Use of Bayesian networks in evaluation

This note introduces some of the basic knowledge to understand Bayesian networks in an evaluation setting. It draws on work conducted previously under a DFID-funded research project (Henderson and Burn, 2004)[1]. It starts with a brief discussion of some key features of Bayesian networks. It then goes on to give an example to demonstrate their potential. It concludes with some thoughts on the relevance to the evaluation setting.

Bayesian networks and what's interesting about them

Bayesian networks provide a graphical way of modelling situations in which causality plays a part but where our understanding is incomplete. They are in essence a "logic model" but one that is probabilistic, reflecting this uncertainty. Two key features are their ability to produce more realistic and interactive logic-models and to systematically incorporate subjective views about causality.

The logic models we normally encounter, of which the log-frame is the most common, are deterministic. The causal connections (or 'arrows') in those models imply "IF -THEN" causal certainty, (i.e. "IF that happens, THEN this *will* occur"), even if everyone accepts that is unrealistic. Some examples include moderating ASSUMPTIONS, but they are nevertheless silent about the actual moderation effect on the causal relationship and still imply a central, causal role for the elements of interest – typically, the programme intervention. That is, it is assumed that the programme is certainly necessary, even if not sufficient. It is a necessary element of a VfM analysis to test that assumption.

In contrast, the links in Bayesian networks are *not* absolute. Instead, they relate to the *likelihood* of a particular result for a given set of causal events or factors. Furthermore, they need not rely on simple, linear cause-and-effect relationships, but can accommodate multiple combinations of factors causing the same or similar outcomes.

Unpacking the causal connections, however, needs information about the relationship between events and outcomes. In situations where we do not typically have this information in an established evidence base[2], we are often forced to rely on subjective assessments. This can variously be the judgement of the evaluator herself based on all the evidence or the collective view of a group of expert stakeholders tasked with making the assessment. Either way subjectivity is an important element in the exercise. What is critical is that the approach and assessments are defensible. Bayesian networks enable us to incorporate these judgements explicitly and systematically into the underlying logic model and explore their implications for how interpret events.

An example

The scenario:

A donor provided electoral assistance in the run-up to national elections. This comprised capacity building support to the national electoral commission as well as support for a community awareness programme. A team of international observers declared the subsequent elections "free and fair". The

[1] Henderson, JS and R Burn (2004) Uptake Pathways: the potential of Bayesian belief networks to assist the management, monitoring and evaluation of development research. Agricultural Systems 79, 3-15

[2] In other fields (medical diagnosis or engineering), there may be a body of statistical evidence available: e.g. in 7 out of 10 cases, presence of symptoms 'A' and 'B' indicates disease 'X'; or, working from the other end, 60% of serious power failures are because of a faulty fuel pump, 20% are because of a blocked air intake and the remaining 20% are because of a number of different reasons.

donor is now commissioning an evaluation to assess effectiveness and determine what difference their assistance made.

For the purpose of this example, let's define 'free and fair' elections as ones with:

- an electorate who are willing (and able) to exercise their choice effectively at the ballot box;
- functional polling stations so that the electorate can vote; and
- a legal/policy framework that provides equality of votes among the electorate, so that all those entitled can vote and each vote is of broadly the same weight in determining the result.

These conditions themselves depend upon a number of other contributing factors, only a few of which were targeted for donor assistance.

Figure 1 sets out a basic logic model developed for the above scenario. It should be noted that the example is meant to be realistic but simple enough to aid understanding and that the arrows connecting different factors suggest causal relationships (or "plausible associations") but as yet these are undefined.

