22nd National Conference of the Australian Society for Operations Research, Adelaide, Australia, 1–6 December 2013
www.asor.org.au/conferences/asor2013

Balancing the Validity and Viability of Bayesian Belief Networks for the Study of National Strategic Decisions

Andrew Coutts

Defence Science and Technology Organisation
Email:andrew.coutts@dsto.defence.gov.au

Abstract: National strategic decisions to employ military force are characterised by high levels of uncertainty regarding threat and actor motivation; political sensitivity regarding consequences; and a lack of clarity regarding influences and objectives. By understanding how these factors have interacted historically to impact on decision outcomes in different situations, strategic analysts are able to provide more informed analysis of current strategic options and therefore better advice to decision makers. Bayesian Belief Networks (BBNs) provide a powerful method to study past decisions as they capture both qualitative insights through concept mapping regarding decision maker rationale, as well as their likely impact on national objectives. Additionally, they readily incorporate subjective human input in both the construction of concept maps and through conditional probability assessments. However, there are a number of problems experienced in constructing BBNs to study these decisions. Of particular concern is the viability of the BBN structure for data capture. If a BBN is too complex a busy decision maker may be required to make an unacceptably high number of probability assessments. In order to make BBNs viable in this setting, an approach is required that simplifies the data capture from the point of view of the decision maker and manages the impact on model validity within a defendable and transparent process.

This paper takes a practitioners view on the application Bayesian Belief Networks (BBN) to study past national strategic decisions in order to inform future decision making. The intent is to propose a framework to structure BBN studies in an environment in which subject matter experts are dispersed, requiring time and money to meet them for data elicitation interviews, and where they are often busy and only able to provide a limited time for data elicitation. The focus of the framework is therefore on how to develop a BBN to study national strategic decisions in such a way so as to minimise the elicitation burden while managing the impact on the validity and integrity of the modelling. First, the advantages and data elicitation challenges of using BBN for studying political-strategic decisions are discussed. Next, various strategies and techniques from the literature regarding reducing the data elicitation burden are reviewed and illustrated, where appropriate, by examples from a recent data elicitation activity to develop a BBN on political-strategic decision making. Finally, selected strategies and insights from the recent data elicitation activity are combined into a proposed data elicitation framework that manages this trade-off between viability and validity.

Keywords: *Bayesian Belief Networks, National Strategic decisions, Data elicitation*

1. INTRODUCTION

The construction of analytical models is central to the application of operational analysis (OA) support to problematic decisions. The modelling process can help decision makers to: visualise the problem; comprehend large quantities of information; understand the influencing factors; investigate options and identify a preferred option; and finally explore 'what if' situations (NATO RTO, 2012, pp. 2-10). While the time constraints of supporting short notice decisions can limit some of the benefits of OA support, these limitations can be mitigated where more detailed modelling of similar historical decisions has previously occurred (Auerswald, 2004).

This modelling of relevant past decisions in slower time will ensure that a more comprehensive coverage of situations and associated influences is available to support immediate decisions. In particular, it will support more confident 'what if' analyses of future decision options at the cost of time and effort by the analyst and without the immediacy of a decision to support. This lack of immediacy further complicates the modelling process as it may lower stakeholder and decision maker engagement[1].

The benefits of detailed historical modelling are particularly relevant for national strategic decisions which may be subject to a wide range of structural, political, economic and cultural influences and may seek to address multiple goals towards improving the national security interests (Auerswald, 2004). The methods employed to model these past decisions must therefore be capable of addressing multiple influences on the decision process, under conditions of uncertainty, which in turn may contribute to multiple higher level goals. Additionally, such models must support a 'what if' analysis of the impact of changing the situation and associated influences on the higher level goals.

A review of the international relations and operations research literature indicates that, in general, three types of models have been employed for this type of strategic analysis: conceptual models that encode assumptions on human behaviour such as those based on rational choice and strategic culture (Lane and Ersson, 2005) models; qualitative-graphical models such as influence diagrams (Coyle, 2004); and structured and hierarchical models based on multi-attribute utility theory (Saaty et al., 1982).

While each of these models provide benefits to the modeller, all have deficiencies in supporting the 'what if' analysis required to inform future decision making. Such analysis requires both the ability to: account for multiple decision goals, measures and complex chains of influence on the decision process; and identify likely outcomes for different situations (Cain, 2001). Conceptual behavioural models provide a theoretical basis for explaining the primary motivation behind strategic decisions but do not attempt to account for multiple goals and secondary influences. For example the rational choice (Lane, et al., 2005) framework posits that strategic decision makers maximise self interest in their decisions. While this model attempts to account for the primary influence and rationale guiding the decision maker (if the premise is accepted), it does not on its own account for indirect and secondary influences and goals that may have a significant effect on the decision outcome and/or national interest.

