

Aldo Benini

Bayesian belief networks- Their use in humanitarian scenarios

An invitation to explorers

15 July 2018, Version 1.0

© Aldo Benini 2018

Email: abenini@starpower.net

Web: www.aldo-benini.org

Contents

- Acknowledgement 4
- Summary 5
- Introduction..... 6
 - The setting for our experiment 7
 - Belief networks for scenario development 7
 - A simple model 8
- Elements of the network..... 10
 - Nodes and arcs..... 10
 - Layers of variables..... 10
 - Causal assumptions..... 11
 - Probabilities 12
 - Conditional probability tables (CPTs)..... 12
 - Nodes with states defined by an equation..... 14
- Scenarios 17
 - Forward and backward-looking 17
 - Updated networks 17
 - Outcome synopsis 21
 - A backward-looking scenario 22
 - Starting point 22
 - Conditions for reduced displacement 23
 - More specific “What .. if?” 25
- Belief networks in humanitarian scenario development 27
- Software 29
- Appendix 29
 - Calculating marginal probabilities from the Conditional Probability Table..... 29
 - The briefing note: “Nigeria – Displacement in Plateau State” 30
- References..... 31

Tables and Figures

Table 1 Potential objectives for belief networks.....	6
Table 2: A Conditional Probability Table (CPT).....	13
Table 3: Equation for the node "Potential IDPs, Oct. 2018"	15
Table 4: Outcomes of three forward-looking scenarios	21
Table 5: Calculating joint and marginal probabilities	29
Table 6: Updating joint and marginal probabilities	30
Figure 1: Belief network with probability bar charts (segment)	5
Figure 2: Simulation set-up and initial probability distributions.....	9
Figure 3: Updating beliefs with new evidence	14
Figure 4: A node with states defined by an equation.....	14
Figure 5: A network segment with one updated node.....	16
Figure 6: Network for scenario #1	18
Figure 7: Network for scenario #2	19
Figure 8: Network for scenario #3	20
Figure 9: Starting point of the backward-looking scenario exercise	23
Figure 10: Backward-looking scenario	24
Figure 11: A more specific "What-if" belief network	26

Acknowledgement

The models described in this note were created using the GeNIe Modeler, available from BayesFusion, LLC, <http://www.bayesfusion.com/> .

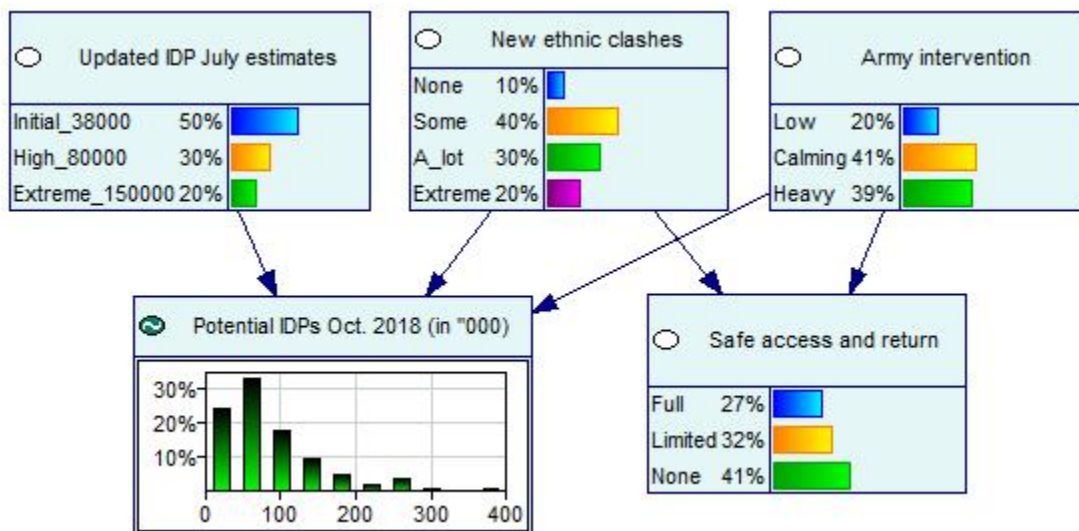
Alex Odlum, Senior Analyst, ACAPS, and unnamed analysts at ACAPS (<https://www.acaps.org/>) and START Network (<https://startnetwork.org/>) authored the briefing note on Displacement in Plateau State, Nigeria, that inspired our demonstration models.

Summary

This is an invitation for humanitarian data analysts and others – assessment, policy and advocacy specialists, response planners and grant writers - to enhance the reach and quality of scenarios by means of so-called Bayesian belief networks. Belief networks are a powerful technique for structuring scenarios in a qualitative as well as quantitative approach. Modern software, with elegant graphical user interfaces, makes for rapid learning, convenient drafting, effortless calculation and compelling presentation in workshops, reports and Web pages.

In recent years, scenario development in humanitarian analysis has grown. Until now, however, the community has hardly ever tried out belief networks, in contrast to the natural disaster and ecological communities. This note offers a small demonstration. We build a simple belief network using information currently (mid-July 2018) available on a recent violent crisis in Nigeria. We produce and discuss several possible scenarios for the next three months, computing probabilities of two humanitarian outcomes.

Figure 1: Belief network with probability bar charts (segment)



We conclude with reflections on the contributions of belief networks to humanitarian scenario building and elsewhere. While much speaks for this technique, the growth of competence, the uses in workshops and the interpretation of graphs and statistics need to be fostered cautiously, with consideration for the real-world complexity and for the doubts that stakeholders may harbor about quantitative approaches.

This note is in its first draft. It needs to be revised, possibly by several authors, in order to connect to progress in humanitarian scenario methodologies, expert judgment and workshop didactics.

Introduction

This is an invitation for humanitarian data analysts and others – assessment, policy and advocacy specialists, response planners and grant writers - to enhance the reach and quality of scenarios by means of so-called Bayesian belief networks. Those positions, and plausibly others in their stakeholder audiences, are exposed to, and frequently build, scenarios. Familiarity, and eventually active competency, with belief networks will take building and critiquing scenarios to a higher level.

What are Bayesian belief networks (BNBs)? Založnik et al (2018:7) offer an excellent nutshell definition:

BBNs are probabilistic representations of the direct and indirect influences between a set of variables. Developed in the 1980s, they have found numerous applications in fields as diverse as artificial intelligence, medical diagnostics, forensic science, computational biology, and environmental sciences. What these applications all have in common is that they deal with complex systems that incorporate at least some degree of uncertainty and BBNs are extremely well suited for this purpose. That, and the fact that they neatly integrate qualitative and quantitative data, make them particularly suited for tackling complex multidisciplinary questions.

A search on Google Scholar returns over 13,000 document references. This number drops to 80 when the term “humanitarian” is added; many of the references, however, have nothing to do with the humanitarian domain as we understand it (with rare exceptions from among supply chain researchers). Among neighboring professional communities, two have made greater use of belief networks: natural disaster and ecological researchers¹.

Marcot (2017:7), an ecological modeler, lists these potential objectives for Bayesian belief networks:

Table 1 Potential objectives for belief networks

Objective	Description
Prediction	Determine possible future outcomes based on initial conditions
Forecast	Determine the most likely future outcome based on initial conditions
Projection	Determine possible future outcomes based on changing future conditions
Scenario planning	Peg the corners of the implications of hypothetical situations
Represent knowledge	Synthesize what we think we know
Identify uncertainties and key data gaps	Identify factors or interactions with the greatest influence on outcomes; sensitivity analysis
Diagnosis	Determine potential causes of a known or specified condition or outcome
Mitigation	Identify alternative conditions that could lead to a desired outcome
Group involvement	Aid individual or collaborative decision-making; engage stakeholders

While for this author scenario planning as an objective is distinct from others, humanitarians may hold notions in which scenarios shade seamlessly into several of the listed objectives and uses. There is no harm in this; the point to stress is the versatility of the instrument.