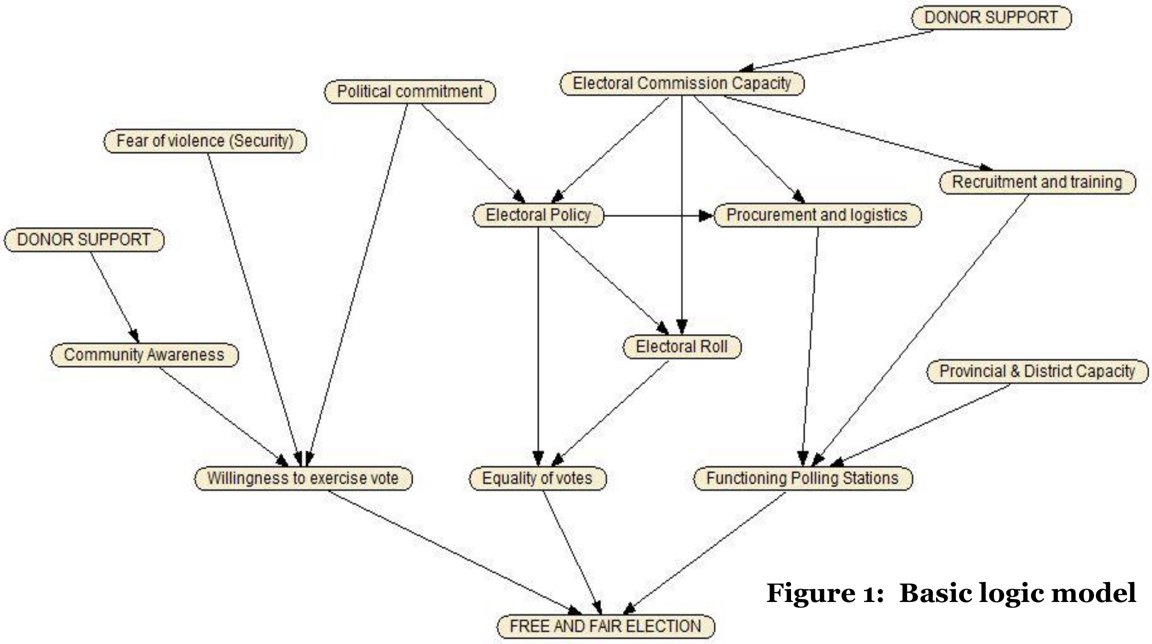

Figure 1: Basic logic model

Populating the model:

In order to establish, using this methodology, the importance of, say, Political Commitment[3] in the elections, we need to define the relationships between the factors linked in the model. For the purposes of illustration, we just focus on the left-hand side of the model, relating to Willingness to exercise vote. To further aid simplicity for this example, we define the possible outcome for each factor in binary terms –Community awareness is either 'High' or 'Low'; and similarly for the other factors as shown in Figure 2[4]. In practice, this could be relaxed to accommodate other possible outcomes if appropriate.

[3] For our purposes, this refers to the willingness of politicians to modify their behaviours, play by the rules, advocate real policies, commit to the system of democracy prescribed by the constitution, and so on.

[4] The 50:50 entry for the majority of factors in the model simply reflects complete uncertainty at this stage about possible outcomes.