Qualitative models based on influence diagrams, provide a powerful means to capture and display complex influences and goals impinging on the decision process but are less useful in identifying the most likely outcomes given a set of start conditions. Finally, multiple attribute utility theory (MAUT) based models can account for multiple goals and influences within their hierarchical structure and provide a means to determine a most likely outcome for different start situations, but in a highly constrained manner. For example, in an AHP hierarchy, lower order attributes can only contribute to one higher order attribute whereas in many real world situations these relationships will be more complex with some influences impacting on multiple goals. The Analytical Network Hierarchy (ANP) (Saaty, 2004), which represents a generalised version of AHP, attempts to remove some of the constraints of hierarchical decision networks and allow a lower order attribute to affect multiple higher level attributes. However, there are few examples in the literature discussing how the data elicitation burden may be managed to suit time constrained experts.

A fourth method that combines aspects of influence diagrams together with a means to capture relevant probability data from decision makers is Bayesian Belief networks (BBN) modelling. Bayesian Belief Networks (BBN) arguably provide a more useful analytical tool for strategic 'what if' analysis under uncertain conditions. They combine many of the qualitative benefits of influence diagram analysis with a probabilistic means to assess preferred options in different situations. Additionally they readily incorporate subjective input from decision makers both in the construction of the network and through conditional

[1] Based on the assumption that busy decision makers will more likely participate in model building if there is a clear and immediate benefit to themselves.

probability assessments. This process provides a useful structure for data elicitation interviews with experts[2] and the resulting network provides powerful means to represent and share the knowledge gained. For these reasons, this paper considers how BBNs may be employed to study past political strategic decisions despite challenges in data elicitation.

2. PROBLEMS WITH DATA ELICITATION

However BBNs also present challenges to a study of national strategic decisions. In particular there are implicit trade-offs between the viability and validity that can be achieved in eliciting data from busy decision makers. This is true both for the construction of the concept map and the probability assessments.

The greatest challenge in constructing BBNs for strategic decision making is that the majority, if not all of the data, must be obtained via subject matter experts, particularly from the busy decision makers themselves. As noted by one author, "it is a tedious job to perform evidence transmission even for simple Bayesian networks" (Jensen, 1996, p. 33). This tedium is exacerbated by the large number of concepts that may influence strategic decision making and the resulting large number of nodes requiring conditional probabilities. However, of greater concern is the exponential growth in the number of conditional probability assessments required as the number of parents, and the number of states that they can hold, increase. For example, a child node with three possible states that is dependent on three parents each with three possible states will require 54 conditional probability assessments. This not only leads to problems associated with the sheer number of assessments to be made, but also with the ability of a subject matter expert to provide coherent assessments for the level of detail required (Wisse et al., 2008). Such challenges to probability assessments would be extremely difficult to overcome given a substantial amount of time with a technically trained practitioner let alone a busy strategic decision maker. It is clear that the "elicitation task thus is considered a major obstacle in the use of BBNs" (Wisse, et al., 2008, p. 1) and that if Bayesian modelling is to be an effective tool to investigate strategic decision making, then strategies must be found to simplify the process of eliciting conditional probabilities.

3. STRATEGIES IN THE LITERATURE TO REDUCE DATA ELICITATION BURDEN

There are four general approaches to simplifying the problem of eliciting conditional probabilities in the literature that are relevant to the study of national strategic decisions, each of which trade-off some model accuracy in order to increase the viability of data elicitation. The first approach, referred to here as **simplifying structure**, focuses on restructuring the underlying belief network by reducing, where possible, the number of parents for each child. A reduction in the number of parents can dramatically reduce the number of conditional probabilities that need to be elicited. A second approach, referred to here as **simplifying elicitation**, focuses on employing techniques to make it easier for a subject matter expert to make probability estimates. A third approach, referred to here as **exploiting causal independence**, identifies instances of causal independence between parents and exploits this independence to reduce the number of conditional probabilities that need to be elicited. Finally, I propose a fourth strategy in which the data burden can be reduced through a process of **targeted elicitation.** In targeted elicitation, individual SMEs interviews are tailored to address specific sections of the BBN that are most relevant to their experience and knowledge. The premise of this paper is that very substantial reductions can be made in data elicitation without substantial reduction in model validity by applying these strategies in a structured way.