¹ For a good example combining belief networks, expert judgment and decision support (volcanic eruption risk in Greece), see Aspinall and Woo (2014), available at <https://appliedvolc.springeropen.com/track/pdf/10.1186/s13617-014-0012-8>.

In recent years, scenario development in humanitarian analysis has grown, as seen in guidance (ACAPS 2012) and application documents (for a representative example, see: Nigeria INGO Forum and ACAPS 2016). One can wait for exemplary studies by other disciplines to percolate into the humanitarian community. Else, one can try to accelerate diffusion and adoption by promoting experiments that speak the language of humanitarians. This is what this note intends.

Technical introductions to Bayesian belief networks are plentiful, but most of them lead on a steep learning curve that may soon discourage the novice. Charniak (1991) must have had a reason when he gave his short article for the “probabilistically unsophisticated” the title “Bayesian Networks without Tears”². Some minimal understanding of probability concepts will sooner or later be imperative, but one of the blessings of this current era is that the learner can acquire it, actively and pleasantly, from the manuals of user-friendly and powerful belief network software while creating his/her own projects.

The remainder of this note does just that, leading the reader through an experiment, a journey into a simple network created with minimal information about a recent violent crisis in Nigeria. Only at the end, in hopes that we made a decent case for this analytic technique, do we revert to the question of its place in humanitarian scenario development. Even if we succeed in creating interest, laying out a path for learners as well as developing and testing templates of practical value – all that will take time.

The setting for our experiment

Armed clashes between migratory herdsman and sedentary farmers are increasingly common in the so-called Middle Belt of Nigeria³. On June 23, 2018, Fulani herdsman attacked local farmers in Plateau State, triggering the displacement of at least 38,000 persons. On July 10, the START Network and ACAPS released a briefing note “Nigeria – Displacement in Plateau State”, collating the available information (START Network and ACAPS 2018: see appendix). In a paragraph on the anticipated scope and scale of the emergency, the authors highlight worsening conditions in shelter, WASH facilities and food security. The rainy season increases the risk of disease outbreaks in the 31 IDP camps known to exist so far.

There is vast uncertainty in the information extant. IDP estimates as of July 9 ranged from 38,000 to 144,000. As for food security, estimates are available only for the entire state of Plateau, with 1,517,886 people in Integrated Food Security Phase 1 (Minimal), 804,412 people in IPC Phase 2 (Stressed) and 191,957 people in IPC Phase 3 (Crisis) (for definitions, see: IPC Global Partners 2012). Nonetheless, the authors classify the humanitarian impact as “moderate”, and the currently felt need for international assistance as somewhere between moderate and “major”.

Belief networks for scenario development

Although the current information on this crisis is poor, forward-looking scenarios are important in deliberating needs and opportunities for humanitarian activities. The information and the uncertainty

² Available at <http://www.aaai.org/ojs/index.php/aimagazine/article/download/918/836>.

³ See the recent piece in The Economist, “Mayhem in the Middle Belt. Fighting between Nigerian farmers and herders is getting worse. An age-old rivalry has grown deadlier, thanks to climate change, bad government and plentiful guns”, June 7, 2018, <https://www.economist.com/middle-east-and-africa/2018/06/07/fighting-between-nigerian-farmers-and-herders-is-getting-worse>.

that surrounds it need to be structured for rational analysis. Here we make a modest attempt to model the uncertainty and create scenarios quantifying impacts in the Plateau State crisis area. We choose a short time horizon, the next three months, until October 2018.

To do so, we draw a causal nexus from current conditions and near-future events to outcomes using a simple Bayesian belief network. In its initial configuration, the network assumes that all events are probabilistic. We populate it with probabilities that function as simulated expert judgments, but de facto they reflect nothing but common-sense assumptions made up from a remote reading of the situation.

The probabilities are of two kinds: First, those of initial conditions and external trigger events (or multiple events). Second, the probabilities of intermediate states and outcomes, given combinations of causes. As time progresses, more information will become available. Notably, assessments will produce more accurate estimates of displaced persons. As conditions and events are closely observed, either on the ground or through updated expert opinions, the belief network will produce more precise, although not necessarily more accurate, predictions. While waiting for those updates, we can recast the probability distributions of the outcome variables by simulating conditions and triggers.

The purpose of the exercise is to give a demonstration of Bayesian belief networks. Of interest are their utility for humanitarian analysis and scenario building, and the ease and comfort with which modern software lets us construct and compute them. This note may or may not succeed in the first aspiration. To fulfill the second, it takes active doing as a modeler at the computer screen. We cannot help the awe and joy when we watch a scenario update itself graphically within a fraction of a second when we click on a trigger event or change an element of a probability table.

A simple model


Belief networks have been built for complex investigations. Reportedly, a diagnostic model for Diesel locomotives boasts over 2,000 variables and 6,000 parameters. Bravo and safe journey!

Our model courts the other extreme. It limits the complexity to a didactic (and policy-wise unrealistic) small set of cause-effect assumptions. The model produces probabilities for two outcomes only: the number of IDPs in October, and the number of camps that experience disease outbreaks between now and October. We do not, in this version, model the number of food-insecure persons.

An easy way for the reader to familiarize with the setup is to browse back and forth between the network diagram on the next page and the explanations starting on the following page.

Scenarios of humanitarian outcome in Plateau State, Nigeria, by October 2018

Figure 2: Simulation set-up and initial probability distributions



Elements of the network

Nodes and arcs

The network diagram has ten boxes each reserved for a variable. In network parlance, these are known as “nodes”. Variables thought to exert a causal influence on other variables are linked by arrows that point from cause to effect. An arrow stands for an “arc”, a term from Graph Theory denoting the link between a pair of nodes. If our count is correct, we have drawn 14 arrows. Note that there are no causal cycles – Nowhere is there a sequence of arcs that leads back to the first variable. Note further that, theoretically, among 10 variables, and avoiding cycles, a maximum of $9 + 8 + \dots + 2 + 1 = 9 * 5 = 45$ arcs can be mapped. In other words, the model assumes that only about a fourth of the possible *direct* causal relationships matter. Assumptions of causal independence are as much a critical element of network models as the specific causal relationships. Technically, they are a key precondition that Bayesian algorithms can work.

Before we look inside the boxes, we go through the layers of variables. Also, we describe, qualitatively, the causal assumptions shaping the intermediate and outcome variables before we illustrate how their probability distributions are calculated.

Layers of variables

There appear to be four layers of variables. The graphic arrangement suggests so. Basically, however, there are three:

- Starting nodes are nodes with no incoming arrows. We have five of them.
 - They include the row of variables expressing initial conditions in July and trigger events expected between then and October in the top row.
 - They also include the variables “International assistance”. This decision variable is external; it is placed at the bottom for space reasons and to underline its function⁴.

At this point, it is not yet known in which states these five variables are. The states will be increasingly known over the next few months. Meanwhile we can simulate knowledge of their states.

- Intermediate states: The layer below is marked as “intermediate states”. These variables cannot be directly observed; their states are *inferred*, as probabilities, from their parent variables. They are needed because conceptually they are believed to be aggregators and transmitters of causal influences. They determine the probabilities of the outcome variables being in certain states.

⁴ The designation as a “decision variable” is motivated by the humanitarian perspective; it is not the usual sense of classic belief network terminology. This holds that a decision node enumerates decision options without probabilities. Such a node feeds into a “utility node” that computes (monetary or other) pay-offs. These depend on the decision and the probabilistic outcomes captured in other nodes. In this note, we avoid discussing these features.

