Of Confidence, Control, and Cause: Using Bayesian Networks for Management Decisions

Bruce G. Marcot, Ph.D.
Research Wildlife Biologist
U.S. Forest Service

This file is an annotated version of, and can be cited as:

Marcot, B. G. 2014. Of confidence, control, and cause: using Bayesian networks for management decisions. Presented 26 November 2014 at the Sixth Annual Conference of the Australasian Bayesian Network Modelling Society, Rotorua, New Zealand. [Invited keynote address].

(Further information on Bruce Marcot's work on Bayesian network modeling can be found on his web page: http://www.spiritone.com/~brucem/bbns.htm)

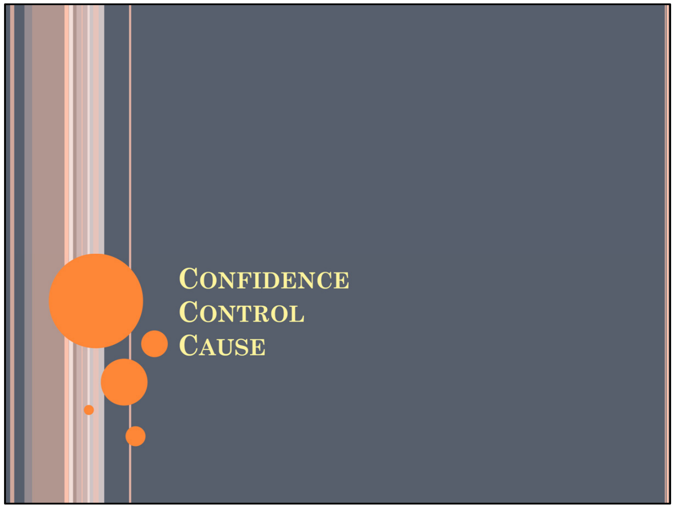

This presentation will delve into three key areas pertaining to Bayesian network modeling: confidence, control, and cause.

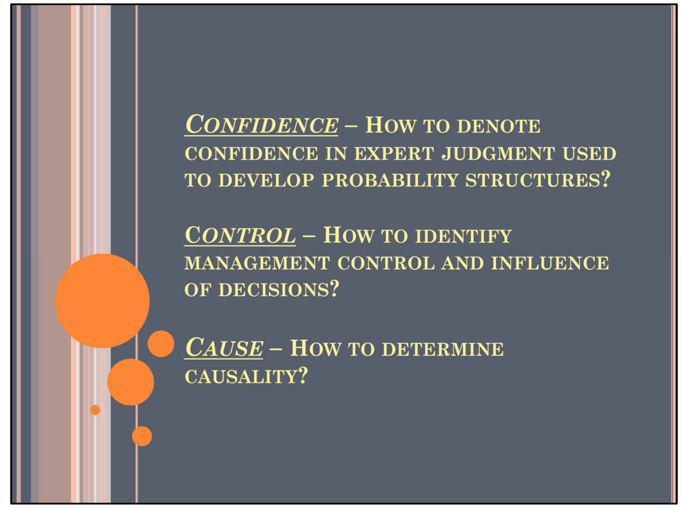

Confidence here refers to using expert judgment for developing Bayesian network probability structures.
Control here refers to identifying how management can know about the degree to which they can control outcomes of their decisions.
Cause here refers to how Bayesian network models can truly depict causal structures.

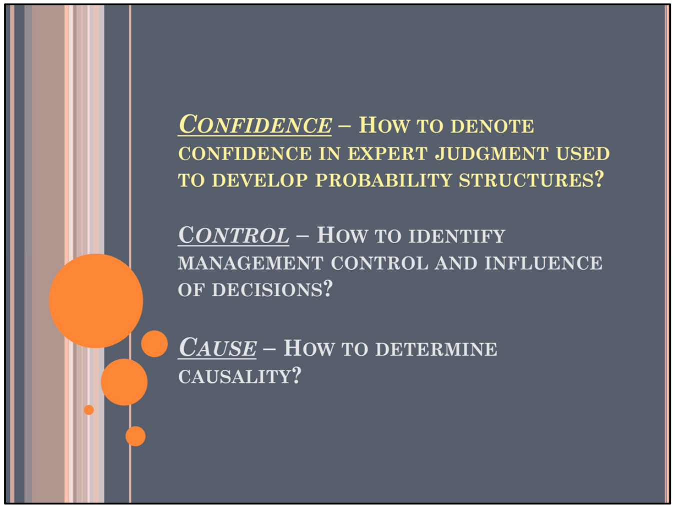

Let's begin by exploring the question of *confidence.*

I would like to lobby for use of the term "expert knowledge" instead of "expert opinion." If we build Bayesian network models at least in part from expert's input, the models should be rigorous, testable, credible, and hold up to peer review scrutiny. I have come to avoid using the term "opinion" in this context, which sounds far more capricious and arbitrary.

Key First Step: What's the objective?

The major first stage in Bayesian network modeling (or any modeling project) should always be to clarify the objective.

Key First Step: What's the objective?

- **prediction** - *possible future outcomes based on initial conditions*
- **forecast** - *most likely future outcome based on initial conditions*
- **projection** - *possible future outcomes based on changing future conditions*
- **scenario analysis** - *implications of hypothetical situations*
- **diagnosis** - *determine potential causes of a known or specified condition or outcome*
- **data mining** - *find patterns in big data*
- **summarize knowledge** - *synthesize what we think we know*
- **identify key data gaps** – *factors, interactions with greatest influence on outcomes*
- **mitigation** – *i.d. alternative conditions that could lead to a desired outcome*
- **aid individual or collaborative decision-making** – *risk analysis & risk management*

There can be a very wide array of potential objectives and uses for Bayesian network models. And no one model can do it all. I have told my students and colleagues that they should pick their top three objectives, then prioritize those, then pick the top ONE from that list.

Next Steps in Model-Building

Next, let's consider a general framework for model-building steps and stages …

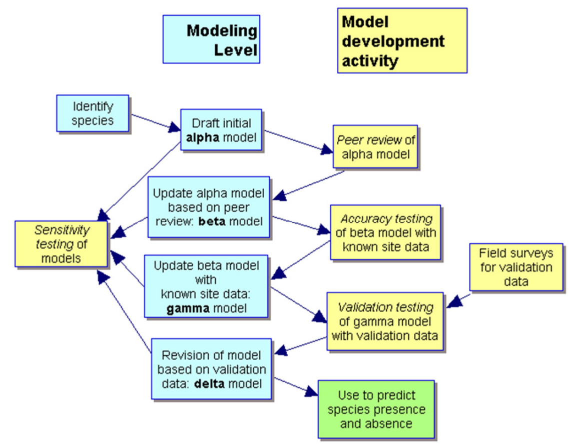

This graphic depicts a general framework for model-building steps that I had evolved over time, and published. It entails use of expert knowledge to build the first ("alpha") level model, then subsequent uses of peer review, and testing of model sensitivity, accuracy, and validation.

Source: Marcot, B. G. 2006. Characterizing species at risk I: modeling rare species under the Northwest Forest Plan. Ecology and Society 11(2):10. [online] URL: http://www.ecologyandsociety.org/vol11/iss2/art10/.

NEXT STEPS IN MODEL-BUILDING

- **Develop a simple influence diagram representation of the key factors and linkages**

Let's go through the various steps in building a Bayesian network model under this framework. (The text in these next series of slides provides the main points.)

NEXT STEPS IN MODEL-BUILDING

- **Develop a simple influence diagram representation of the key factors and linkages**
- **Document the influence diagram**
  - **data, info sources, pubs, consultations used**
  - **identify what each arrow represents – correlation, causation**
  - **peer review**

Next Steps in Model-Building

- Develop a simple influence diagram representation of the key factors and linkages
- Document the influence diagram
  - data, info sources, pubs, consultations used
  - identify what each arrow represents – correlation, causation
  - peer review
- Clearly define each variable (node) & their states
  - input nodes
  - intermediate nodes (latent variables, summary nodes)
  - output nodes

Next Steps in Model-Building

- Develop a simple influence diagram representation of the key factors and linkages
- Document the influence diagram
  - data, info sources, pubs, consultations used
  - identify what each arrow represents – correlation, causation
  - peer review
- Clearly define each variable (node) & their states
  - input nodes
  - intermediate nodes (latent variables, summary nodes)
  - output nodes
- Document probability values
  - basis, source of info, methods used to derive
  - what do they represent? – frequencies, relative outcomes

Next Steps in Model-Building

- Sensitivity analysis
  - underlying probability structure of the model (inputs set to default prior probabilities)

Next Steps in Model-Building

- Sensitivity analysis
  - underlying probability structure of the model (inputs set to default prior probabilities)
- Influence analysis
  - inputs w/ most influence on outcome (for given scenarios)

Next Steps in Model-Building

- Sensitivity analysis
  - underlying probability structure of the model (inputs set to default prior probabilities)
- Influence analysis
  - inputs w/ most influence on outcome (for given scenarios)
- Calculate metrics of model complexity
  - no. variables (nodes)
  - no. of links
  - no. of node cliques
  - no. of probabilities

Next Steps in Model-Building

- Sensitivity analysis
  - underlying probability structure of the model (inputs set to default prior probabilities)
- Influence analysis
  - inputs w/ most influence on outcome (for given scenarios)
- Calculate metrics of model complexity
  - no. variables (nodes)
  - no. of links
  - no. of node cliques
  - no. of probabilities
- Calculate metrics of model performance
  - classification error rates (confusion tables)
  - many metrics available (AUC, TSS, spherical payoff, etc.)