Example BBN Structure for National Strategic Decision Making

These strategies will be illustrated through an example from a recent BBN data elicitation activity conducted by the author. In this activity, a number of SME were interviewed to develop a BBN on post-cold war Australian national strategic decisions on the employment of military force. This activity relied on a preliminary BBN structure that was constructed prior to SME interviews.

The preliminary BBN model structure was developed through an analysis of relevant decision maker statements in textual sources made at the time of the decisions. This analysis identified decision maker objectives and influences for a range of similar decisions that were then structured into a set of hierarchical objective categories (Coutts, 2013). These influences and lower and higher level objectives were then linked

[2] Recent experience of the author suggests that where the SME accepts the logic of the BBN, the 'what if' data elicitation questions can provide a greater level of insight into the decision processes than an open ended question. In effect this suggests that the comments made by the SME in reaching their probability assessment may be just as useful a source of data to the wider research project.

via causal statements in the texts to structure the preliminary model using a method similar to Nadkarni and Shenoy (2004)[3].

A preliminary model was necessary for this study in order to help SMEs engage more quickly with the problem given the limited time available for interviews. The model provided an initial frame of reference[4] for the SME thereby allowing them to both grasp the intent of the modelling and identify areas for improvement more quickly. Within this paper, this recent BBN study is referred to as the **National Strategic Decision** (**NSD) Study**. A sub-model within the wider BBN, labelled as **government priorities**, is used in this paper to illustrate the application of strategies to reduce the data elicitation burden.

The original sub-model structure prior to the application of the elicitation reduction strategies is detailed in Figure 1. The oval nodes represented with double lines are deterministic nodes while the other nodes are referred to as chance nodes. Deterministic nodes will be set to a constant value to analyse a particular decision while chance nodes will contain conditional probability tables to determine possible output states.

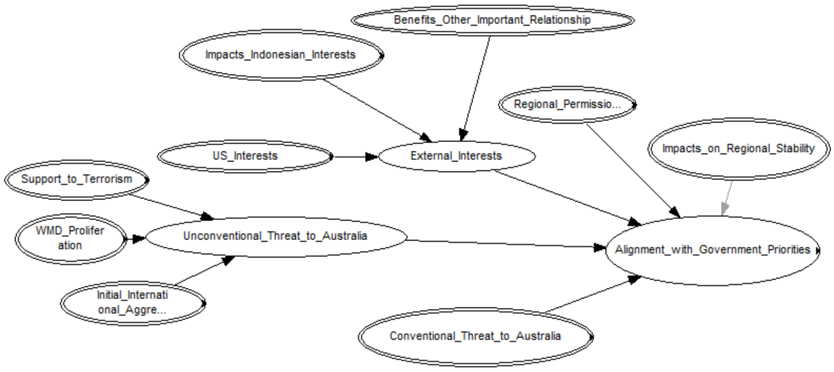

Figure 1 Original structure of the 'government priorities' sub-model

The definitions and possible states for the key variables for this sub-model are defined in Table 1. In the table, a '?' symbol indicates that probabilities will need to be elicited for these combinations of parent states and child state. The child states are detailed in the bottom three rows, titled low, medium and high.

Table 1 Nodal definitions for the 'government priorities' sub-model

Node Label	Possible States	Bayesian Definition
Alignment_with_Government_Priorities	High, Medium, Low	Indicates the current level of threat to Australia's direct security
Conventional_Threat_to_Australia	Yes, No	Indicates whether a conventional threat to Australia and/or Australian key interests (sea lanes) are involved.
Unconventional_Threat_to_Australia	High, Low	Indicats the level of unconventional threat to Australia or Australians involved with the proposed deployment
Impacts_on _Regional_Stability	Yes, No	Indicates whether the proposed deployment is addressing a problem with regional stability
External_Interests	High, Medium, Low	Indicates the level of External interests in favour of the deployment
Regional_Permissions	Yes, No	Indicates whether regional permissions exist for the proposed deployment

Table 2 Portion (1/9th) of the 'alignment with government priorities' node's CPT

Conventional Threat to Australia	*Low*					
Unconventional Threat to Australia	*Low*					
Impacts on Regional Security	*No*					
External Interests	*Low*		*Medium*		*High*	
Regional Permissions	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Low	?	?	?	?	?	?
Medium	?	?	?	?	?	?
High	?	?	?	?	?	?

As can be seen from Table 2, which displays a 1/9th of the entire conditional probability table, 108 probability elicitation questions are required to populate just this node, even after allowing that only two of

[3] The method used to construct the preliminary BBN from the influences and objectives identified in textual sources will be described in more detail in a separate paper.