- Outcomes: The variables of greatest interest are the two outcomes marked with red title banners⁵. Naturally, these outcomes are observable – they will be observed between now and October.

Causal assumptions

We made the following assumptions behind the five intermediate and outcome variables:

- Potential IDPs by October: The number of persons who could be IDPs by October is modeled as the product of the number of IDPs in July and factors depending on the levels of further ethnic clashes and of the Nigerian army response. “No new clashes” or just “some” make the potential shrink, “a lot” or “extremely many” make it explode. Similarly, if the army interposes itself between Fulanis and farmers in a calming manner, it reduces the potential for IDPs. A heavy-handed response, such as by aerial bombardments, increases it.
- Safe access and return: In a more realistic model, this variable would be divided into two. It encompasses the safety of people and relief agencies to work together (“humanitarian access”) as well as the safety of the displaced to return to their places of residence and work. It is entirely determined by the levels of new clashes and army response.
- “Humanitarian efficacy” is a variable at a loss for a good name. It packs the causal power of “safe access and return” and of the “rainy season” before their influence is passed on to the outcome variables. “No safe access and return” maps to “very poor efficacy”, regardless of the level of precipitation. Normal rainfall decreases efficacy slightly, floods do so majorly. Humanitarian efficacy is a moderator variable on the effects on the outcome variables with a subtle nuance: One might expect that it moderate only the effects of the “significant, but late” and “significant and timely” levels of the international assistance. However, there are Nigerian relief organizations working in the affected area (little is known about them). They too are subject to efficacy limits. The arcs from efficacy, therefore, go directly into the outcome nodes.
- The expected number of IDPs is the product of potential IDPs and mitigating factors that depend of the levels of humanitarian efficacy and international assistance. If efficacy is very poor, and the international response is insignificant, the expected IDPs equal the potential ones. If efficacy and/or international assistance improve, the expected number decreases. At these better levels, conditions for return are better, international agencies contribute more resources, and both national and international actors work more effectively in assisting IDP.
- Camps with disease outbreaks: More potential IDPs go hand in hand with more camps. The risk of outbreaks is reduced by higher efficacy and international assistance in the same logic that works for IDPs. In addition, drier weather lowers the risk via more salubrious camp conditions.


⁵ However, note this inconsistency, if not of substance, at least of style: Both “Expected IDPs” and “camps with outbreaks” receive causal arcs from two kinds of variables. Both are influenced by the potential for displacement and by something called “humanitarian efficacy”. These are unobservable intermediaries. In addition, the extent and timeliness of international assistance will have a mitigating effect on displacement and disease. Moreover, there is a direct arrow from “rainy season” to “camps with outbreaks”. The causal reasoning about these links from observables to outcomes is intuitive, but the model would be more consistent if the effects of weather and international assistance were first captured in intermediate variables. Simplicity and space limits forces this compromise.

This indicator may be acceptable for this demonstration; in real-world epidemic monitoring, the number of camps experiencing some kind of outbreak of undefined nature, duration and incidence would not be useful.


Probabilities

Three of the ten variables – potential and expected IDPs as well as camps with disease outbreaks – are count variables, conspicuous for their vertical histograms. By the same logic, the estimated IDPs in July should be a count variable. Artificially, we reduced it to three categories, each given a numeric interpretation (38,000 the initial; 80,000 a middling; 150,000 the highest estimate so far – ranges would be more intuitive, 30,000 – 60,000 etc.). The purpose is to demonstrate that continuous and count variables can (and in Bayesian belief networks often should) be treated as sets of ordered ranges, i.e. as categorical variables. The categorical character of the other variables is more or less straightforward. Their ordinal interpretation is less clear for some. From the humanitarian viewpoint, a heavy armed response is not preferable to calming interposition. A dry rainy season may be better for humanitarian logistics and disease control than normal precipitation levels, but worse for farmers and hence for food security. A more refined model would take these opposite effects into account.

All start point variables are categorical, recognized by horizontal bar graphs inside the node boxes. At the outset, the bars (and the percentages to their left) express the probabilities that experts have given each state to be the current, as in this image of the node for new ethnic clashes:



Once the state of the variable has been observed, or we pretend to know it for a simulation, the bar chart changes. The observed or simulated state becomes certain (100% probability); the others are excluded (0%).




In network lingo, the observed or simulated states are referred to as the “evidence”.

What about the probabilities estimated for the states of the intermediate variables? We distinguish between categorical and continuous/count variables.

Conditional probability tables (CPTs)

We exemplify the categorical case with “Safe access and return”. This node is a child of “New ethnic clashes” and “Army intervention”; these are its parents, in the network lingo⁶. Before any evidence is entered into the model, the unobserved state probabilities are:



How do we know that, before evidence, the probability that IDPs and humanitarians enjoy full access and safe return is 26 percent, etc.? These figures are the results of expert judgment too, but in an indirect way. Experts, by judgment or relying on prior evidence elsewhere,

⁶ In turn, “Safe access and return” has one child node, “Humanitarian efficacy”, but this is not important for the explanation of the probabilities of its states.

determine the probabilities of "Full", "Limited" and "None" from so-called conditional probabilities. These are the probabilities that this variable is in a certain state given the states of its parent nodes (new clashes and army intervention). The judgments or estimates come in so-called "conditional probability tables". This is the CPT for "Safe access and return":


Table 2: A Conditional Probability Table (CPT)

Determined by:		Determining probabilities of:			
Army intervention	New ethnic clashes	Safe access and return			Sum of p
		Full	Limited	None	
Low	None	0.8000	0.1000	0.1000	1.0000
Low	Some	0.4000	0.5000	0.1000	1.0000
Low	A lot	0.0500	0.5000	0.4500	1.0000
Low	Extreme	0.0010	0.3000	0.6990	1.0000
Calming	None	0.9500	0.0030	0.0470	1.0000
Calming	Some	0.7000	0.2000	0.1000	1.0000
Calming	A lot	0.3000	0.5000	0.2000	1.0000
Calming	Extreme	0.0010	0.2000	0.7990	1.0000
Heavy	None	0.2000	0.5000	0.3000	1.0000
Heavy	Some	0.1000	0.4000	0.5000	1.0000
Heavy	A lot	0.0010	0.3000	0.6990	1.0000
Heavy	Extreme	0.0001	0.1000	0.8999	1.0000

Note: Probabilities expressed as fractions

The unconditional probability of each state of "Safe access and return" that we see in the bar chart is the sum of these conditional probabilities multiplied by the unconditional probabilities of the parents. Interested readers find the calculations in a spreadsheet in the appendix. The result is what we see in the node bar chart: a 26 percent probability of full access and return, 31 percent of limited, etc. As we get to know more about further ethnic clashes and the way the army responds, we can enter this information as evidence, with a click on one bar in each of the trigger event nodes, e.g. "A lot of clashes" and a "Heavy intervention". The Bayesian algorithm immediately updates the probabilities, using the same CPT.

