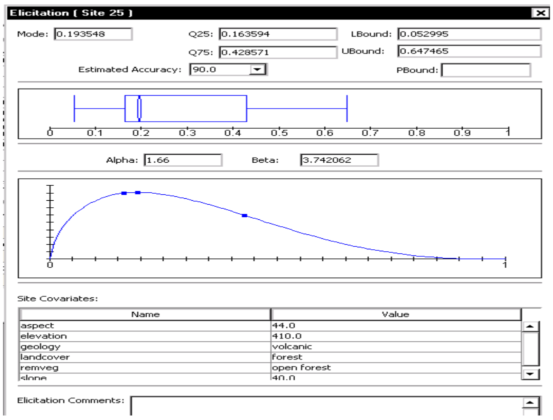

Figure 3. Elicitation window in *Elicitator*.

Cognitive services. *Elicitator* also supports *active questioning*, *cognitive exploration*, *education* and *learning* in two ways. Firstly when encoding information and providing feedback one case at a time (Step [S2c]), the tool addresses experts who think in different *modes* by providing alternative ways of eliciting and communicating assessments (E5). To encode the Beta distribution in Eqn 1 (Figure 3), the distributional parameters, quantiles and a density plot are shown (similar to Denham and Mengersen, 2007), together with a boxplot (Leal et al., 2007) and the mode. These are all *instantly updated* when the expert changes one. Secondly when combining information across cases to encode the prior distribution's parameters, regression diagnostics are provided in step [S4]. This helps experts to switch between conceptualization of case-by-case conditional probabilities of success given case-specific covariates $P(Z_k|X_{1k},X_{2k}, \ldots, X_{Jk})$ and predicted response curves providing the conditional probability of success across values for a particular covariate $Pr(Z \mid X_j, X_{-j} \text{ fixed}, \beta)$. This *feedback* also helps ensure *accuracy* (E5) during encoding (E4).

Automation services. *Elicitator* provides all the previously mentioned benefits of *automating* encoding, both underpinning case-by-case encoding of the probability of success (Step [S2b]) and combined-case encoding of the regression coefficients (Step [S2b,S3]). Thus the elicitor need only focus on the interview (E6) and accuracy (E5) of encoding whilst implementing the encoding method embodied in the software. Moreover this cyclic elicitation-feedback process releases the expert from the pressure of making every assessment right first time (E4,E5). This is particularly important during the initial phases of elicitation when revisiting assessments may be more common: the expert may still be refining their understanding of what is required, and the elicitor may still be guiding the expert on the quantities required and the mechanics of using the software. The process is *repeatable*, assuming that preparatory steps (training, conditioning to biases) and the interview transcript are also documented, and that elicitors are sufficiently trained in the software, the elicitation method and general statistical and probability concepts.

The modular implementation underpinning this tool also provides some *options* in the encoding algorithm, and the ability to extend the tool with more options. The expert may be able to quantify their certainty (E5) in each case's assessment, which when rescaled can be used (in a weighted regression) to ensure the overall expert model fits better to cases that the expert is most sure about.

Accuracy services. Feedback is a key ingredient in ensuring accuracy. In addition, *Elicitator* is unique among elicitation tools, since it explicitly manages *multiple elicitations* (either from a single expert in different sittings, or from several experts). This is managed in the project setup [Step S1], where the elicitor can create or access a phase of an elicitation project comprising: a dataset of covariates for each case, a list of experts together with their elicited information for each case (blank initially), and the expert model diagnostics for each regression coefficient (also blank initially). The underlying relational database provides *data persistence* and *data integrity*, which in *Elicitator* is complicated given the dynamic and interactive interface and the two feedback loops. Once elicitations have been obtained from several experts and encoded into individual priors, their mathematical combination can be achieved in another package (O'Hagan et al., 2006). Small teams could use *Elicitator* to record consensus opinions to provide a psychological combination (O'Hagan et al., 2006). The project database provides up-to-date documentation and therefore transparency.

Use of a tool to consistently elicit opinions in a standard way from each expert enhances comparability and *consistency* amongst opinions. Expert opinions may also be *calibrated*, using a purposively selected calibration dataset, for comparing elicitations to a gold standard. Where calibration data is not available, the tool may also be used to obtain elicitations from a highly regarded 'super' expert.