[4] Essentially a decision frame as described in the problem structuring literature (von Winterfeldt et. al. 2009).

the three child state probabilities need to be elicited for each combination of parent states[5]. Given that this is just one node of a sub-model within a much larger BBN with multiple sub-models, it was clear to the author that direct elicitation of the required data from busy SME was not feasible.

3.1. Strategy 1 - Simplifying Structure

An obvious area to simplify elicitation is to simplify the underlying belief network that dictates the number of conditional probabilities that are required. In producing guidelines for employing Bayesian Belief Networks (BBN) to support water sector development, Cain (2001) points out that by simplifying the underlying belief network the analyst not only reduces the number of conditional probability assessments required but also increases the ability of those not involved in its construction to understand the network as a model. Cain recommends that the analyst should be clear about the aim of the network and in so doing, remove nodes and node states that are unlikely to be reached or unlikely to affect the outcome. Cain[6] also recommends "Divorcing" nodes as a way to reduce the number of parents contributing to a child and therefore simplifying the process of eliciting conditional probabilities. This process involves identifying where some contributing parents can be logically grouped together as contributing to an intermediate child. For example, if three of six parents to a child node titled 'Crop Yield' involve different aspects of employing fertiliser, an analyst might instead insert an intermediate child node titled 'Fertiliser Application' and direct the fertiliser related parents to this node instead. The intermediate child nodes can then be directed as contributing to the 'Crop Yield' child, thus reducing maximum number of contributing parent nodes and simplifying the data collection process[7]. Cain warns that care must be taken not to dilute " the impact of the interventions on the objectives" (Cain, 2001, p. 26). Druzdzel (1996) suggests one method to achieve this simplification is to substantiate variables wherever possible. That is, replace variables with observed values. This would reduce the complexity of a given belief network and therefore also simplify the process of eliciting probabilities.

Within the NSD study, the network structure in Figure 1 was reviewed by a strategic analyst SME prior to the main data elicitation activity. The key problem with the sub-model was that there were too many parents for the child node, **Alignment with Government Priorities**. On review, it was determined that the **Conventional Threat to Australia**, together with the **Unconventional Threat to Australia** nodes could be divorced from **Alignment with Government Priorities** by combining it with a new node titled **Threat to Australia**. Similarly, the node **Regional Permissions** was seen to be more applicable to a higher level sub-model within the BBN and could be removed as a parent for **Alignment with Government Priorities**. This resulted in the revised simplified structure in Figure 2 and substantially reduced CPT in Table 3. These changes reduced the maximum number of probability elicitation questions required to 36. Given that these changes are in line with the SME understanding of the wider influences, validity was maintained through this process and, arguably, the clarity of the model improved.

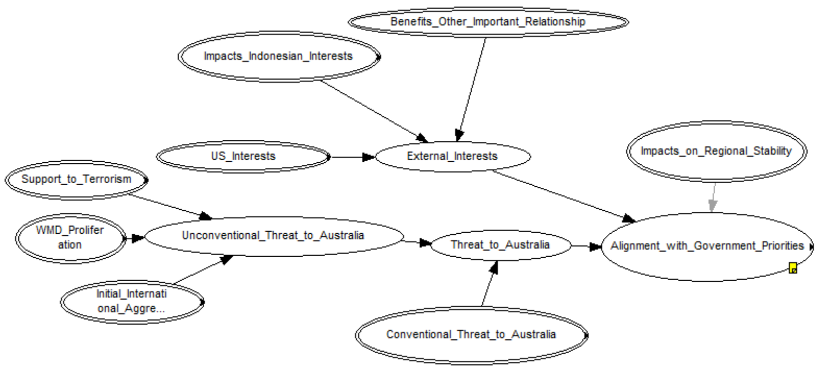

Figure 2 Revised 'Government Priorities' sub-model structure

[5] The third state can be determined as the sum of the probabilities of the three states must equal 1.
[6] see also (Jensen, 1996)
[7] See Figure 3.7 in (Cain, 2001, p. 25)

Table 3 Portion (1/3rd) of the revised conditional probability table 'alignment with government priorities' node

Threat to Australia	*Low*					
External Interests	*Low*		*Medium*		*High*	
Impacts on Regional Security	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Low	?	?	?	?	?	?
Medium	?	?	?	?	?	?
High	?	?	?	?	?	?