Figure 3: Updating beliefs with new evidence

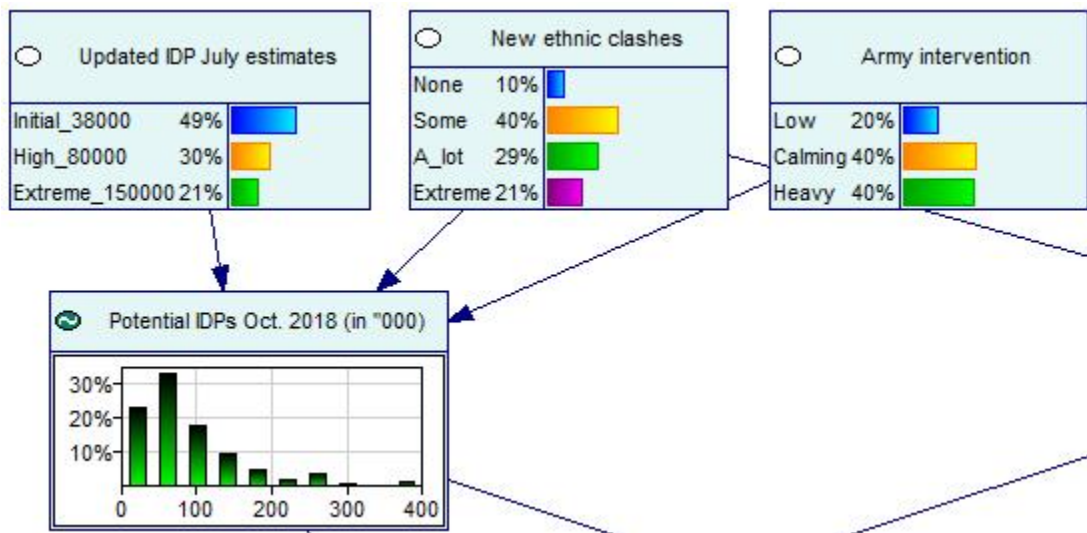


As one would expect, given the new information, the chance of full access and return is entirely gone, and “no safe access or return” await as the most probable outcome between now and October.

Nodes with states defined by an equation

Equations are a plausible way to model the probability distribution of a continuous or count variable. We exemplify with the equation for the potential IDPs by October. This node has three parents:

Figure 4: A node with states defined by an equation



The exact median of potential IDPs is not available in the histogram; by adding the values on the y-axis for the leftmost two bars (circa 20 + 30 = 50 percent probability), we guess that the value must be somewhere near 50,000. We cannot get the median directly; instead, the Bayesian application tells us to expect the arithmetic mean of this distribution to be near 93,000.

Mean:	StdDev:	Skewness:	Kurtosis:	Minimum:	Maximum:
93.0811	72.5595	1.85635	6.95092	18.24	390

This is an uninformative estimate because no evidence has as yet been entered into the parent nodes. The range goes all the way from 18,000 to 390,000 IDPs. Clearly, we want better.

But first, how is this "Potential IDPs" distribution computed?

It is the outcome of an equation, weighted by the parent node probabilities. We write the unweighted part in tabular form

Table 3: Equation for the node "Potential IDPs, Oct. 2018"

Potential IDPs by Oct. =

IDPs_in_July estimate	Choose one of the "IDPs in July" estimates: 38,000; 80,000; 150,000
-----------------------	------------------------------------------------------------------------

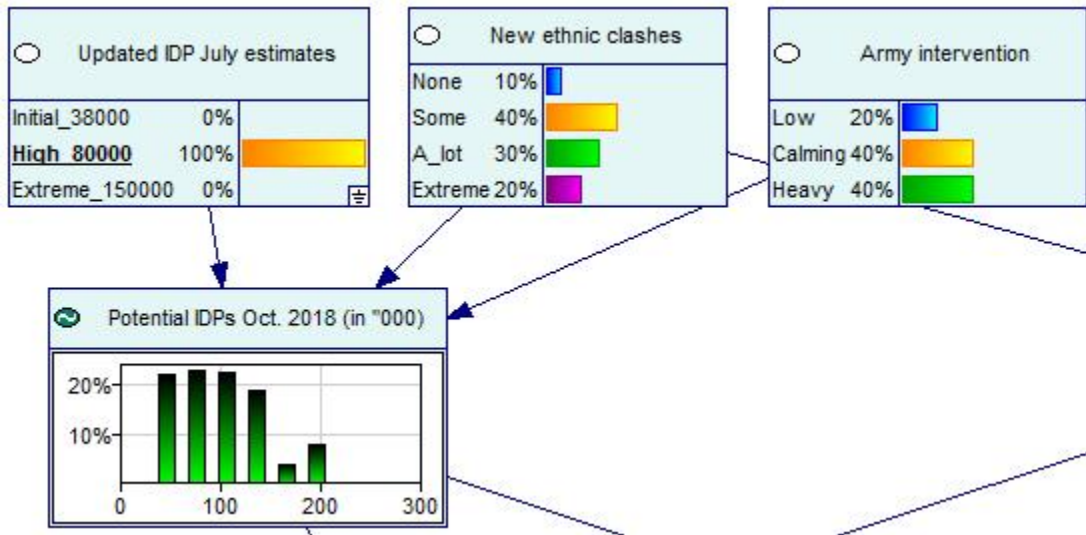
Multiply by	If the level of "New ethnic clashes" =
0.6	None
0.9	Some
1.3	A lot
2.0	Extreme

Multiply that result by	If the level of "Army intervention" =
1.0	Low
0.8	Calming
1.3	Heavy

That operation yields $3 * 4 * 3 = 36$ estimates; each is weighted by the joint probability of its parent node states. Together, they form the distribution that the rather imprecise histogram renders.

In practice, the authorities and agencies in Plateau states will most likely have a much better idea of the number of IDPs, long before it is clearer what the army and the local belligerents will be doing between now and October. Let us assume that the new IDP estimate is 80,000.

Figure 5: A network segment with one updated node



The shape of the distribution has changed, and so have the potential IDP statistics:

Mean:	StdDev:	Skewness:	Kurtosis:	Minimum:	Maximum:
101.011	44.4966	0.976327	3.34996	38.4	208

The expected number has gone up slightly, from 93,000 to 101,000. The standard deviation has gone down substantially, from 73,000 to 44,000. There is greater certainty. The remaining uncertainty is due to the as yet unknown behavior of the fighting factions and the army.

This is just one illustration of creating the probability distribution for a child node by means of an equation. We leave this technicality behind us, except for one nuance. This is often overlooked; it plays a role also in qualitative debates in scenario workshops; it can frustrate soft verbal and hard formal modeling when not properly considered. In short, there are conceptual differences depending on the location of the uncertainty:

- The causal mechanism is known with certainty. The initial conditions and triggering events are uncertain. This is the case of the potential IDP equation. The equation itself is deterministic. The inputs are uncertain until fully observed.
- Initial conditions and triggering events have been fully observed, or are certain to materialize as expected. The causal mechanism is poorly understood.

- Both types of uncertainty operate. This is the case of the number of “camps with disease outbreaks, July – Oct.” The equation is stochastic (a random element is built in)⁷, reflecting our ignorance about the behavior of epidemics and the emerging pattern of camps. The initial and evolving conditions too are uncertain. The uncertainty decreases, but does not fully go away, when the parent node values are fixed.

The third type of uncertainty is predominant. It easily overburdens qualitative reasoning. And it tempts quantitative analysts to pretend that they know more, or know it more accurately, than is feasible, given limited information and understanding. Bayesian belief networks, combined with sensitivity analysis, can sharpen our understanding of the uncertainty and thereby the limits of our knowledge.

It is time now to present some of the scenarios calculated with this Bayesian model.

Scenarios

Forward and backward-looking

We present three forward-looking scenarios and one of the backward-looking variety (these terms are ours). The forward-looking ones, as the reader may expect, fix all or most of the initial conditions and trigger events. The statistics of interest are those of the outcome variables, i.e. how many IDPs we are to expect, and how many camps are likely to report outbreaks.