Next Steps in Model-Building

- Update part or all of the model using case data

Next Steps in Model-Building

- Update part or all of the model using case data
- Calculate metrics of uncertainty in posterior probability distributions
  - Bayesian credible intervals
  - normalized Gini index
  - other metrics (posterior probability certainty index, certainty envelope)

Next Steps in Model-Building

- **Update part or all of the model using case data**
- **Calculate metrics of uncertainty in posterior probability distributions**
  - **Bayesian credible intervals**
  - **normalized Gini index**
  - **other metrics (posterior probability certainty index, certainty envelope)**
- **Compare alternative posterior probability distributions**

Perhaps not every model-building project needs to go through all of the steps listed in this slide sequence, but the main steps entailing peer review and model testing and updating most lend to ensuring scientific credibility of the final product.

Using Expert Judgment in Model-Building

So how can, or should, expert judgment (or expert knowledge) be used in building Bayesian network models?

Using Expert Judgment in Model-Building

- for building the initial influence diagram
- for identifying conditional probabilities or other variable relationships
- iterative procedures
  - internal team reviews
  - external peer reviews

Depending on the type of model to be developed, and the availability of empirical data (or lack thereof) by which to structure and parameterize the model, expert judgment could play a key role in building the initial influence diagram (the "boxes and arrows" stage), and then to identify probability structures. The more that a Bayesian network model depends on expert knowledge, the more the model-builder should take the time and trouble to seek internal and external peer reviews, and then reconcile the reviews clearly by updating the model as deemed necessary.

Using Expert Judgment in Model-Building

- state the goal, objective, purpose, scope, intended use, and intended audience for the model …
- … then … if building the model from multiple experts … decide:

In working with a TEAM of domain experts, or at least with MULTIPLE experts, when building an expert-based model, first clearly state the modeling goals etc., and then …

Using Expert Judgment in Model-Building

- state the goal, objective, purpose, scope, intended use, and intended audience for the model …
- … then … if building the model from multiple experts … decide:

One model

Solicit probability values from the panel via:
- group consensus (e.g., Delphi method)
- guided brainstorming
- polling individuals, then combine results
  - simple average
  - weighted average or sums
- parcel out submodels to specific experts, then combine into one model

… then decide if you want to end up with ONE model only, that would represent the multiple experts' collective knowledge in a single network … OR …

Using Expert Judgment in Model-Building

- state the goal, objective, purpose, scope, intended use, and intended audience for the model …
- … then … if building the model from multiple experts … decide:

One model	> One model
Solicit probability values from the panel via:	Represent differences in expert knowledge
- group consensus (e.g., Delphi method)	- develop range of model structures
- guided brainstorming	- develop range of probability values
- polling individuals, then combine results	- compare results to denote variation in expert judgments
- simple average	
- weighted average or sums	
- parcel out submodels to specific experts, then combine into one model	

… or if you want to potentially develop more than one model so as to depict the range of the experts' ideas and experience on model structure and probability parameters.

Knowledge elicitation

- rich literature
  - artificial intelligence, expert systems programming (1970s - present)
- group elicitation – expert panels
  - biases (e.g., motivational bias, knowledge parity, facilitator bias)

There is a rich history of publications and practices associated with expert elicitation methods.

Knowledge elicitation – expert panels

- How to denote degree of confidence in probability values, in Bayesian networks built from expert judgment?
  - feathering CPT values is an expression of uncertainty – but not expressing confidence in the supplied probability values
  - expert systems have used a confidence indexing approach – [0,1] or [0 , 100%]
  - BaysiaLab uses a confidence factor approach outside the BN
  - Netica uses "experience tables" – but only when inducing probability structures from mostly empirical case files

There is also a great deal of subjectivity used in denoting expert-judgment "confidence levels."

Knowledge elicitation – expert panels

- **Real-world example: Polar Bear Science Team** *(7 members)*
  - **needed clear direction on resolving differences of opinions**
  - **consensus or competing models?**
  - **“round-robin” reviews and edits of CPTs**

http://www.intuitor.com/statistics/SmallGroups.html

In one of my real-world examples, I led a team of 7 researchers (including wildlife biologists, statisticians, & climate scientists) in the modeling of global polar bear populations. We had differing ideas of the overall model structure and its probability tables. From the start, we decided that we wanted to create one model representing the consensus of our team, so we instituted a procedure of a “round robin” series of team member reviews of the model, each reviewer adding their ideas and edits as it was passed around. The final decision on amending the model was up to the team leader and myself.

Source: Atwood, T. C., B. G. Marcot, D. C. Douglas, S. C. Amstrup, K. D. Rode, G. M. Durner, and J. F. Bromaghin. 2015. Evaluating and ranking threats to the long-term persistence of polar bears. U.S. Geological Survey, Open-File Report 2014-1254. http://dx.doi.org/10.3133/ofr20141254. Anchorage, Alaska. 114 pp.

Knowledge elicitation – expert panels

- **Real-world example: Polar Bear Science Team** *(7 members)*
  - **needed clear direction on resolving differences of opinions**
  - **consensus or competing models?**
  - **“round-robin” reviews and edits of CPTs**
  - **no. of interactions as a function of size of the team – factorial calculations**

http://www.intuitor.com/statistics/SmallGroups.html

In general, however, note that as the size of an expert panel or modeling team increases, the number of potential interactions (and “camps” and debates) has the potential to balloon enormously …

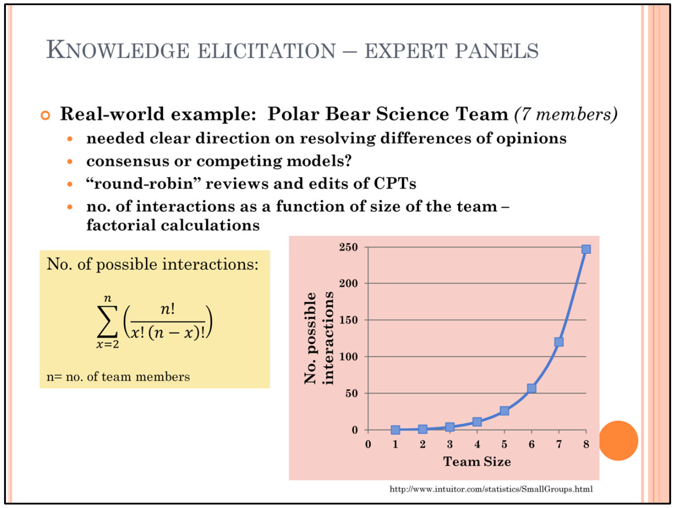

… as depicted with this simple factorial calculation. So it's critical to be very clear at the onset on team objectives and methods. Also, I have strived to keep such teams or panels to about 5 members, beyond which the debates and differences can quickly grow out of hand, UNLESS one asserts a strict structure for querying members, for panel discussions, and for closing discussions.

In fact, in some Bayesian network applications in management, I have co-led panels up to 10 members strong with great success, by clearly structuring discussions and meeting agendas.

Source: Marcot, B. G., P. A. Hohenlohe, S. Morey, R. Holmes, R. Molina, M. Turley, M. Huff, and J. Laurence. 2006. Characterizing species at risk II: using Bayesian belief networks as decision support tools to determine species conservation categories under the Northwest Forest Plan. Ecology and Society 11(2):12. [online] URL: http://www.ecologyandsociety.org/vol11/iss2/art12/.

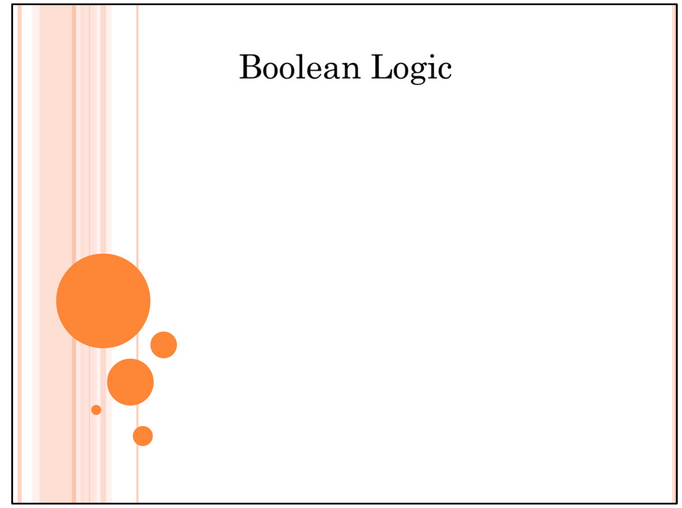

For some expert-based models, I have used Boolean logic as the basis for the conditional probability tables.

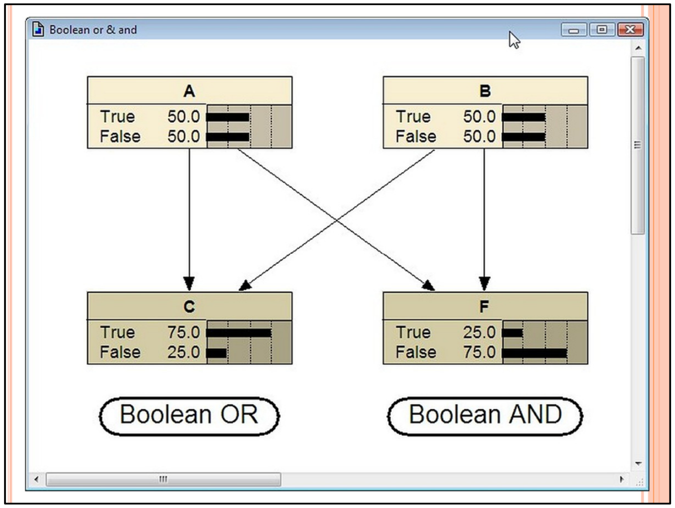

You can easily represent Boolean “or” and “and” functions in conditional probability tables …

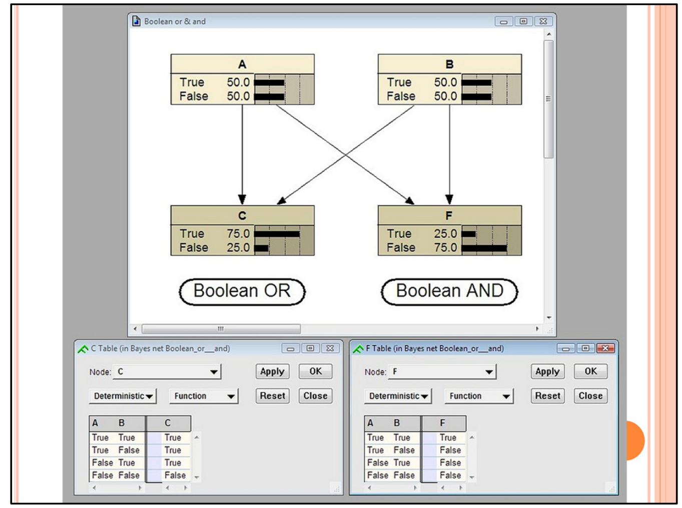

… as shown here with simple, deterministic CPTs.