3.2. Strategy 2 - Simplifying Elicitation

A common way to reduce the elicitation burden placed on subject matter experts is to simplify the human process of estimating probabilities (Wisse, et al., 2008). Many authors observe that in general, human reasoning is rarely based on probabilities and have proposed various means to overcome this reality (Jensen, 1996; Kjaerulff, et al., 2008; Meyer and Booker, 2001; Von Winterfeldt and Edwards, 1986).

Given that small probability differences are unlikely to substantially impact on model outputs, some authors recommend qualitative approaches to elicit data. Qualitative approaches can put the subject matter expert at ease and enable them to focus on the "causal relations and relative preferences associated with decision problems"(Kjaerulff, et al., 2008, p. 164).

One qualitative approach is the probability wheel (Jensen, 1996; Kjaerulff, et al., 2008). This is essentially a pie chart divided into n partitions, where n is the number of possible states of the variable of interest[8]. The domain expert is invited to adjust the angular size of each partition to reflect the relative probability of a given set of inputs producing each variable state. Another method uses verbal statements that are mapped to numerical probabilities. For example, very unlikely might be mapped to a probability of 0.05[9]. A further method proposed involves a gambling analogy in which the expert judges how many 'winning' red balls must be included in a 100 ball lottery so that the chances of winning the lottery match the chances of a set of inputs producing a particular output state in a variable.

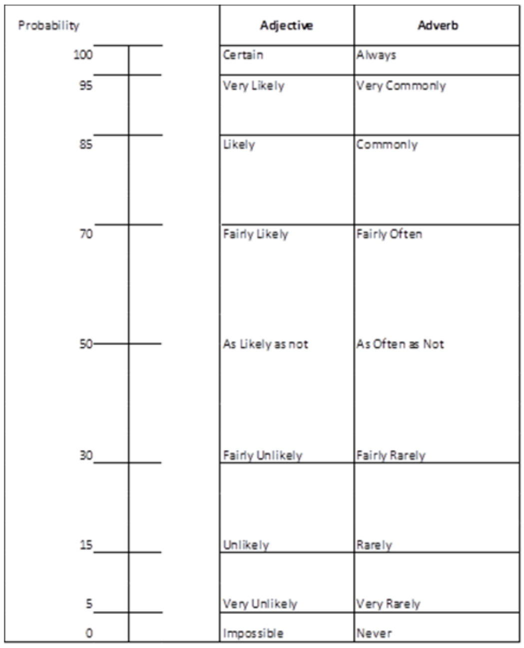

Figure 3. Data Elicitation Tool - Qualitative Statements Mapped to Probabilities - based on Kjaerulff et al (2008, p. 164) and Druzdzel (1996, p. 6).

In considering which of these methods might best support eliciting probabilities from decision makers, two factors are paramount: accuracy and feasibility. That is achieving an acceptable level of accuracy in the probability assessments while ensuring that the method adopted will be suitable and acceptable to the decision maker. While there is some evidence that assessments using the gambling analogy method are initially more realistic, thereby leading arguably to more accurate subjective assessments, respondents tend to quickly see past the analogy and revert to giving direct assessments (Von Winterfeldt, et al., 1986). Additionally, attempts to employ the pie chart method during the NSD study suggest that the method is not immediately intuitive to SME and slowed data elicitation. Overall, a mapping of qualitative statements to numerical probabilities, based on Kjaerulff et al (2008, p. 164) and Druzdzel (1996, p. 6), is preferred due to its simplicity and the expectation that the simpler approach will encourage more consistency in data collection (see Figure 3).

The experience of conducting the elicitation with SME using both methods during the NSD study tends to confirm that the mapping of qualitative statements to probability levels was the best method to engage SME, particularly policy advisers and analysts.

[8] See Figure 6.14 in (Kjaerulff, et al., 2008, p. 164)
[9] See Figure 6.15 in (Kjaerulff, et al., 2008, p. 164)

3.3. Strategy 3 - Exploiting Causal Independence

Having simplified the structure of the BBN and the elicitation process, the literature next recommends the identification and exploitation of causal independence. Causal independence describes a particular type of relationship between parents and their joint effect on children and, where such relationships exist, can dramatically reduce the elicitation burden (Jensen, 1996; Pearl, 1988; Zhang and Poole, 1996).

Zhang and Poole (1996, p. 307) state that parents "are causally independent [with respect to] their common effect if individual contributions from different [parents] are independent and the total influence on [a child] is a combination of the individual contributions"[10]. (Zagorecki, 2010) identifies the family of models with these properties as independence of causal influence (ICI) models.