Backward-looking ones fix an outcome value and, if at all, only a small subset of the initial conditions and trigger events. The interest is with the distributions of the free-node variables. The fixed outcome, by its repercussions on node layers above it, changes the probabilities of other outcomes as well. Of prime concern is the question what changes in the probabilities of trigger variables are needed if the outcome is to be a particular value. The scenario that we want to calculate concerns a decrease in displaced persons to below 50,000 by October.

For all four scenarios, we set the current IDP figure to 80,000, assuming that an assessment will be completed in July and will produce an estimate close to that.

A synoptic table gathers outcome statistics of the forward-looking scenarios.

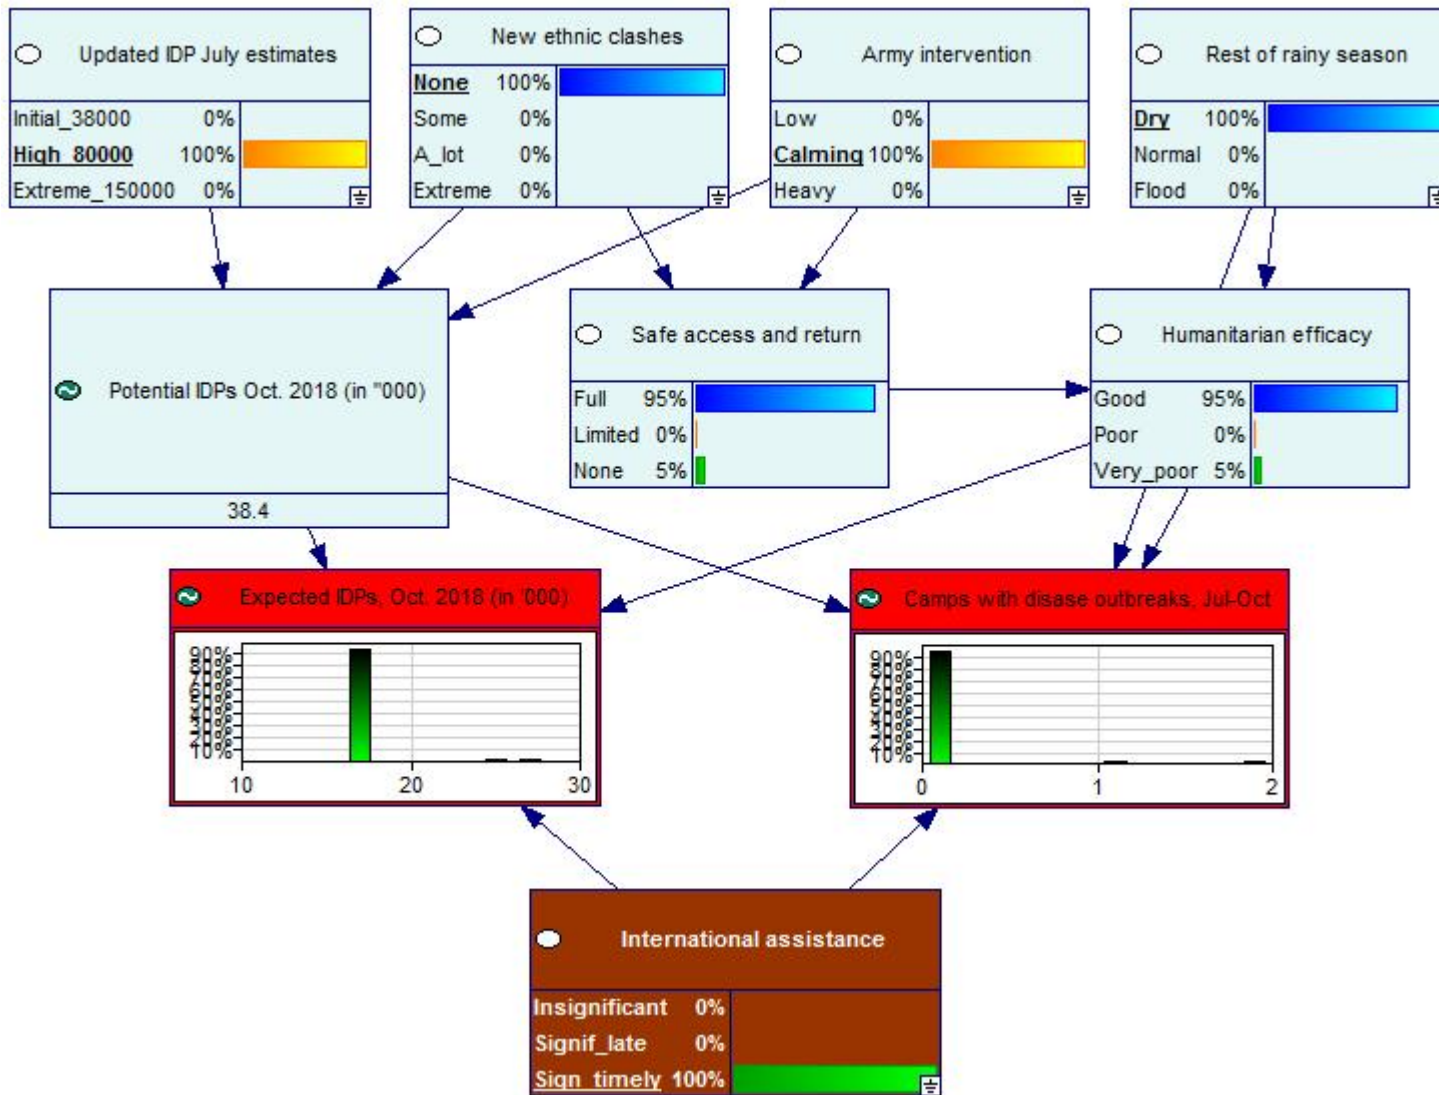
Updated networks

See the following pages.

⁷ The number of camps with outbreaks arrives from a binomial distribution, with the number of trials depending on the potential IDPs, and the success probability depending on the weather, on humanitarian efficacy as well as on international assistance. A technical discussion would exceed the ambition of this note.

Scenarios #1: Extremely favorable developments

Figure 6: Network for scenario #1



Initial condit. & trigger events

Intermediate states

Outcomes of interest

Decision variable

Scenarios #2: Extremely adverse developments

Figure 7: Network for scenario #2



Scenarios #3: Mixed developments, with unforeseeable herdsman attacks, but less than extreme

Figure 8: Network for scenario #3



Note that the probability of extreme ethnic clashes has been fixed to zero, with the other state probabilities adjusted proportionately.

Outcome synopsis

The table allows us to compare the outcome across scenarios as well as against the model set-up.

Table 4: Outcomes of three forward-looking scenarios

Domains and variables	Model set-up	Scenarios		
		#1: Favorable	#2: Adverse	#3: Middle of the road
Initial conditions				
IDPs in July	Not fixed	80,000	80,000	80,000
Trigger events				
New ethnic clashes	Not fixed	None	Extreme	Less than extreme
Army intervention	Not fixed	Calming	Heavy	Low
Rainy season	Not fixed	Dry	Flood	Normal
Decisions				
International assistance	Not fixed	Significant and timely	Insignificant	Significant, but late
Outcomes				
IDPs in October				
Min	7,700	16,100	125,000	25,900
Mean	79,200	16,700	207,000	63,800
[as % of IDPs in July]		21%	259%	80%
Max	390,000	26,900	208,000	93,600
SD	66,400	2,300	6,200	21,800
[Uncertainty:] C.o.V.	0.84	0.14	0.03	0.34
Camps with outbreaks				
Min	0	0	0	0
Mean	1.318	0.039	6.533	0.782
Max	16	2	11	5
SD	1.836	0.199	1.658	0.845
C.o.V.	1.39	5.10	0.25	1.08

Scenario #1 and 2 are of interest only in two respects. First, the difference in the outcomes, seen in the means, is huge – roughly 1 : 12 for the IDPs in October, and 1 : 170 for camps with outbreaks. Favorable conditions cause the majority of IDPs to return home; extremely adverse ones multiply the numbers in camps over the next three months. The tendencies are unsurprising; the proportions shock.