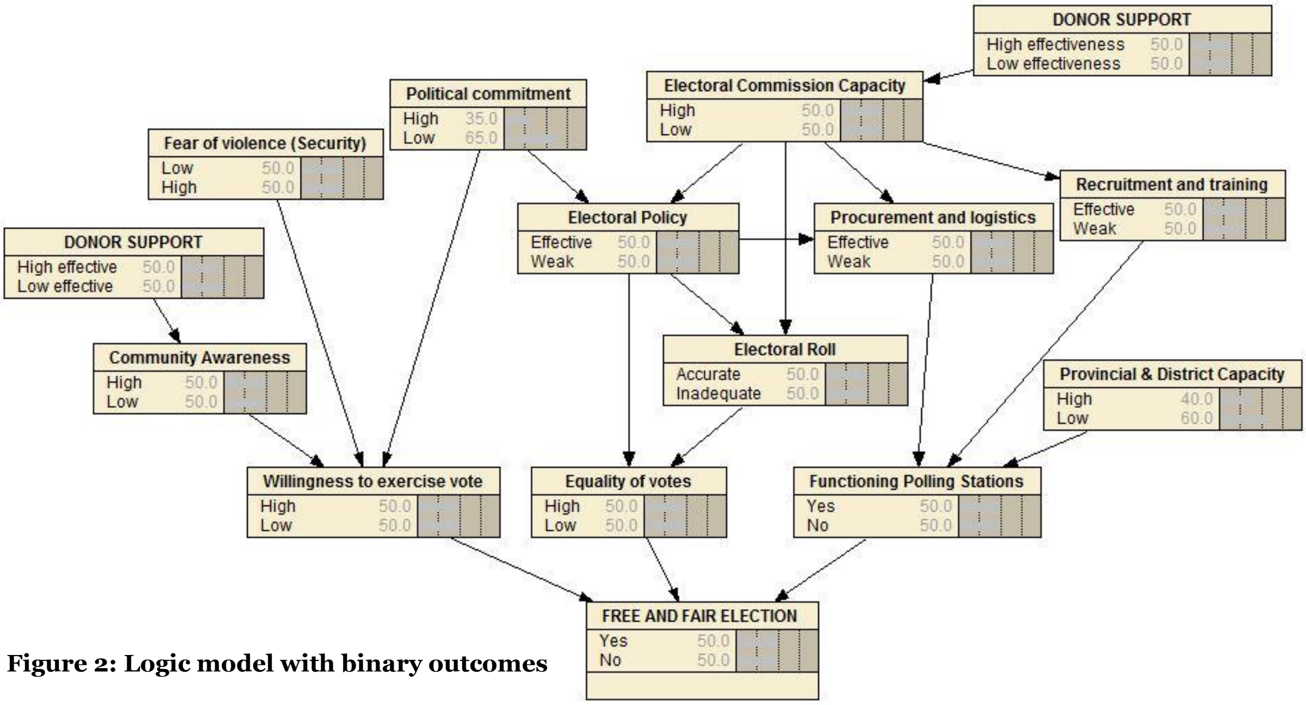

Figure 2: Logic model with binary outcomes

The model proposes that Willingness to exercise vote is a function of three 'parent' factors:

- Political Commitment: do politicians' behaviours and policies lend the necessary credibility and integrity to the process to encourage people to vote?
- Fear of violence: are people afraid to turn out?
- Community awareness: have people heard about the election and electoral process?

The donor has only supported the last of these three.

As this is an *ex post* exercise, the evaluator already knows that the election was declared "Free and Fair". The question to be examined then is "what difference did donor support make?" The first step towards this is to understand how the different possible outcomes among 'parent' factors affect the likely outcomes for 'child' factors, and ultimately the final outcome – the freeness and fairness of the election.

To do this, the evaluator has to define the *conditional probabilities* for each factor, based on the possible outcomes of its 'parent' factors. For factors with no 'parents', she defines the 'prior probabilities' – e.g. what's the likelihood there will be a High Fear of violence or Political Commitment or effective Donor Support and so on.

It is important to note that all the prior and conditional probabilities in this example are based on *subjective* assessment – i.e. people's beliefs and opinions. The methods of eliciting such information varies greatly according to context, looking at Table 1 can give the reader immediately several ideas about how this could be done[5], but the essential point is that the probabilities are all derived from people's judgement.

Table 1 presents the conditional probabilities table for the factor Willingness to exercise vote. It sets out the judgements about the chances (%) of a High or Low turnout, *given the different possible configurations of the 'parent' factors*. So for example, in case 1, when the outcomes of all three

[5] For example in Table 1, line 1 is 'the best of worlds' where the factors are all favourable. It is easy to imagine asking for the outcome probabilities from this point. Line 8 is 'the worst of all worlds' with the factors all set unfavourably. Again eliciting a response for this is easy to imagine. Other lines can be seen as moving one factor, or combinations of factors, from good to bad, or the reverse.

parent factors are assumed positive, the chance that Willingness to exercise vote will be 'High' is judged very likely (95%).

Table 1: Conditional probabilities for Willingness to exercise vote

	Parent factors	Possible outcomes	*Chances* **Willingness to exercise vote** *will be..?*	
			High (%)	*Low (%)*
1	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	**High** **Low** **High**	95	5
2	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	High Low **Low**	60	40
3	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	High **High** **High**	40	60
4	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	High High **Low**	30	70
5	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	**Low** **Low** **High**	70	30
6	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	Low Low **Low**	50	50
7	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	Low **High** **High**	35	65
8	**if...** Political Commitment = **and...**Fear of violence = **and...** Community awareness =	Low High **Low**	10	90

In contrast, in case 8, when the outcomes of all three parent factors are assumed negative, the chance of a high turnout is judged very unlikely (10%), though not impossible due perhaps to other motivational factors not explicitly listed in the model. In the middle are the cases for the other possible combinations of parent outcomes. So, for example, the chances expressed in case 3, suggests that in the view of the assessor(s), the positive effect on turnout of a high Political Commitment and Community awareness can be undermined, at least partially, by a high Fear of violence (less than 50% likelihood of a 'high' turnout).

Having defined the conditional probabilities for Willingness to exercise vote, the evaluator then does the same for the factors further up in the chain. Although this sounds complicated it should be noted that Table 1 (with three parent factors) is the most complicated that this model gets. Table 2 sets out the conditional probabilities for the factor Community awareness. With only one parent factor, the task is noticeably simpler.