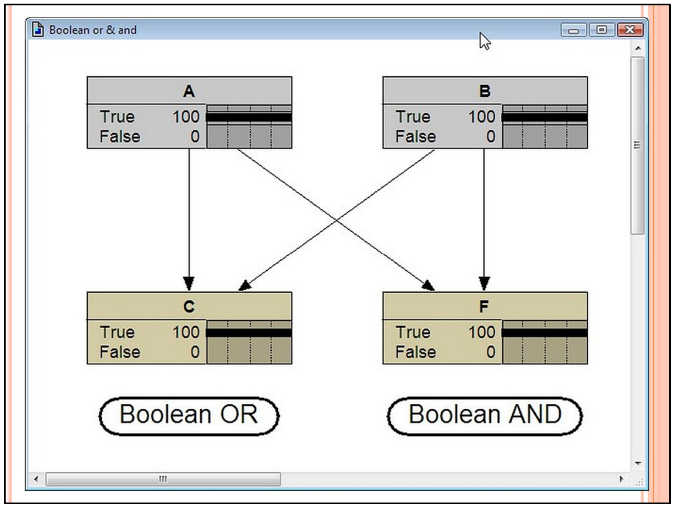

This simple model would then work as shown here and in the subsequent several slides. When A is true and B is true, then A or B is true, and A and B is true.

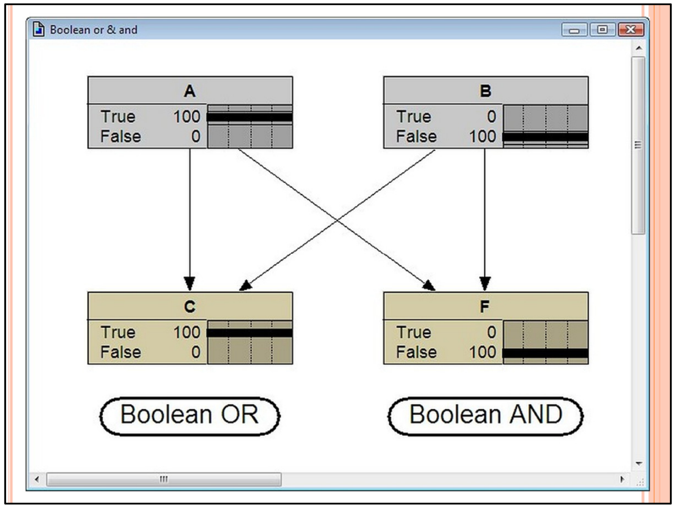

When A is true and B is false, then A or B is true, and A and B is false.

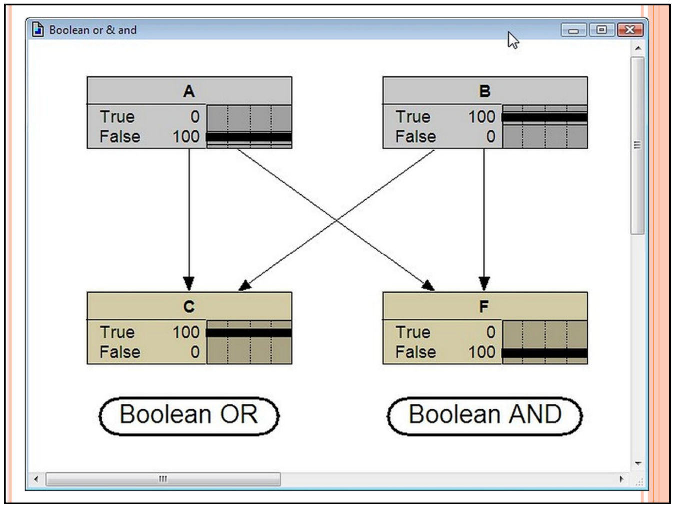

When A is false and B is true, then A or B is true, and A and B is false.

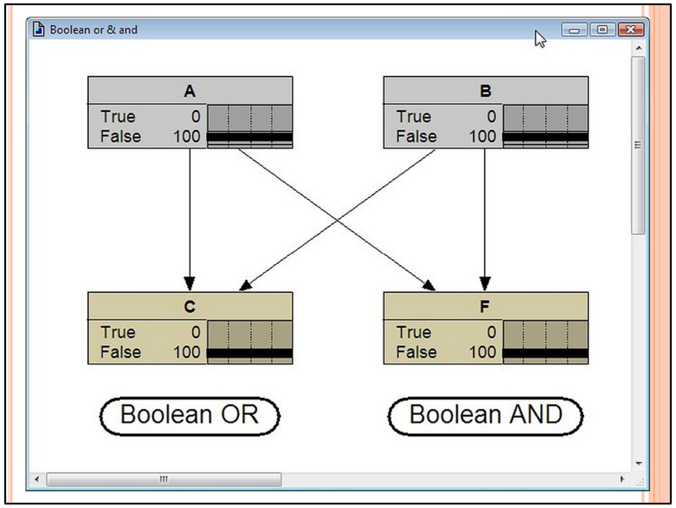

And when A is false and B is false, then A or B is false, and A and B is false.

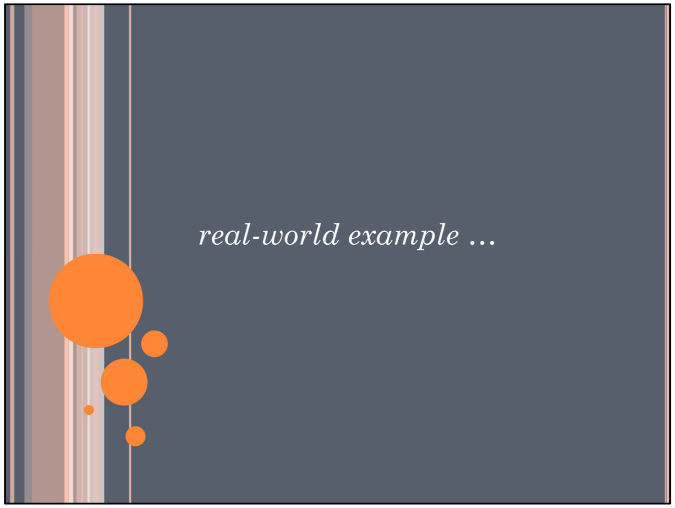

Here’s a real-world example of a complex Bayesian network model I developed that is essentially entirely based on deterministic conditional probability tables representing Bayesian “or” and “and” functions.

The project pertained to holding 10-person expert panels in “Annual Species Reviews” to determine the appropriate management category for each of dozens of plant or animal species …

Northwest Forest Plan

- Guidelines for managing late-successional and old-growth forest-related species within the range of the Northern Spotted Owl

… under the "Northwest Forest Plan," a multi-agency forest management and conservation plan covering western Washington, western Oregon, and northwestern California, in the United States. The Northwest Forest Plan was established to conserve the Northern Spotted Owl, salmonid fishes, and the full array of species associated with old-growth conifer forests of the region.

Northwest Forest Plan

- Annual Species Review: expert panels to determine species conservation status

Part of the regulations under the plan entailed holding "Annual Species Reviews" by a panel of 5 biologists and 5 managers (10 total) to advise the decision-makers on the most appropriate conservation and management status of each species.

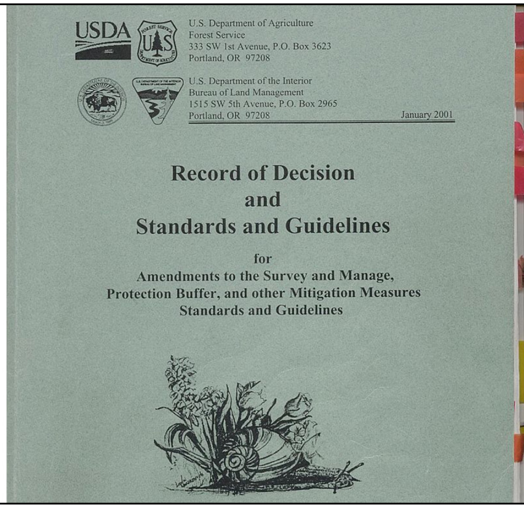

U.S. Department of Agriculture
Forest Service
333 SW 1st Avenue, P.O. Box 3623
Portland, OR 97208

U.S. Department of the Interior
Bureau of Land Management
1515 SW 5th Avenue, P.O. Box 2965
Portland, OR 97208

January 2001

Record of Decision
and
Standards and Guidelines

for
Amendments to the Survey and Manage,
Protection Buffer, and other Mitigation Measures
Standards and Guidelines

The criteria by which the panel was held and evaluated the status of each species were provided in (complex) details in a published regulation.