4. DISCUSSION AND CONCLUSIONS

The elicitation tool *Elicitator* has been designed to incorporate many software services (Section 2), to enable elicitation (Table 1). These naturally fall into the four main service areas of communication, cognitive benefits, automation and accuracy (Section 3.2). Nevertheless, as shown in Table 3, there is some overlap in the four service areas. All four service areas are addressed by providing feedback, a major benefit and education, a spin-off. Similarly exploration and interaction are facilitated to some degree by all four areas, particularly automation. Communication and cognitive services are strongly linked to familiar context and provision of alternative thinking modes. Both automation and accuracy services enable consistent consultation of multiple experts.

Table 4. Features of *Elicitator* used to support elicitation (column 1), and how these achieve the elicitation services: [Com]munication, [Cog]nitive benefits in thinking and learning, [Aut]omation and [Acc]uracy. Contributions are denoted ★ for major and • for minor.

Elicitation service	**Com**	**Cog**	**Aut**	**Acc**
Familiar context	★	★		•
Thinking modes – alternatives	★	★		•
Thinking modes – combination		★		★
Exploration	•	★	★	•
Interaction	★	•	★	•
Educate users	•	•	•	•
Streamline session			★	•
Adapt to incremental changes		•	★	
Instantaneous calculation	•		★	•
Feedback	★	★	★	★
Help consistency, repeatability			★	★
Standardize & promote good design				★
Support calibration			•	★
Multiple experts			•	★

Elicitator therefore provides a modern example of the multifarious ways in which technology can support elicitation. This is enabled by the computing platform, especially dynamic and interactive use of GUI in the elicitation and feedback windows, the underlying relational database system and project management features. Implementation in open-source software and libraries increases the flexibility, stability and robustness of the tool.

In some cases elicitation of responses under various scenarios (covariates) can require enormous input from experts. In Landrum and Normand (1999), clinical experts were asked to estimate the log odds of survival for nearly 900 different "indications" (i.e. covariate sets). When constructing Bayesian Networks from expert knowledge, constraint is often advised: at most three parent nodes per child, up to five categories per parent node, and only a few scenarios requiring elicitation. Elicitator can simplify and streamline elicitation for these types of problems. Statistical design principles can also guide selection of elicitation cases. For instance stratified sampling ensures adequate coverage of covariates (Murray et al, *in press*). With limited time, the choice of elicitation cases may optimally cover covariate space (Kadane et al. 1980). Alternatively elicitation can optimize accuracy by targeting cases that experts know well (Denham and Mengersen 2007).

Elicitator offers proof-of-concept for generalizing the elicitation method to other distributional choices beyond logistic regression. Easily accessible and user-friendly software tools to support elicitation using robust methods, like *Elicitator*, may make elicitation more accessible, particularly on modern computing platforms. Such tools also have the potential to influence methodological choices across a range of applications: enhancing rigour in expert-derived models and informative priors in Bayesian statistics.

5. ACKNOWLEDGMENTS

Funding for the first and third authors was provided under Australian Research Council Discovery Grant DP0667168, and for the second author by the Collaborative Centre for Complex Dynamic Systems and Control. We thank Justine Murray and her experts and ecological application which helped shape *Elicitator*. We thank Robert Denham for inspiration provided by his elicitation tool, and opportunity to review his code.

REFERENCES

Al-Awadhi, S. A. and Garthwaite, P. H. (2006). Quantifying expert opinion for modelling fauna habitat distributions. *Computational Statistics*, 21:121–140.

Bedrick, E. J., Christensen, R., and Johnson, W. (1996). A new perspective on priors for generalized linear models. *Journal of the American Statistical Association*, 91(436):1450–1460.

Chaloner, K., Church, T., Louis, T. A., and Matts, J. P. (1993). Graphical elicitation of a prior distribution for a clinical trial. *The Statistician*, 42(4):341–353.

Chaloner, K. M. and Duncan, G. T. (1983). Assessment of a Beta prior distribution: PM elicitation. *The Statistician*, 32:174–180.

Denham, R. and Mengersen, K. (2007). Geographically assisted elicitation of expert opinion for regression models. *Bayesian Analysis*, 2(1):99–136.

Du Mouchel, W. (1988). A Bayesian model and a graphical elicitation procedure for multiple comparisons. In Bernardo, J. M., de Groot, M. H., Lindley, D. V., and Smith, A. F. M., editors, *Bayesian Statistics 3*, pages 127–145. Oxford University Press: Oxford, UK.

Garthwaite, P. H. and O'Hagan, A. (2000). Quantifying expert opinion in the UK water industry: an experimental study. *The Statistician*, 49:455–477.