Table 2: Conditional probabilities for Community awareness

	Parent factors	Possible outcomes	*Chances* **Community Awareness** *will be..?*	
			High (%)	*Low (%)*
1	**if...** Donor Support effectiveness =	**High**	85	15
2	**if...** Donor Support effectiveness =	**Low**	60	40

A reflection on table 2 suggests that the conditional probabilities indicate that effective donor support for raising awareness among communities *can* materially improve the likelihood of high levels of awareness, but it is likely not the only source of information available to communities. Even without donor support, there is still a better than evens chance of high levels of awareness among communities.

Compiling the network:

Figure 3 shows the compiled network for the example:

Figure 3: Compiled network

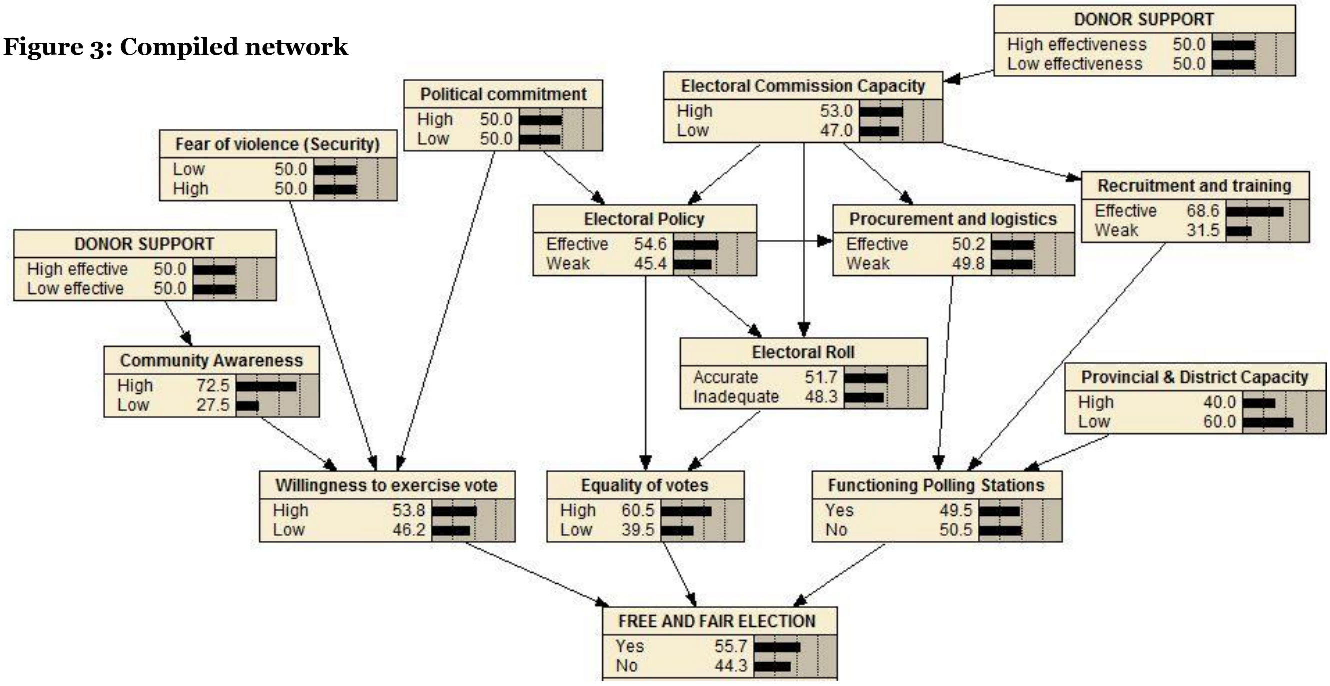

Compiling the network involves using an appropriate software[6] package to combine the prior probabilities and conditional probabilities using the basic (logical) rules of probability.

The compiled network suggests the most likely outcomes for the different factors, *given* the prior and the conditional probabilities that were entered. In this case, it suggests that the most likely (final) outcome is a Free and fair election but it is far from certain i.e. close to 50:50. This is informative but the network offers much more information by using it to answer 'what if?' type questions.