Second, the uncertainty of the IDP forecasts, expressed by the coefficients of variation (C.o.V. = SD / mean), is strikingly lower than in the set-up, in which all variables are probabilistic. Although in # 1 and

#2 two of the three intermediate states are not fully determined, in both of them one level is starkly dominant. For the number of camps with outbreaks, it is harder to find a valid uncertainty measure.

Scenario #1 and 2 form a kind of frame within which we expect most of the possible scenarios with less extreme trigger events to fall. Therefore, our major substantive interest lies with scenario #3. Formally, one could call it a “middle of the road” choice. From a humanitarian perspective, “benign neglect” fits better. Neither army nor international agencies are doing much in the area. By end of August, #3 assumes that the rainy season was normal; Plateau State was spared the extremes of drought and flood. The major uncertainty persisting is in the behavior of hostile ethnic groups. Key informants and experts rule out that the violence will climb to extreme levels. “Lots of attacks” are still a strong possibility.

This reduced uncertainty is enough to rule out massive increases in displacement and to cherish, as the lead hypothesis, a minor reduction. The same holds for disease outbreaks.

Still, the uncertainty is at levels that would make response planning difficult. The C.o.V. of the IDP forecast is 0.34, much higher than for scenarios #1 and 2. Regarding the camps with outbreaks, assessment analysts would surely feel more comfortable with a forecast range of [0, 2] (scenario #1) than [0, 5] (#3), which in turn is much narrower than [0, 11] (#2)⁸.

If indeed extreme ethnic violence is ruled out, this confidence must be based on some knowledge of the belligerents’ relationship that is not yet part of the belief network. One wonders whether “international assistance” should be replaced with two nodes: one capturing the level of physical relief, the other the level of protection activities. These can be performed by national and/or international agencies. *Avant la lettre* of Bayesian networks, and with primitive tools, Benini (1993) simulated such a causal network, distinguishing between protection and assistance efforts, in an ethnic conflict in Mali. Their synergetic interaction should be amenable to belief networks. The current network is too simplistic to investigate such questions.

A backward-looking scenario

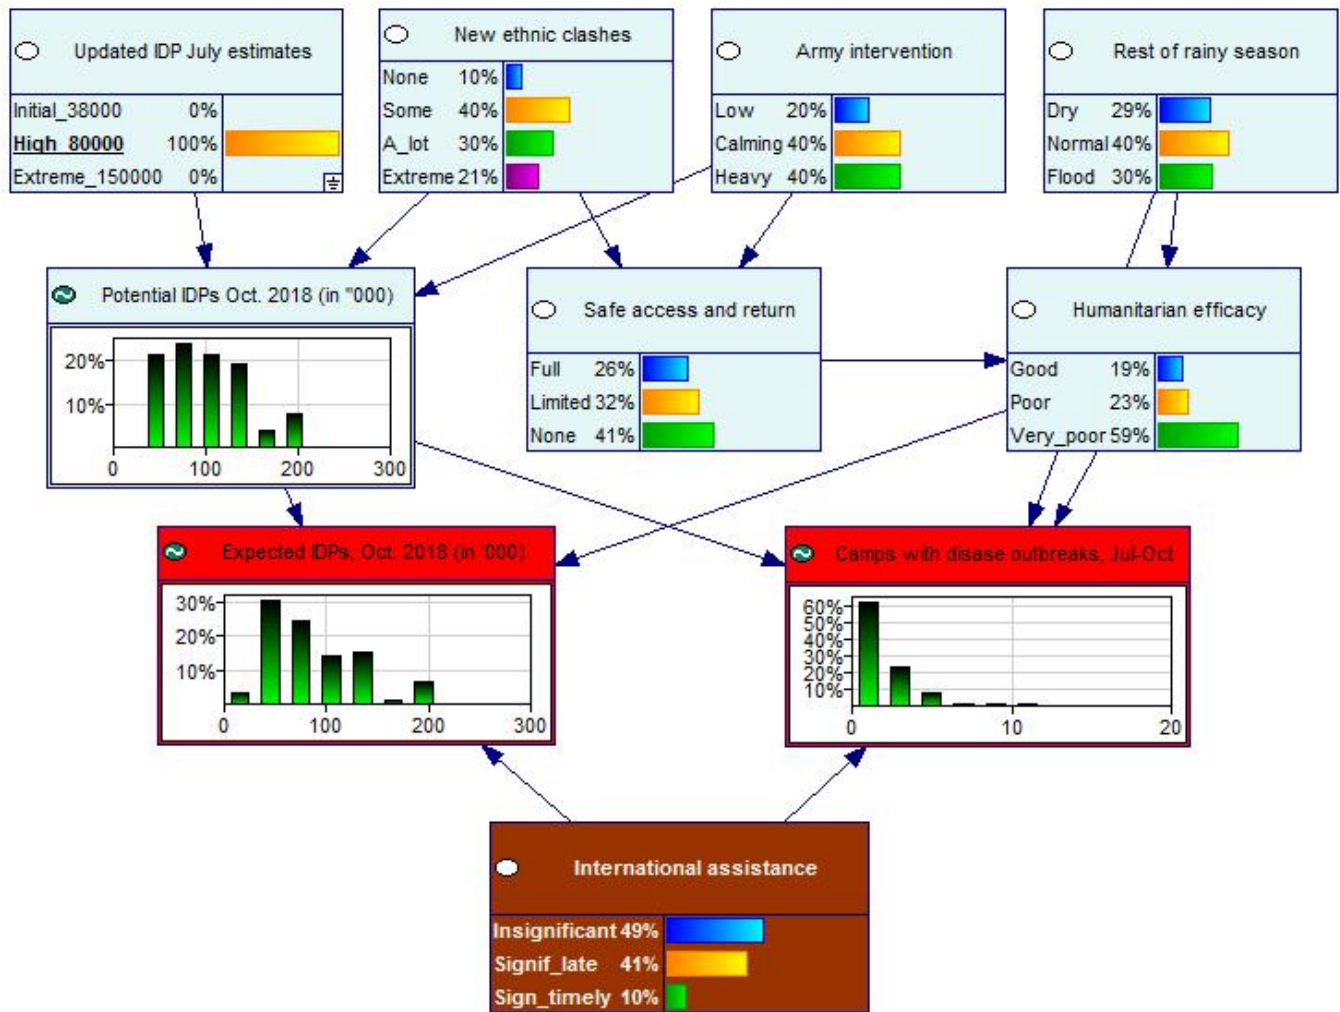
Starting point

Our starting point is a belief network in which the only evidence posted is the updated IDP estimate for July, which we set to 80,000 persons, as in the other scenarios. All the other starting nodes are in their usual probability distributions, as in the set-up. Looking into the outcome nodes, we find the mean of IDPs expected by October to be 87,500, and the mean of camps with outbreaks = 1,473.

This is the calculated network from which we start:

⁸ This raises an interesting point for further work: one should compute an additional scenario #3b, the same in all input variables as #3, except that the probability distribution of new ethnic clashes is returned to the initial set-up. The comparison of the changes in outcomes to the changes in the probabilities of further ethnic clashes would be a *sensitivity measure* of interest. We have not yet performed such tests.