A common means to exploit causal independence where it exists is through ICI models such as the noisy OR and nosy AND functions (Zhang, et al., 1996). Pearl (1988) indicated that the conditions for disjunctive interaction must be met in order to be able to assume a noisy OR relationship exists between causally independent parents that impact on a common child. Disjunctive interactions occur provided two conditions are met (Pearl, 1988, p. 185): Accountability and Exception Independence. Accountability requires that an event (consequence) be considered false if all of the causes that make the event true are false, as would be expected with an OR function. This condition can be described as the distinguished state of parent and child nodes, that is for a noisy OR model, the distinguished states of the parent nodes must be low[11] which corresponds to a distinguished state of the child as also low. A popular ICI model, which is essentially an extension of the noisy-OR model to non-binary parents and children, is the noisy-Max. Zagorecki (2010) introduced the concept of probabilistic independence of causal influence (PICI) models in which some of the assumptions required of ICI models are relaxed, allowing more flexible modelling of interdependencies. For example, models such as the noisy average (Zagorecki, 2010) allow distinguished states for parent and child other than the low state.

While a range of models were considered and tested in pilot elicitation activities, the most useful model for NSD study was found to be the noisy-Max. This was based on its prominence in the literature, its ability to replicate the interaction of parent and child nodes in line with the SME expectations in addition to the greatly reduced elicitation burden. Out of the 47 non-deterministic nodes in the NSD study BBN, thirteen were found suitable through interview to be treated as noisy-Max nodes. This included most of the nodes requiring large CPTs.

In applying the noisy-Max model[12] to the example in Figure 2, having confirmed through interview the assumptions for accountability and exception independence, resulted in a substantial reduction in the required number of elicitation questions as shown in Table 4. The table indicates the number of elicitation questions required in order for the noisy-Max model to calculate a conditional probability table (essentially Table 3). The distinguished states of the parents and child nodes are bolded and underlined. Elicitation is only required for each of the non-distinguished states of each parent where it is assumed the remaining parents are in their distinguished state. Noting that probabilities are required for only two out of the three child states, the noisy-Max model results in a reduction in the required elicitation questions from 30 down to 10.

Table 4 Complete elicitation table for 'alignment with government priorities' using noisy-Max model

Parent	Threat Level			External Interests			Impacts on Regional Security	
State	High	Medium	**Low**	High	Medium	**Low**	**No**	Yes
Low	?	?	**1**	?	?	**1**	**1**	?
Medium	?	?	**0**	?	?	**0**	**0**	?
High	?	?	**0**	?	?	**0**	**0**	?

3.4. Strategy 4 - Targeted Elicitation

Meyer and Booker (2001, p. 232) suggest that the structure of Bayesian networks offer a means to reduce the problem of eliciting conditional probabilities. I propose that by splitting data capture between experts and decision makers and further tailoring elicitation on their individual areas of expertise that substantial reductions to elicitation burden can be achieved. Using this approach in the NSD study, reductions were

[10] This is similar to the concept of non-modifying parents put forward by Cain (2001) in his guidelines for developing and employing Bayesian Belief networks.

[11] The noisy-OR and noisy-Max models require that the distinguished states of parent and child be a low or false state. This in turn requires an ordering of states from a low/inactive state to a high/active state (Zagorecki, 2010).

[12] The model was implemented in Genie, which is a software package developed by the Decision Systems Laboratory to develop and analyse Bayesian networks (Druzdzel, 1999).

achieved in the data elicitation burden placed on both decision makers and analysts. For example, the entire 'government priorities' sub-model (Figure 2) was reviewed only by selected analysts and policy advisers, while decision makers were able to focus on more value focused judgements in higher sub models which considered other elements of national interest. This essentially reduced the data elicitation burden (in terms of the number of nodes requiring elicitation) by up to half, while ensuring that data was elicited from SMEs within their areas of competence.

4. DISCUSSION AND CONCLUSION

The four strategies presented in the previous section and employed in the NSD study substantially reduced the data elicitation burden placed on SME and hence increased the viability of data collection and model building. However, the aim of this paper is to combine these strategies into a viable data elicitation framework that balances these elicitation benefits against any likely impact on model validity. Consequently, the following paragraphs review the validity and viability issues that were experienced in employing these strategies, leading to a suggested data elicitation framework.