Goldman, L., E.F. Cook, D.A. Brand, T.H. Lee, G.W. Rouan, M.C. Weisberg, D. Acampora, C. Stasiulewicz, Walshon, J., Terranova, G., Gottlieb, L., Kobernick, M., Goldstein-Wayne, B., Cohen, D., Daley, K., Brandt, A. A., Jones, D., Mellors, J. and Jakubowski, R. (1988) A computer protocol to predict myocardial infarction in emergency department patients with chest pain. *New Eng. J. Med.*, 318, 797-803.

Kadane, J. B., Dickey, J. M., Winkler, R. L., Smith, W. S., and Peters, S. C. (1980). Interactive elicitation of opinion for a normal linear model. *Journal of the American Statistical Association*, 75:845–854.

Kadane, J. B. and Wolfson, L. J. (1998). Experiences in elicitation. *The Statistician*, 47:3–19.

Kuhnert, P. M., T. G. Martin, K. Mengersen and H. P. Possingham. (2005). Assessing the impacts of grazing levels on bird density in woodland habitat: a Bayesian approach using expert opinion. *Environmetrics* 16:717–747.

Kynn, M. (2006). *Designing elicitor: Software to graphically elicit expert priors for logistic regression models in ecology*. http://www.winbugs-development.org.uk/elicitor/files/designing.elicitor.pdf.

Kynn, M. (2008). The "heuristics and biases" bias in expert elicitation. *Journal of the Royal Statistical Society, Series A*, 171(1):239–264.

Landrum, M. B. and Normand, S.-L. T. (1999). Applying Bayesian ideas to the development of medical guidelines. *Statistics in Medicine*, 18:117–137.

Leal, J., Wordsworth, S., Legood, R., and Blair, E. (2007). Eliciting expert opinion for economic models: An applied example. *Value in Health*, 10(3):195–203.

Lehmann, A., Overton, J. M., and Leathwick, J. R. (2002). GRASP: generalized regression analysis and spatial prediction. *Ecological Modelling*, 157:189–207.

Low Choy, S., Murray, J., James, A. and Mengersen, K. (*in press*). Indirect elicitation from ecological experts: from methods and software to habitat modelling and rock-wallabies. In *Handbook for Applied Bayesian Analysis*, eds A. O'Hagan and M. West, Oxford University Press: OUP.

Low Choy, S., O'Leary, R., and Mengersen, K. (2009). Elicitation by design for ecology: using expert opinion to inform priors for Bayesian statistical models. *Ecology*, 90:265–277.

Martin, T. G., Kuhnert, P. M., Mengersen, K., and Possingham, H. P. (2005). The power of expert opinion in ecological models: A Bayesian approach examining the impact of livestock grazing on birds. *Ecological Applications*, 15(1):266–280.

Murray, J.V., A.W. Goldizen, R.A. O'Leary, C.A. McAlpine, H.P. Possingham and S. Low-Choy (*in press*), How useful is expert opinion for predicting the distribution of a species within and beyond the region of expertise? A case study using brush-tailed rock-wallabies (*Petrogale penicillata*), *J. Appl. Ecol.*

O'Hagan, A. (1998). Eliciting expert beliefs in substantial practical applications. The Statistician, 47:21–35.

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, R., Garthwaite, P., Jenkinson, D., Oakley, J., and Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts' Probabilities.* John Wiley & Sons, Inc.

O'Hagan, A. (1997). The ABLE story: Bayesian asset management in the water industry. In *The Practice of Bayesian Analysis*, eds French, S. and Smith, J. Q., pp173–198. London: Arnold.

O'Leary, R. A., Low Choy, S. J., Murray, J. V., Kynn, M., Denham, R., Martin, T. G., and Mengersen, K. (2008). Comparison of three expert elicitation methods for logistic regression on predicting the presence of the threatened brushtailed rock-wallaby (*Petrogale Penicillata*). *Environmetrics*, 19:1–20.

Spetzler, C. S. and Stäel von Holstein, C.-A. S. (1975). Probability encoding in decision analysis. *Management Science*, 22(3):340–358.

Uusitalo, L. (2007). Advantages and challenges of Bayesian networks in environmental modelling, *Ecological Modelling*, 203: 312–318.

Willems, A., Janssen, M., Verstegen, C., and Bedford, T. (2005). Expert quantification of uncertainties in a risk analysis for an infrastructure project. *J. Risk Research*, 8:3–17.