Figure 9: Starting point of the backward-looking scenario exercise



Conditions for reduced displacement


The question that we want to ask the Mighty Bayesian Ghost in the Bottle is:

What are the conditions under which the IDPs in October will be below 50,000?

To get the Ghost to murmur an answer, we enter a value < 50 (in thousands) as evidence in the Expected IDPs by October node, such as 49.9.

The picture changes:

Figure 10: Backward-looking scenario



The appearance of the outcome nodes has greatly changed. Bright red warns us that at least one of the outcome nodes has been fixed, and the 100% bar value in the “Expected IDPs node” makes it clear it is indeed this node. The variable type too has been changed – from continuous counts to intervals⁹.

What we want to understand is the meaning of the changed probabilities in the trigger events:

- New ethnic clashes: A reduction of the IDPs to below the 50,000 will not happen amid extreme ethnic violence even if all other triggers are favorable (probability of “Extreme” = 0%, to be precise: = 0.03187%; i.e. virtually impossible). It is unlikely to happen when there is “a lot of” violence (15%). However, the continuation of “some” clashes does not necessarily derail the reduction; the chances are, so to speak, fifty/fifty. All depends on what the army, the weather and the aid agencies do.

⁹ The operation of “discretization”, which preceded this change, is beyond the scope of this note. We only show the result.


- Rest of the rainy season: Floods will likely frustrate the hoped-for return of 30,000 or more IDPs. Floods depress humanitarian efficacy, which leaves more IDPs unassisted.
- Army intervention: Our hypothesis that the number of IDPs will drop below 50,000 is improbable if the army makes little effort intervene or, conversely, does so with a heavy hand. A calming interposition would make it much more probable.
- International assistance: This is perhaps the most interesting result. Yes, significant and timely assistance is helpful (the probability of seeing the international agencies operating at that level has risen from 10% in the starting point to 18%, almost a doubling). 18% is still very low. Significant and timely international assistance is not a major precondition to the IDP return success.

More specific "What .. if?"

But what if the international community mobilizes itself and assists the affected populations in a significant and timely manner? Let us assume also that the level of army intervention will be low, and the rainy season has been normal. Thus, four out of the five starting nodes have been fixed by assumption. The "Expected IDPs" node too has been fixed because we are investigating conditions for the drop in the IDP number. Only "New ethnic clashes" remains probabilistic.

The resulting network returns that distribution for this node.

Figure 11: A more specific "What-if" belief network



It is clear that, even at that level of international assistance, the reduction cannot be achieved if the ethnic violence continues at a significant level ("A lot" or "Extreme"). The persistence of sporadic clashes ("Some") still presents a strong risk for the return of IDPs (100% - 71% = 29%), almost a one-in-three chance to fail. But the optimist now sees a glass two-thirds full.

How do we need to understand such probabilities in a "What - if?" scenario like that? What does it mean that when the IDP numbers drop the probability of no further ethnic clashes is 27%, of some is 71%, and of a lot a mere 2%?

It helps to reformulate the probabilities in a number of "possible worlds" clauses (Booth, Rowlinson et al. 2009, Bruner 2009, Barma, Durbin et al. 2016):

IF

- our understanding of the crisis in Plateau State is correct, and
- to the extent that our model encodes that knowledge correctly, and

- the history of the world repeated itself millions of times, and
- in thousands of those histories
 - The number of IDPs in July was 80,000, and in the following three months
 - The level of army intervention was low, and
 - The rainy season was normal, and
 - International assistance arrived in significant and timely manner

THEN

- in every 100 histories that also see IDPs drop below 50,000 by October
 - In 27 we would find no further ethnic clashes between now and then,
 - in 71 some clashes, and
 - in 2 a lot of them.

Think “possibilities” instead of “histories”, and it is clear that this multi-condition scenario is (almost) impossible to materialize while significant ethnic violence continues.

That is exactly what the Bayesian algorithm does. It samples possible worlds – 10,000 per run, to be precise – and computes the probability of any type of world within the domain that we define for it.

By extension, it is obvious that backward-looking scenarios reveal trade-offs among initial conditions and trigger events for a fixed outcome. This ability is helpful in humanitarian response planning. It encourages multi-option thinking of the kind: “We can achieve this objective if ABC, or BCD, or perhaps if ACD, but not likely so if only AB, BC or CD.”

Belief networks in humanitarian scenario development

This brief demonstration of what Bayesian belief networks can do originates from a limited understanding of a little known humanitarian crisis, as reported in one short briefing note. Neither have we given the reader an idea of the full potential of this analytic technique, nor does our model adequately mirror the complexity of that crisis. It is a purely didactic demonstration to rouse interest among analysts who, at least occasionally, are asked to contribute to the development of humanitarian scenarios.

Typically, scenario thinking is stimulated in workshops that bring together field workers, area managers and consultants from a variety of frontline, technical and donor organizations. Experience shows that participants come with greatly different understandings of cause-effect models and of probability concepts. This concerns not only substantive understandings – in which diversity is desirable, in fact necessary for productive debate -, but also the elementary notions of probability.

Given the typical composition of participants, it is unlikely that a technical vehicle like Bayesian networks can help to organize, channel and focus the conduct of the workshop itself. In fact, the mere demonstration of a calculated network could have an exclusionary effect and create the suspicion of foregone conclusions, a sure killer of the kind of open debate that produces fresh scenarios. It would

also go against current philosophies of humanitarian analysis that favor taking the effort closer to the affected population¹⁰.

Against those reservations, one may hold a more optimistic view. Založnik, Bonsall and Harper (2018) built Bayesian belief networks of a qualitative kind with focus groups in rural Vietnam. They found it helpful to divide the work into two phases. In the first, several groups would each build a provisional network structure of the problem at hand (fertilizer and irrigation management) during a two-hour session. The researchers would consolidate the networks. In the second round, they would present a common structure to all the groups, who would then – separately - elaborate the framework further. The researchers obtained two workable versions from which they were able to collect quantitative data as the basis for the needed conditional probability tables.

If it works with farmers on the other side of a language barrier, one is tempted to think it should enthruse workshop participants who are more highly educated, and most of whom speak the same language. This much seems to have happened in some professional communities. Notably, environmental scientists have advanced the methodology of modeling with stakeholders (Voinov and Bousquet 2010, Voinov, Kolagani et al. 2016, Gray, Voinov et al. 2018). In some near future, something from this treasure house of participatory experience will, hopefully, flow into the humanitarian community.

Meanwhile, we have to be honest and humble and, from personal experience, admit that Bayesian belief networks are not for the faint of heart. The analyst who listens to a field worker detailing the evolution of a crisis must be able to outline conditional probability tables in the back of his mind. And, come time, ask questions that elicit bits of “after this, that” and “when X, then Ys and Zs happen, roughly in this proportion”, “if X stopped, Y would likely continue, but not Z”, etc. And sit down with knowledgeable participant in order to map these relationships more completely.

We see two likely productive uses of belief networks in the near future of humanitarian scenario development:

- In generating interest in, and support for, scenario building at the invitation and agenda-building stage, regardless of whether the plan is for a physical workshop or for some form of remote cooperation. Similarly in the “drawing it all together in an orderly form” phase after workshops or sequential consultations.
- In alerts and briefing notes that seek to attract the attention of donors, relief agencies, researchers and media to emerging or persistent crises, using the beliefs of the authors as a basis for small network models that capture different scenarios. Likewise in funding requests, particularly those of which “theory of change”-based rationales are expected.

The linchpin in all of that, of course, is the individual analyst who embarks on a learning curve that at the start may look daunting, but, as he/she will discover, is surprisingly gentle. It quickly leads to green

¹⁰ See, e.g., <http://www.jips.org/en/news/latest-news/improving-data-to-put-people-first-in-idp-situations>.

pastures crisscrossed by a network of paths that lead to fresh discoveries and joyful encounters with others who dare to look deeply into the humanitarian landscape. So we believe.