Decision Criteria

1. Geographic range
2. Old growth association
3. Plan provides for persistence
4. Data sufficiency for management
5. Survey practicality
6. Species rarity

The criteria for evaluating each species' status fell under 6 main headings shown here.

Survey and Manage Categories

Redefine Categories Based on Species Characteristics			
Relative Rarity	**Pre-Disturbance Surveys Practical**	**Pre-Disturbance Surveys Not Practical**	**Status Undetermined**
Rare	**Category A** - 57 species • Manage All Known Sites • Pre-Disturbance Surveys • Strategic Surveys	**Category B** - 222 species • Manage All Known Sites • N/A • Strategic Surveys	**Category E** - 22 species • Manage All Known Sites • N/A • Strategic Surveys
Uncommon	**Category C** - 10 species • Manage High-Priority Sites • Pre-Disturbance Surveys • Strategic Surveys	**Category D** - 14 species[1] • Manage High-Priority Sites • N/A • Strategic Surveys	**Category F** - 21 species • N/A • N/A • Strategic Surveys

[1] Includes three species for which pre-disturbance surveys are not necessary

Determining the conservation category for a species … depends on decision criteria in the Record of Decision

The outcome would be for the panel to recommend which of 7 categories each species should be placed into – the 7 being the Categories A through F shown above, plus one for potentially removing the species entirely from the list. Each category depending on the species' relative rarity among many other of the 6 criteria in the preceding slide, and each category carried different implications for follow-up site management and for the extent and type of surveys for the species.

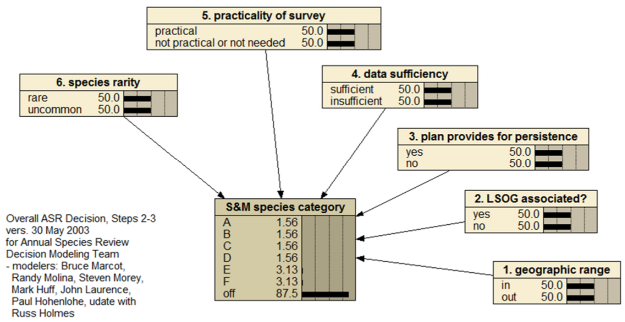

The main Bayesian network model looked like this.

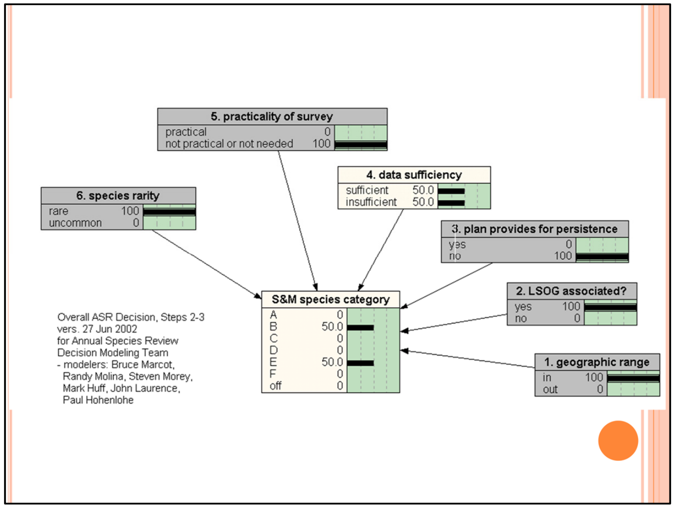

Here's an example of one species for which the Criteria 1, 2, 3, 5, and 6 were judged by the panel, with criterion 4 unknown so in the model that node was left to its default prior probability distribution of complete uncertainty (uniform probability distribution).

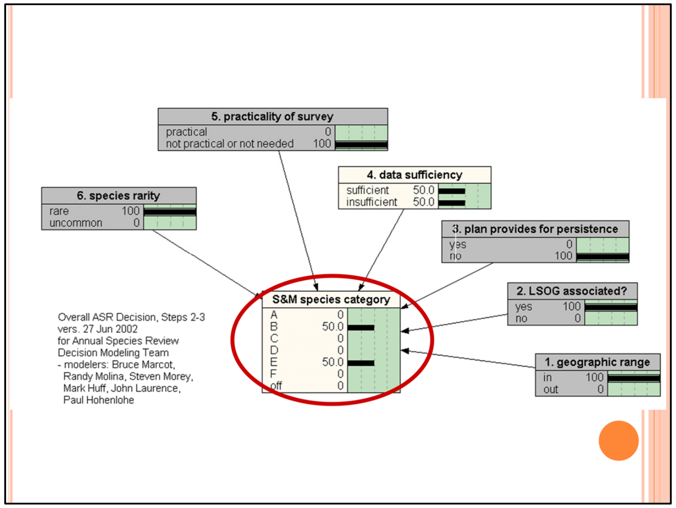

The outcome in this example here was that this particular species might fit under Category B or E.

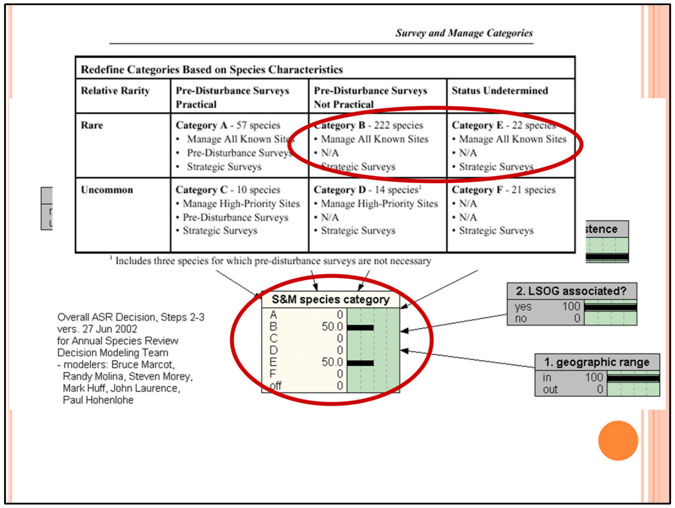

Survey and Manage Categories

Redefine Categories Based on Species Characteristics			
Relative Rarity	**Pre-Disturbance Surveys Practical**	**Pre-Disturbance Surveys Not Practical**	**Status Undetermined**
Rare	**Category A** - 57 species • Manage All Known Sites • Pre-Disturbance Surveys • Strategic Surveys	**Category B** - 222 species • Manage All Known Sites • N/A • Strategic Surveys	**Category E** - 22 species • Manage All Known Sites • N/A • Strategic Surveys
Uncommon	**Category C** - 10 species • Manage High-Priority Sites • Pre-Disturbance Surveys • Strategic Surveys	**Category D** - 14 species[1] • Manage High-Priority Sites • N/A • Strategic Surveys	**Category F** - 21 species • N/A • N/A • Strategic Surveys

[1] Includes three species for which pre-disturbance surveys are not necessary

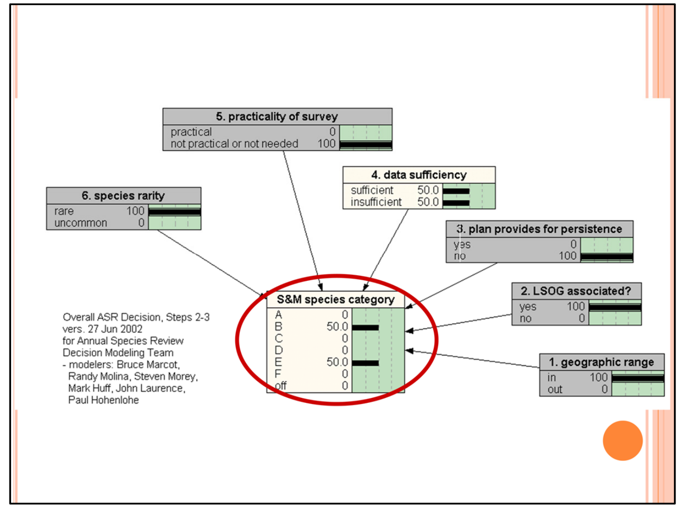

Again, this model was built with deterministic Boolean logic-based conditional probability tables …

conditional probability table

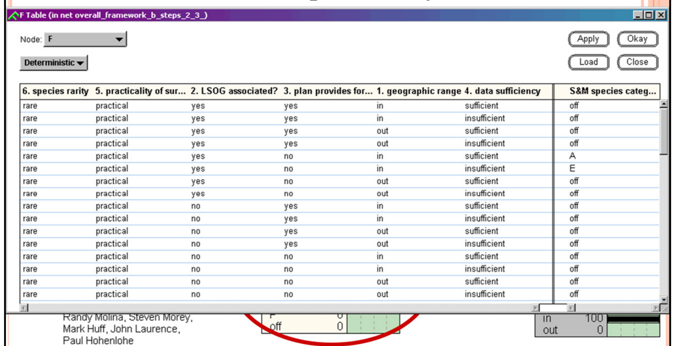

F Table (in net overall_framework_b_steps_2_3_)

Node: F

Deterministic

Apply Okay Load Close

6. species rarity	5. practicality of sur...	2. LSOG associated?	3. plan provides for...	1. geographic range	4. data sufficiency	S&M species categ...
rare	practical	yes	yes	in	sufficient	off
rare	practical	yes	yes	in	insufficient	off
rare	practical	yes	yes	out	sufficient	off
rare	practical	yes	yes	out	insufficient	off
rare	practical	yes	no	in	sufficient	A
rare	practical	yes	no	in	insufficient	E
rare	practical	yes	no	out	sufficient	off
rare	practical	yes	no	out	insufficient	off
rare	practical	no	yes	in	sufficient	off
rare	practical	no	yes	in	insufficient	off
rare	practical	no	yes	out	sufficient	off
rare	practical	no	yes	out	insufficient	off
rare	practical	no	no	in	sufficient	off
rare	practical	no	no	in	insufficient	off
rare	practical	no	no	out	sufficient	off
rare	practical	no	no	out	insufficient	off

… such as shown here, in part, for that final output node.