The first strategy, **simplifying structure**, presents little challenge to validity on its own provided the analyst has a means to confirm the legitimacy of the proposed change. This can be managed by establishing a clear logic for the proposed change, and confirming the basis for the change where possible with strategic analysts and policy advisors. In the NSD study the rationale for combining conventional and unconventional threats was suggested in the primary source documents and tested through an interview with a relevant strategic analyst. This highlights the need for a more general discussion about the issues that inform the structure of the model, prior to an interview focused on probability elicitation. In the NSD study, wherever possible, a semi-structured interview was conducted with SMEs a minimum of a day prior to a structured probability elicitation interview. Not only did this facilitate the simplification of the BBN structure, it also demonstrated to the SME that their wider views on the topic were given appropriate consideration prior to structured probability elicitation interviews. The semi-structured interviews also appeared to generate a high level of engagement in the modelling process.

The second strategy, **simplifying elicitation**, aimed to minimise errors in the probability elicitation process by using methods that are conducive to human perception of probability. Hence the use of the data collection elicitation sheet (Figure 3) that combined qualitative statements combined with numerical probabilities. During the example study it was found that this sheet was preferred by SMEs rather than the pie chart. The qualitative statements appeared to make SMEs more comfortable with the elicitation process and resulted in many of them shifting to providing direct numerical probability assessments as their confidence increased.

The third strategy, **exploiting causal independence**, dramatically reduced data elicitation in the NSD study by making assumptions about the interaction between parent and child nodes. While these interactions were discussed during interviews and partially confirmed, it is reasonable to assume that any ICI/PICI model employed may introduce some error into the model representing differences between actual SME belief and the probabilities generated by the models. However, Zagorecki (2010) showed that in a study comparing human assessments of probability using a full CPT and noisy OR models (essentially the binary variable case for the noisy-Max) that the Noisy OR provided equivalent elicitation accuracy. Other findings suggest that noisy-Max models may provide a "surprisingly good fit" for up to half of CPT nodes in practical BBNs (Zagorecki, 2010, p. 99).

Finally **targeted elicitation** envisages that the data elicitation can be divided between SMEs with different competencies, thus reducing individual elicitations burdens. In practice, SME areas of strength were only apparent during interview. Consequently, two interviews were conducted with each SME wherever possible. The first interview was semi-structured and sought to capture qualitative insights and examples of the effects modelled in the BBN. This not only allowed the author to assess the strengths of the SME, and hence which parts of the BBN to elicit probabilities on, it also allowed a level of validation of the BBN structure and tended to give SMEs a sense of ownership of the BBN model which helped in the conduct of the more complicated data elicitation interview.

In summary, the following viable data elicitation framework, employing each of these four strategies, is suggested for studies of this kind. The starting point for this framework is that the analyst has previously: identified influences, decision maker objectives and cause-effect relationships from other sources; and converted these relationships into a preliminary BBN.

Proposed Data Elicitation Framework

1. **Identify SME**. Identify appropriate SMEs within two broad categories: strategic analysts/advisors and strategic decision makers. (In preparation for Strategy 4)
2. **Conduct Semi-Structured Interviews**. Conduct semi-structured interviews with SMEs using relevant objective categories within the BBN as interview prompts. The focus for the strategic analyst/advisor interviews should be on the detail of existing policy and practice, while the emphasis for decision makers should be on higher level strategic objectives. It is important to ensure that, where possible, there is an overlap between the two groups. (Strategy 4)
3. **Analyse Semi-Structured Interviews**. Review assumptions in the BBN structure in light of statements and examples provided in the semi-structured interviews and identify appropriate simplifications and, if necessary, additions. Document the SME statements leading to the change. (Strategy 1)
4. **Prepare Structured Interviews**.
    a. Based on the SME interviews, identify the relevant areas where individual SME have most experience and knowledge. Assign individual BBN nodes and whole sub-models for data elicitation to SMEs based on these strengths. As before, maintain overlap between decision makers and analysts/advisors wherever possible. (Strategy 4)
    b. Also based on the Semi-Structured interviews, identify nodes which are likely, in the opinion of the analyst, to satisfy the conditions for an ICI or PICI model and prepare appropriate data elicitation tables for these nodes. The assumptions for the ICI/PICI models will be tested with the SME during the structured interview and consequently the analyst should also prepare the associated full CPT table in case these assumptions are rejected. (Strategy 3)
5. **Conduct Structured Interviews**. Having provided a high level overview of the BBN modelling approach to the SME, employ the elicitation table recommended in Strategy 3 together with the elicitation tables prepared in step 4b, for each of the nodes selected for that SME. For each node which has been modelled as an ICI/PICI model, test the assumptions made. In the case of the noisy-Max model, this requires confirming with the analyst that they are comfortable that the output will be in the 'low' state when all of the parent nodes are in the low state and that the parents can act independently of each other. (Strategy 2, 3 and 4).