Interrogating the network:

The implications of the probabilities entered in the network can be explored both 'forwards' and 'backwards'. A 'backwards' view is where the fact that the election was deemed "free and fair" is already known. This information is entered directly into the network (called "entering evidence"), which in turn updates the model to show (a) the most likely combination of factors that led to this outcome (indicating 'contribution'); and (b) the relative certainty that those factors explain the outcome. For discussion purposes, Figure 4 shows this working for the left-hand side of our network, when the positive final outcome for the elections is entered:

[6] In this example, I use Netica by Norsys Software Corp: https://www.norsys.com/netica.html

Figure 4: Entering evidence #1

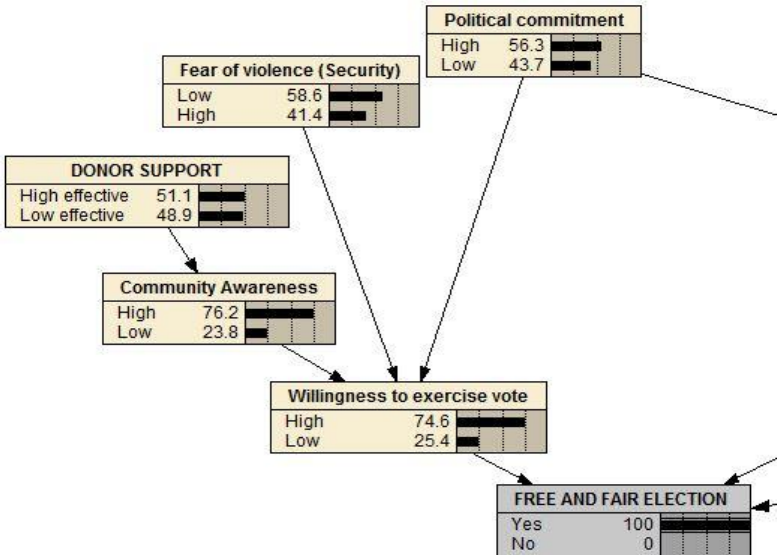

The figure indicates that, when the final outcome *is* positive, it can be expected that all contributing factors had positive outcomes (compared with the compiled network in Figure 3). However, importantly, for some the *relative* likelihood is greater. Compare the change in positive outcome for Fear of violence (from '50' to '58.6') with the change for Donor Support (from '50' to '51.1').

This greater relative likelihood also indicates a greater relative significance for the final outcome – the contribution element.

Application in a monitoring setting

While this paper has considered use in an evaluation setting, it is worth noting that the ability to update probabilities in all directions in a pathway is one of the most powerful features of the updating algorithm included in the software. Thus, for on-going monitoring purposes, evidence can be entered into the 'top' of the network, as and when the outcome of critical factors becomes known, and the effect on the likelihood of success at the end point can be observed (improving or worsening, and by what magnitude). Furthermore, stakeholders might revise their conditional probabilities, real-time, in the light of changing circumstance. Again, the ability to examine the effects, quantitatively, of such revisions on the prospects of success is .

Relevance to a monitoring and evaluation setting

In a number of respects, Bayesian networks appear a good fit with certain types of evaluation work: theory-based, participatory and predominantly qualitative. Although dealing in numbers and probabilities, they are unlike conventional 'black box' statistical enquiry. They engage directly with subjective data and are transparent and easily and intuitively explored. It is better to think of the approach as a tool to explore beliefs, evidence and their logical implications, than as a means to 'prove' something in some absolute sense. They therefore are also useful in producing the balanced judgements required for evaluation in a VfM context.

Rather than representing an entirely new method, Bayesian networks offer the potential to augment and strengthen existing theory-based approaches. For example, accounting for influencing factors and alternative explanations is a key element in Contribution Analysis but "[i]n spite of conceptual advancement… little attention has been awarded to how to examine [these]."[7] Bayesian Networks can be used to explore just part of the theory of change we are examining, used 'privately' to structure and inform the evaluator's own understanding or 'publicly' in a participatory process to stimulate and challenge collective views.

[7] Lemire, S. *et al* (2012) Making contribution analysis work: A practical framework for handling influencing factors and alternative explanations. *Evaluation* vol. 18 no. 3 294-309, July

Bayesian Networks are also relatively user-friendly, practicable and, with the right treatment, can present in an intuitive and graphical way the 'story' behind a finding. In addition, although they can look complex at first glance, only a relatively limited number of conditional probabilities are required to populate the model. In our example, the evaluator specified a total of 51 probabilities for the 15 factors in the logic model. This seems reasonable, given that the complete specification 'contained' in our model totals 32,767 joint probabilities[8].

As indicated at the outset, the potential of Bayesian networks for evaluation work is still to be fully appreciated but their use has been well proven in other very similar fields.

[8] For the simplest case of n random variables with binary outcomes, the complete distribution is specified by 2^n-1 joint probabilities.