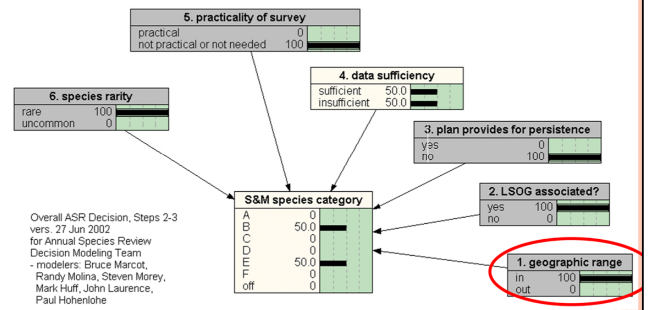

Now, each input node was actually the result of a submodel representing the published guidelines. For example, category 1 “geographic range” …

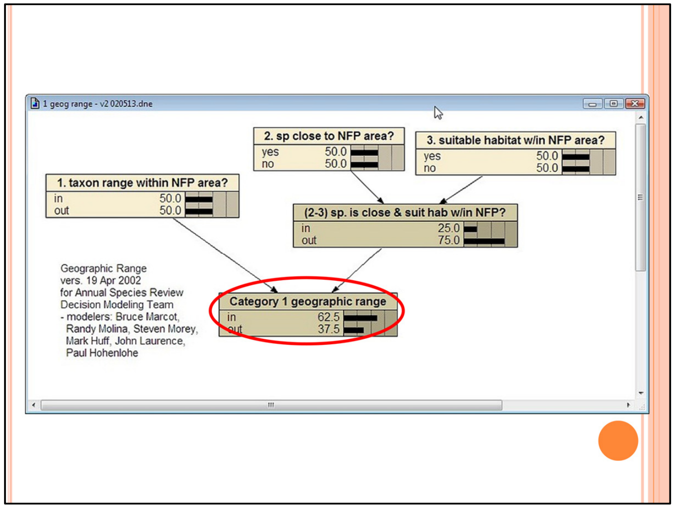

… is actually the result of its submodel. Here, nodes 1, 2, and 3 at the top are the three criteria explicitly listed in the published regulation guidelines for this particular category, by which to determine if a species is within or out of the range of the Northwest Forest Plan.

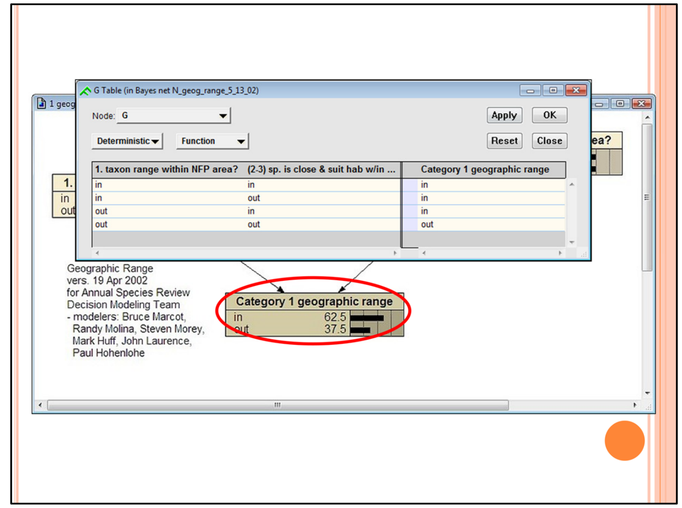

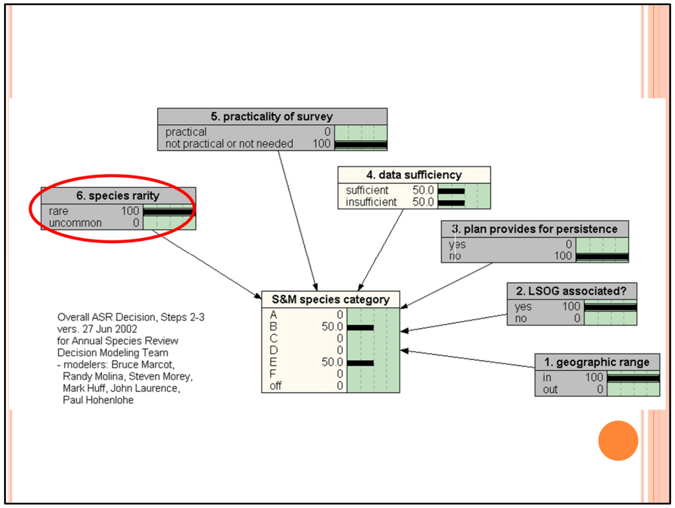

Some of the categories were very complicated, such as the one for ascertaining species rarity …

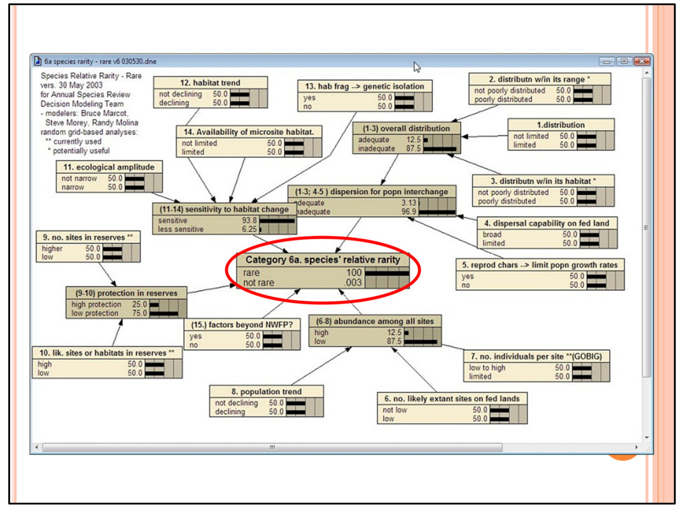

… which entailed specifying some 15 different factors according to the regulation guidelines. Here is the submodel for species rarity, using deterministic Boolean logic-based conditional probability tables throughout.

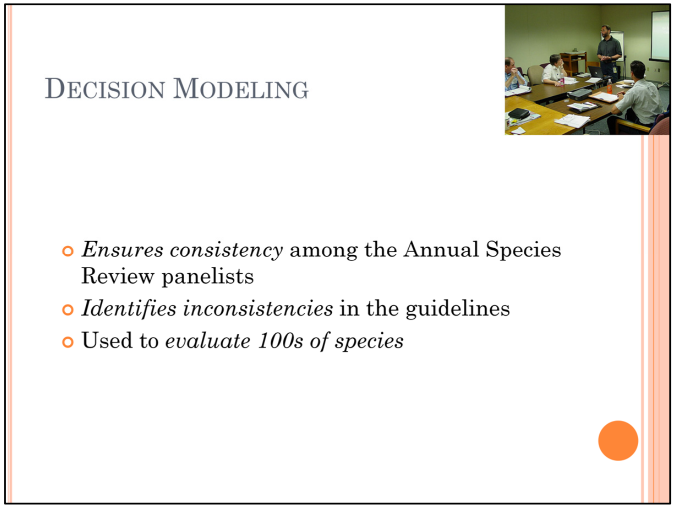

So this Bayesian network model was very helpful in guiding the Annual Species Review panels – who met annually for a number of years running – to navigate through the complex web of regulation guidelines. Use of the model ensured consistency among the panelists, and helped identify some key (but inadvertent) *inconsistencies* hiding in the regulation guidelines and to deal with them in a clear, open manner. The Annual Species Review panels used the model to evaluate hundreds of species, appropriately as a *decision-aiding*, not a decision-making, tool.

Source: Marcot, B. G., P. A. Hohenlohe, S. Morey, R. Holmes, R. Molina, M. Turley, M. Huff, and J. Laurence. 2006. Characterizing species at risk II: using Bayesian belief networks as decision support tools to determine species conservation categories under the Northwest Forest Plan. Ecology and Society 11(2):12. [online] URL: http://www.ecologyandsociety.org/vol11/iss2/art12/.

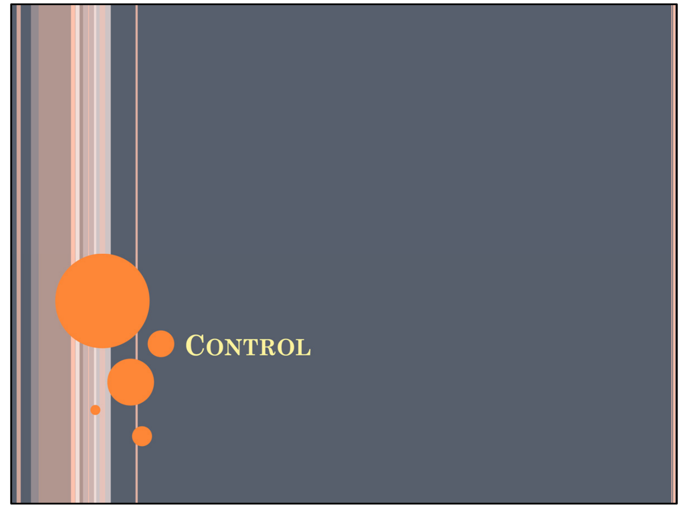

OK, let's move on to our second major consideration in Bayesian network modeling in management …

CONFIDENCE – How to denote confidence in expert judgment used to develop probability structures?

CONTROL – How to identify management control and influence of decisions?

CAUSE – How to determine causality?

… the issue of *control.*

"There's no time for thinking. We have to make a *management decision*."