In applying this approach to the NSD study, the number of data elicitation questions required for individual nodes was reduced by up to 90% via the application of strategies 1 and 3. The individual SME elicitation burden was reduced a further 50% by the application of strategy 4 and, observationally, the difficulty in providing the probability assessments, and therefore the SME's fatigue was reduced through the application of strategy 3. However, maintaining model validity required that these strategies be carefully applied as detailed in the framework. Arguably, the semi-structured interviews were critical to the successful application of strategies 1 and 4. Without these interviews, the process of reducing the complexity of the network (strategy 1) and identifying the strengths of individual SMEs in order to allocate BBN nodes (strategy 4) would have relied solely on the analyst's 'best guess' and this would have reduced the validity of the modelling process. Conversely, the application of the strategies in the suggested two interview framework provides a reasonable approach to gain the benefits of the four strategies in reducing the elicitation burden while managing modelling validity.

Future work will focus on establishing a coherent framework to assess the internal and external validity of similar BBN models and assess their usefulness for informing future decision making. Such work will need to consider all aspects of the modelling process - structure, data, and model outputs - in order to establish a measure for the level of confidence that decision makers and strategic analysts can place in such models.

REFERENCES

Auerswald, D. P. (2004). Explaining wars of choice: an integrated decision model of NATO policy in Kosovo. *International Studies Quarterly, 48*(3), 631-662.

Cain, J. (2001). Planning improvements in natural resource management. Guidelines for using Bayesian networks to support the planning and management of development programmes in the water sector and beyond. Crowmarsh Gifford, Wallingford, Oxon, UK: Centre for Ecology and Hydrology.

Coutts, A. (2013). *A Structured Approach to Extract Strategic Objective Categories From Textual Sources*. Paper presented at the 22nd National Conference of the Australian Society of Operations Research Society (ASOR 2013), Adelaide, South Australia, Australia.

Coyle, G. (2004). *Practical strategy: structured tools and techniques*: Prentice Hall.

Druzdzel, M. J. (1996). Qualitiative verbal explanations in Bayesian belief networks. *AISB QUARTERLY*, 43-54.

Druzdzel, M. J. (1999). *SMILE: Structural Modeling, Inference, and Learning Engine and GeNIe: a development environment for graphical decision-theoretic models.* Paper presented at the AAAI/IAAI.

Jensen, F. V. (1996). *An introduction to Bayesian Networks*. London: UCL Press Limited.

Kjaerulff, U. B., and Madsen, A. L. (2008). *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*: Springer.

Lane, J. E., and Ersson, S. O. (2005). *Culture and politics: a comparative approach*: Ashgate Pub Co.

Meyer, M. A., and Booker, J. M. (2001). *Eliciting and Analyzing Expert Judgement: A Practical Guide*: ASASIAM Series on Statistics and Applied Probability.

Nadkarni, S., and Shenoy, P. P. (2004). A causal mapping approach to constructing Bayesian networks. *Decision Support Systems, 38*(2), 259-281.

NATO RTO, S.-. (2012). NATO Guide for Judgement-based Operational Analysis in Defence Decision Making. Analyst-oriented Volume: Code of best practice for 'Soft' Operational Analysis.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, California: Morgan Kaufmann Publishers, Inc.

Saaty, T. (2004). Fundamentals of the analytic network process — Dependence and feedback in decision-making with a single network. *Journal of Systems Science and Systems Engineering, 13*(2), 129-157. doi: 10.1007/s11518-006-0158-y

Saaty, T. L., Vargas, L. G., and Barzilay, A. (1982). High-Level Decisions: A lesson from the Iran hostage rescue Operation. *Decision Sciences, 13*(2), 185-206.

Von Winterfeldt, D., and Edwards, W. (1986). *Decision analysis and behavioral research* (Vol. 1): Cambridge University Press Cambridge.

Wisse, B., Van Gosliga, S., Van Elst, N., and Barros, A. (2008). *Relieving the elicitation burden of Bayesian Belief Networks.* Paper presented at the Sixth Bayesian Modelling Applications Workshop on UAI.

Zagorecki, A. (2010). *Local Probability Distributions in Bayesian Networks: Knowledge Elicitation and Inference.* Doctor of Philosophy, University of Pittsburgh, Pittsburgh.

Zhang, N. L., and Poole, D. (1996). Exploiting causal independence in Bayesian network inference. *Journal of Artificial Intelligence Research, 5*, 301-328.