Software

The models shown in this note were built with GeNIe, the graphical user interface to SMILE (Structural Modeling, Inference, and Learning Engine), produced by BayesFusion, LLC.

The user manual is available at <https://support.bayesfusion.com/docs/GeNIe.pdf>.

Appendix

Calculating marginal probabilities from the Conditional Probability Table

We calculate them in this table by replacing the verbal labels of the army intervention and ethnic clashes states with their probabilities, then multiplying and summing:

Table 5: Calculating joint and marginal probabilities

Probabilities:							
Unconditional		Conditional			Multiplied (so-called "joint")		
Army intervention	New ethnic clashes	Safe access and return			Safe access and return		
		Full	Limited	None	Full	Limited	None
0.2000	0.1000	0.8000	0.1000	0.1000	0.0160	0.0020	0.0020
0.2000	0.4000	0.4000	0.5000	0.1000	0.0320	0.0400	0.0080
0.2000	0.2900	0.0500	0.5000	0.4500	0.0029	0.0290	0.0261
0.2000	0.2100	0.0010	0.3000	0.6990	0.0000	0.0126	0.0294
0.4000	0.1000	0.9500	0.0030	0.0470	0.0380	0.0001	0.0019
0.4000	0.4000	0.7000	0.2000	0.1000	0.1120	0.0320	0.0160
0.4000	0.2900	0.3000	0.5000	0.2000	0.0348	0.0580	0.0232
0.4000	0.2100	0.0010	0.2000	0.7990	0.0001	0.0168	0.0671
0.4000	0.1000	0.2000	0.5000	0.3000	0.0080	0.0200	0.0120
0.4000	0.4000	0.1000	0.4000	0.5000	0.0160	0.0640	0.0800
0.4000	0.2900	0.0010	0.3000	0.6990	0.0001	0.0348	0.0811
0.4000	0.2100	0.0001	0.1000	0.8999	0.0000	0.0084	0.0756
Summed (so-called "marginals"):					0.2600	0.3177	0.4223
					Same as in the bar char		

Assuming that these columns are numbered 1 through 8 in an Excel spreadsheet, the formulas in the light green cells in columns 6 – 8 are identical: `"=RC1*RC2*RC[-3]"`, with mixed and relative references.

Once evidence has been entered into the "Army intervention" and "New ethnic clashes" nodes, the joint probabilities (light green cells in columns 6 – 8) are positive in one row only; all others are turned to zero.

Suppose we update "Army intervention" = "Low", and "New ethnic clashes" = "None", this is what you would get. The CPT stays the same, the derived joint and marginal probabilities change.

Table 6: Updating joint and marginal probabilities

Probabilities:								
Unconditional		Conditional			Multiplied (so-called "joint")			
Army intervention	New ethnic clashes	Safe access and return			Safe access and return			
		Full	Limited	None	Full	Limited	None	
1.0000	1.0000	0.8000	0.1000	0.1000	0.8000	0.1000	0.1000	
1.0000	0.0000	0.4000	0.5000	0.1000	0.0000	0.0000	0.0000	
1.0000	0.0000	0.0500	0.5000	0.4500	0.0000	0.0000	0.0000	
1.0000	0.0000	0.0010	0.3000	0.6990	0.0000	0.0000	0.0000	
0.0000	1.0000	0.9500	0.0030	0.0470	0.0000	0.0000	0.0000	
0.0000	0.0000	0.7000	0.2000	0.1000	0.0000	0.0000	0.0000	
0.0000	0.0000	0.3000	0.5000	0.2000	0.0000	0.0000	0.0000	
0.0000	0.0000	0.0010	0.2000	0.7990	0.0000	0.0000	0.0000	
0.0000	1.0000	0.2000	0.5000	0.3000	0.0000	0.0000	0.0000	
0.0000	0.0000	0.1000	0.4000	0.5000	0.0000	0.0000	0.0000	
0.0000	0.0000	0.0010	0.3000	0.6990	0.0000	0.0000	0.0000	
0.0000	0.0000	0.0001	0.1000	0.8999	0.0000	0.0000	0.0000	
Summed (so-called "marginals"):					0.8000	0.1000	0.1000	
					Same as in the bar chart			

In the "Safe access and return" node you will now see: Full: 80%, Limited: 10%, None: 10%.

The briefing note: "Nigeria – Displacement in Plateau State"

The note (START Network and ACAPS 2018) is appended.

References

- ACAPS (2012). Scenario Development [Technical Brief, November 2012]. Geneva, Assessment Capacity Project.
- Aspinall, W. P. and G. Woo (2014). "Santorini unrest 2011–2012: an immediate Bayesian belief network analysis of eruption scenario probabilities for urgent decision support under uncertainty." Journal of Applied Volcanology 3(1): 12.
- Barma, N. H., B. Durbin, E. Lorber and R. E. Whitlark (2016). "'Imagine a World in Which': Using Scenarios in Political Science." International Studies Perspectives 17(2): 117-135.
- Benini, A. A. (1993). "Simulation of the Effectiveness of Protection and Assistance for Victims of Armed Conflict: An Example from Mali, West Africa." Journal of Contingencies and Crisis Management 1(4): 215-228.
- Booth, C., M. Rowlinson, P. Clark, A. Delahaye and S. Procter (2009). "Scenarios and counterfactuals as modal narratives." Futures 41(2): 87-95.
- Bruner, J. S. (2009). Actual minds, possible worlds, Harvard University Press.
- Charniak, E. (1991). "Bayesian networks without tears." AI magazine 12(4): 50-63.
- Gray, S., A. Voinov, M. Paolisso, R. Jordan, T. BenDor, P. Bommel, P. Glynn, B. Hedelin, K. Hubacek, J. Introne, N. Kolagani, B. Laursen, C. Prell, O. L. Schmitt, A. Singer, E. Sterling and M. Zellner (2018). "Purpose, processes, partnerships, and products: four Ps to advance participatory socio-environmental modeling." Ecological Applications 28(1): 46-61.
- IPC Global Partners (2012). Integrated Food Security Phase Classification Technical Manual. Version 2.0. Evidence and Standards for Better Food Security Decision. Rome, FAO.
- Marcot, B. G. (2017). "Common quandaries and their practical solutions in Bayesian network modeling." Ecological Modelling 358: 1-9.
- Nigeria INGO Forum and ACAPS (2016). Impact of Insurgency in Northeast Nigeria: Scenarios. Possible developments in northeast Nigeria over the next 9 months (to June 2017) [03 October 2016] Maiduguri, Nigeria and Geneva, Switzerland.

START Network and ACAPS (2018) "Nigeria - Displacement in Plateau State. Briefing note - 10 July 2018." from https://www.acaps.org/sites/acaps/files/products/files/20180710_acaps_start_nigeria_plateau_displacement_0.pdf.

Voinov, A. and F. Bousquet (2010). "Modelling with stakeholders." Environmental Modelling & Software 25(11): 1268-1281.

Voinov, A., N. Kolagani, M. K. McCall, P. D. Glynn, M. E. Kragt, F. O. Ostermann, S. A. Pierce and P. Ramu (2016). "Modelling with stakeholders – Next generation." Environmental Modelling & Software 77: 196-220.

Založnik, M., M. B. Bonsall and S. Harper (2018). "The Qualitative Stage of Building Bayesian Belief Networks in a Focus Group Setting: Decision-Making under Uncertainty among Vietnamese Rice Farmers." Sociological Methods & Research: 0049124118769094.