Key Questions About Control

- How do you know what management can control?
- What are realistic expectations for the degree of management control?

Goes to the heart of inferring causality.

Here are some key questions about management control.

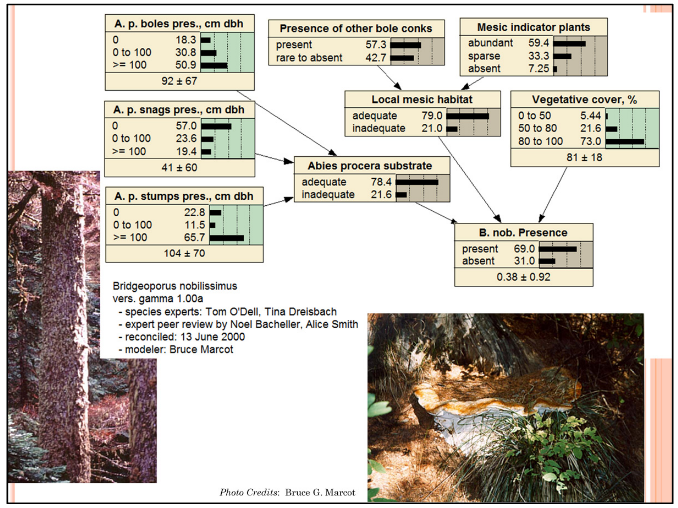

Photo Credits: Bruce G. Marcot

This is a model I built, with mycologist expert input, that predicts habitat suitability and presence of a rare bracket fungus.

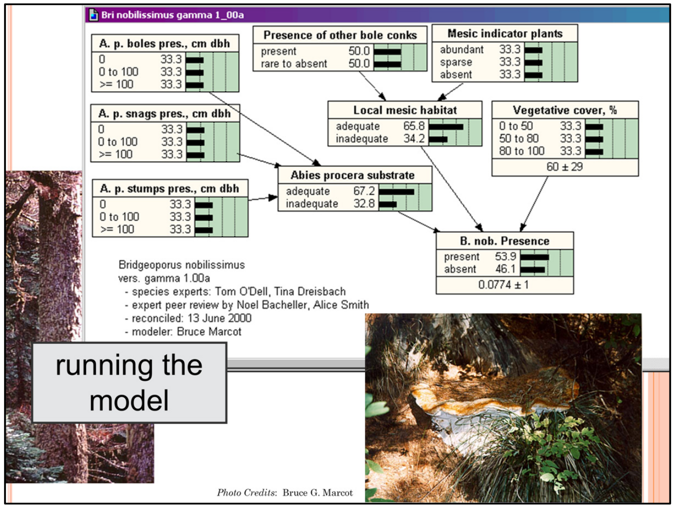

The model is run by specifying values of the input environmental conditions.

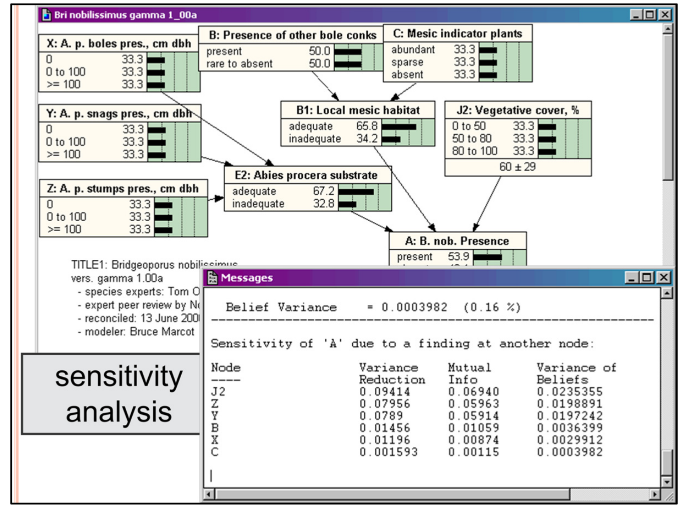

So what about management “control?” What can this model tell us about what parts of forest structures are more important for this species? One answer lies in conducting sensitivity analysis of the model …

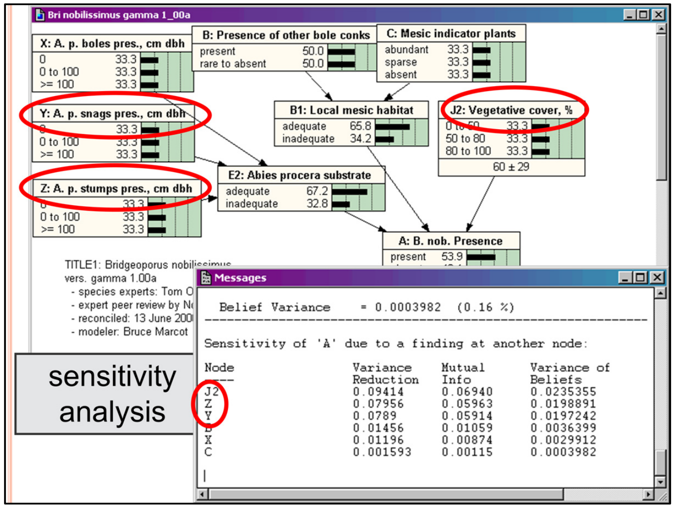

… that has identified that vegetative cover, and large-diameter stumps and snags (standing dead trees) of noble fir, contribute the most to habitat suitability and presence of this rare bracket fungus species.

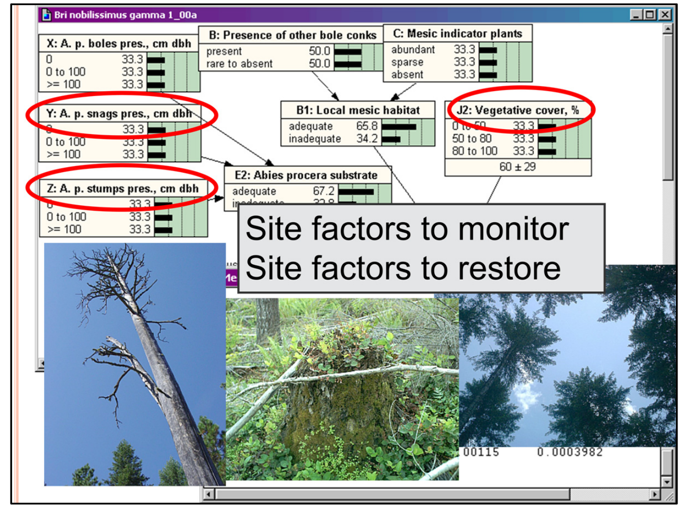

So these three forest stand attributes could the key ones that management might want to monitor, and possibly to restore to create habitat for the species. They are ones also that management could *control* through proper forestry silvicultural prescriptions.

Management and Causality

- *Causality* is central to the management question of controllability.
- What is the cheapest and easiest way to *control* some environmental condition for a desired outcome?
- Use a Bayesian network to determine effect of controllability by use of "*influence runs*."
  - * examples

In this example, it is assumed that these forest stand conditions *cause* the habitat to be suitable for this bracket fungus species. The question of *causality* is central to management controllability. It raises these further questions, as well. And the use of Bayesian network models to conduct what I term "influence runs." Let's look at that for a minute…

sensitivity analysis
v.
influence runs

I differentiate between "sensitivity analysis" and "influence runs" in Bayesian network modeling.

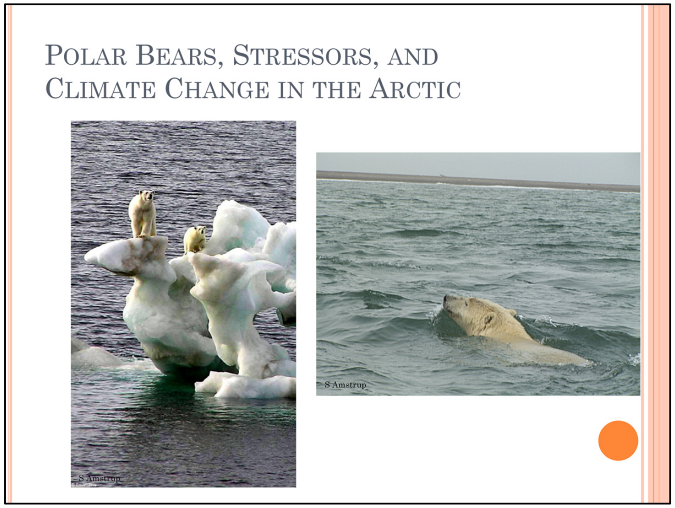

Here's an example with the polar bear model.

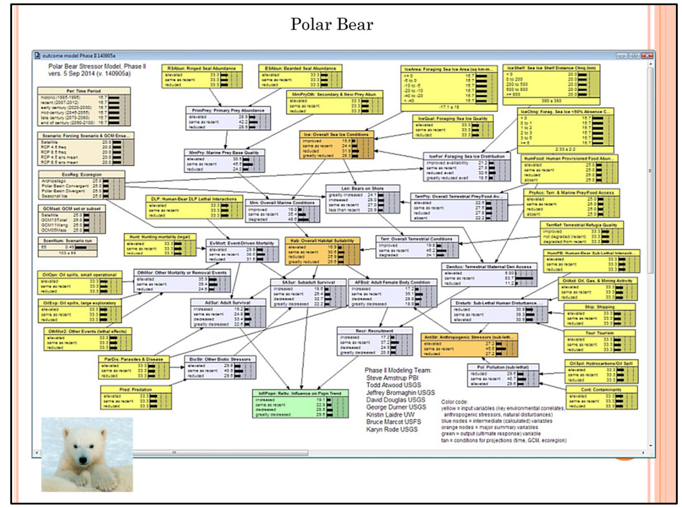

This is the polar bear model mentioned previously.

Source: Atwood, T. C., B. G. Marcot, D. C. Douglas, S. C. Amstrup, K. D. Rode, G. M. Durner, and J. F. Bromaghin. 2015. Evaluating and ranking threats to the long-term persistence of polar bears. U.S. Geological Survey, Open-File Report 2014-1254. http://dx.doi.org/10.3133/ofr20141254. Anchorage, Alaska. 114 pp.

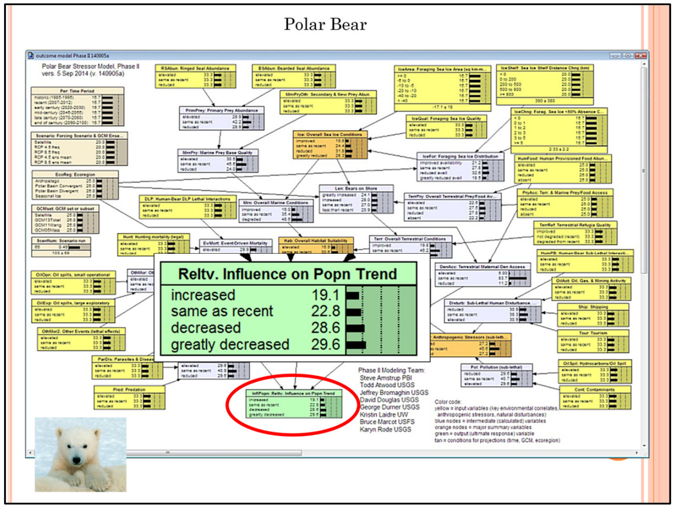

The final output node here depicts the relative influence of climate change on Arctic sea ice, and of other environmental and anthropogenic stressors, on polar bear population trend.

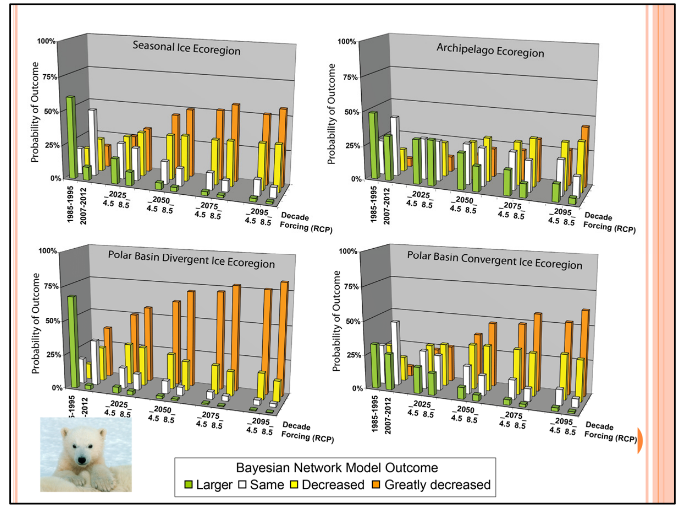

Running the model for each of four Arctic ecoregions, under two climate change scenarios and 6 time period, resulted in probability projections for polar bear population trends, as shown here.

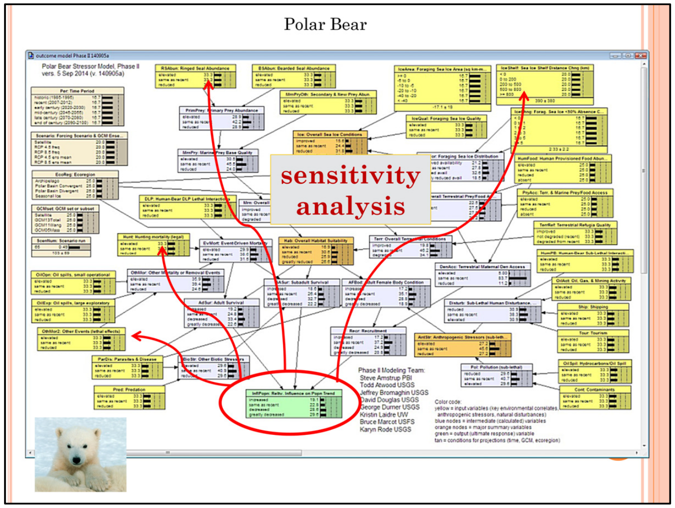

Now, “sensitivity analysis” is run first by setting all inputs (the yellow nodes) to their default prior probability distributions.

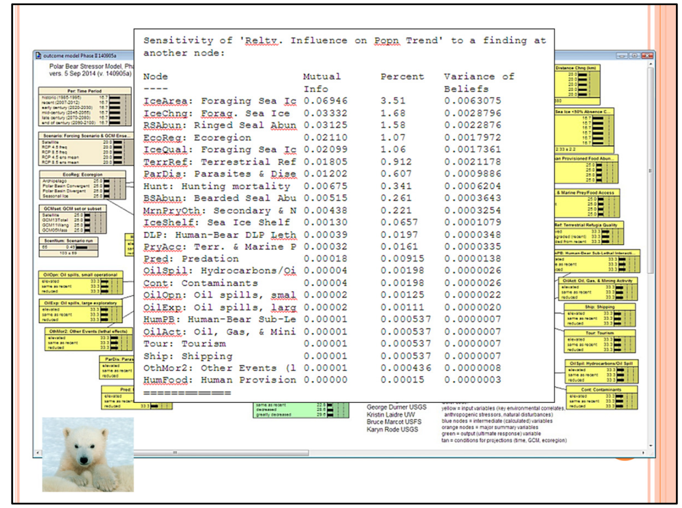

Sensitivity of 'Reltv. Influence on Popn Trend' to a finding at another node:

Node	Mutual Info	Percent	Variance of Beliefs
IceArea: Foraging Sea Ic	0.06946	3.51	0.0063075
IceChng: Forag. Sea Ice	0.03332	1.68	0.0028796
RSAbun: Ringed Seal Abun	0.03125	1.58	0.0022876
EcoReg: Ecoregion	0.02110	1.07	0.0017972
IceQual: Foraging Sea Ic	0.02099	1.06	0.0017361
TerrRef: Terrestrial Ref	0.01805	0.912	0.0021178
ParDis: Parasites & Dise	0.01202	0.607	0.0009886
Hunt: Hunting mortality	0.00675	0.341	0.0006204
BSAbun: Bearded Seal Abu	0.00515	0.261	0.0003643
MrnPryOth: Secondary & N	0.00438	0.221	0.0003254
IceShelf: Sea Ice Shelf	0.00130	0.0657	0.0001079
DLP: Human-Bear DLP Leth	0.00039	0.0197	0.0000348
PryAcc: Terr. & Marine P	0.00032	0.0161	0.0000335
Pred: Predation	0.00018	0.00915	0.0000138
OilSpil: Hydrocarbons/Oi	0.00004	0.00198	0.0000026
Cont: Contaminants	0.00004	0.00198	0.0000026
OilOpn: Oil spills, smal	0.00002	0.00125	0.0000022
OilExp: Oil spills, larg	0.00002	0.00111	0.0000020
HumPB: Human-Bear Sub-Le	0.00001	0.000537	0.0000007
OilAct: Oil, Gas, & Mini	0.00001	0.000537	0.0000007
Tour: Tourism	0.00001	0.000537	0.0000007
Ship: Shipping	0.00001	0.000537	0.0000007
OthMor2: Other Events (l	0.00001	0.000436	0.0000008
HumFood: Human Provision	0.00000	0.00015	0.0000003

============

The result of such a sensitivity analysis is essentially a statement about the underlying probability structure of the model …

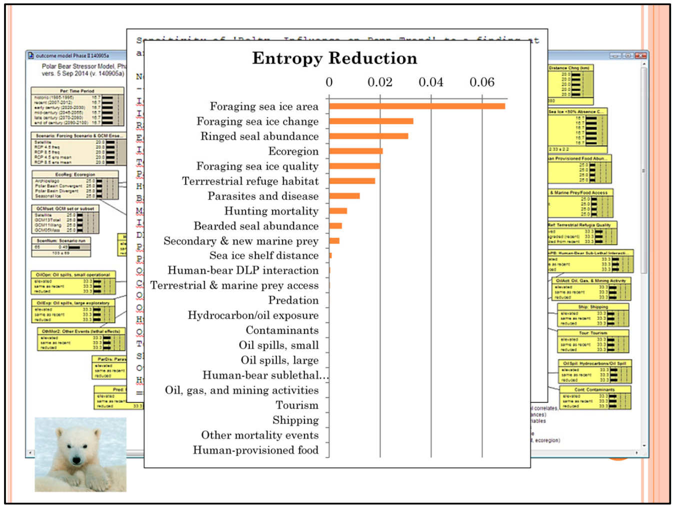

… here depicted more clearly in a bar graph. In this case, it seems that most of the model sensitivity pertains to variables related to climate, sea ice, and their effects on key prey species of the polar bear.

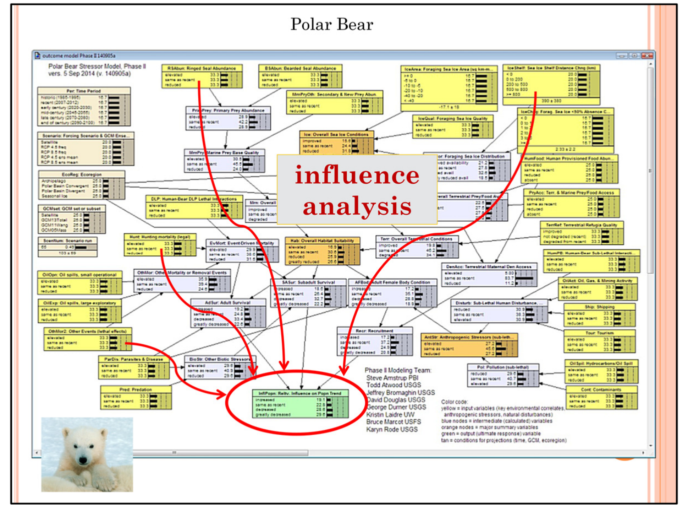

In comparison, what I term an "influence analysis" pertains to determining the influence of one or more inputs on model outcome run under a specified scenario, not under default prior probabilities of the input nodes …

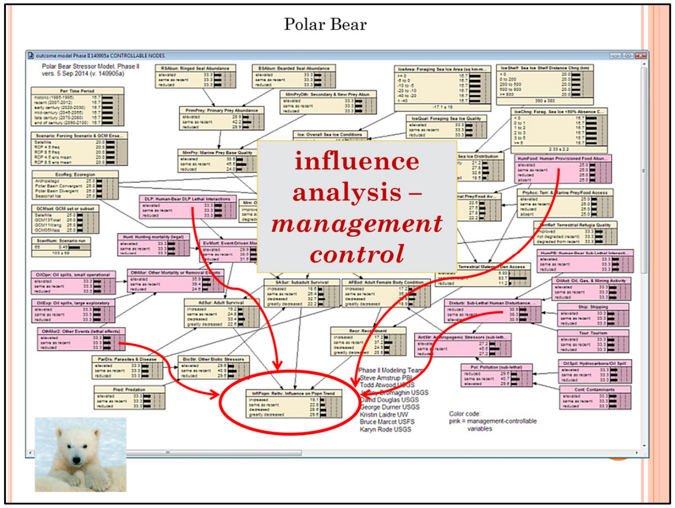

… and by this manner, influence runs can determine the degree to which varying a given input factor can affect the overall outcome. This is then related to the degree of management control on that outcome.

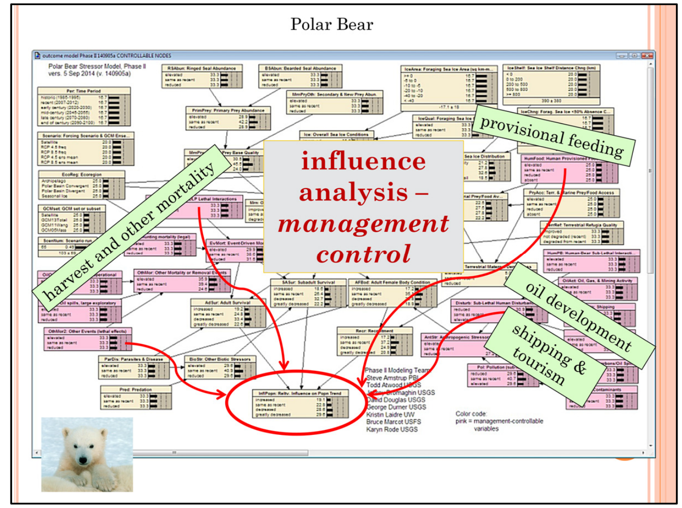

In the polar bear model, sets of inputs were each varied to determine the degree to which each set could help improve (or, contrariwise, further degrade) polar bear futures.

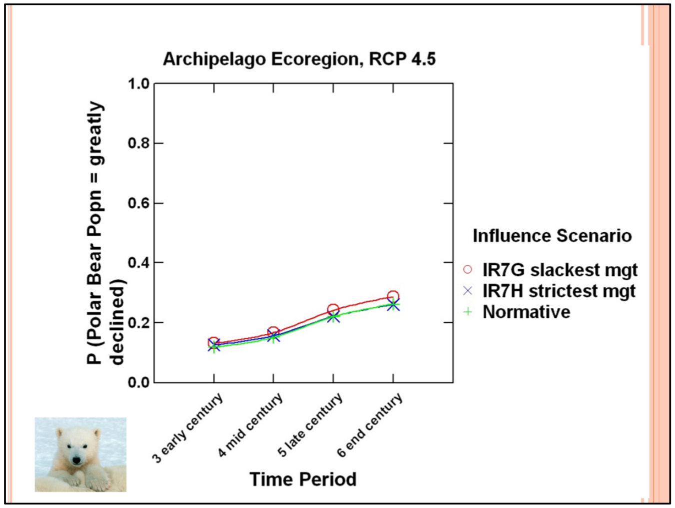

I plotted the results of such influence runs this way. The “normative” or expected outcomes under future scenarios is shown here in green, and the other lines depict results when specific sets of inputs (that theoretically could be controlled by management) were set to their worst or best conditions.

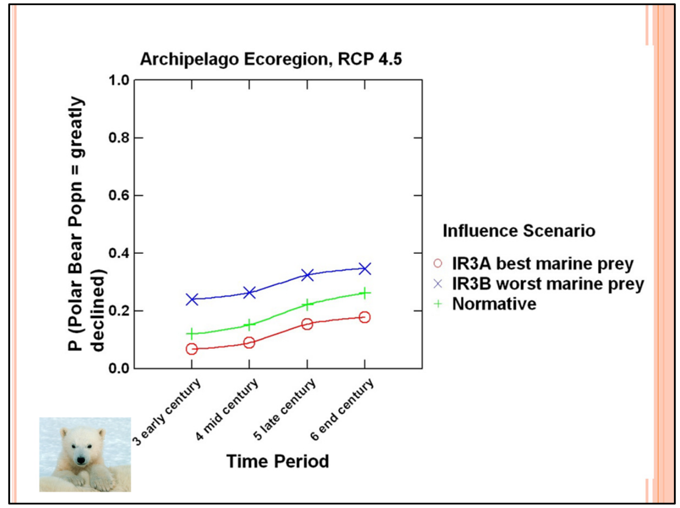

For example, with all else being held to “normative” (expected) future conditions for this ecoregion under this climate change scenario, if marine prey of the polar bear (including marine seals) were to achieve best possible states (the red line), then the probability of future polar bear populations being greatly declined could be lessened by as much as about 10 percent by the end of the century; this could be a significant conservation outcome. If, however, marine prey were to crash, then the probability of polar bears being greatly declined would increase by 10 percent or more. This shows the *influence* of this factor (marine prey) on polar bear outcomes.

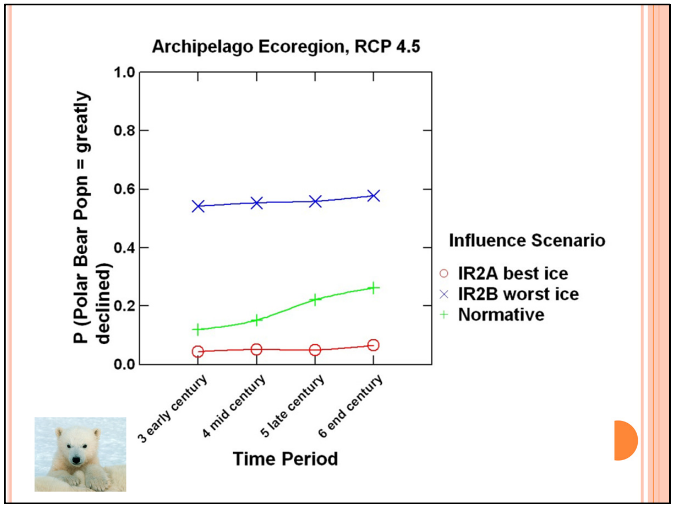

Conducting similar influence runs for the future of Arctic sea ice shows the most extreme *influence* outcome.

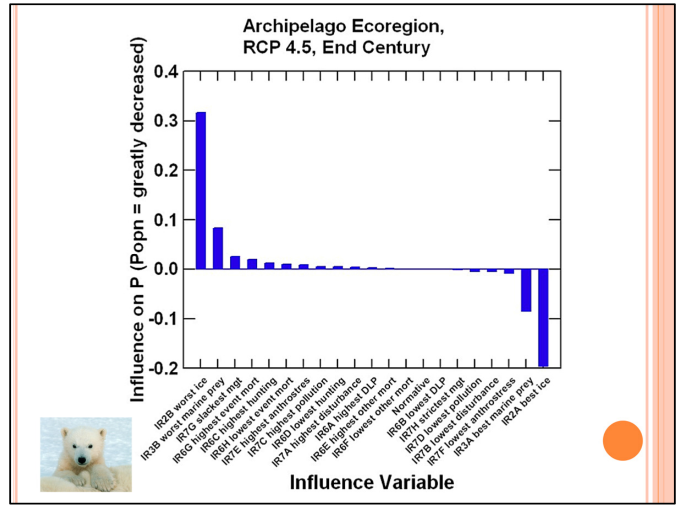

I summarized all influence runs this way, in a form of "tornado diagram" as it's called in the literature. This plot rank-orders the degree of influence – bad on the left (with the bars going up), good on the right (with the bars going down) – to clearly show their relative differences. It's clear from this plot that sea ice and marine prey can have the greatest *influence* on future polar bear populations. Other factors could contribute as well, but not nearly as much. So this provide management with some clear expectations of outcomes should they choose to control certain aspects of the polar bear's environment.

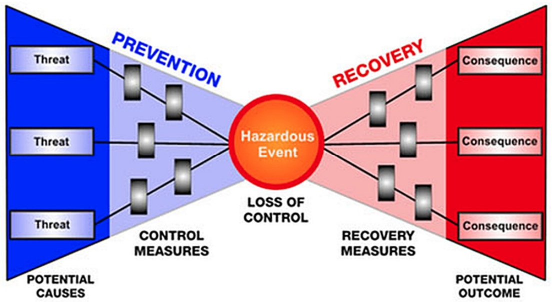

In environmental engineering, such influence results are depicted in a "bow-tie digram" as generalize here.

Structured Decision-Making to Evaluate Species at Risk

General format of risk analysis depictions of BN model results:

Probability		Outcome: Larger	Same as now	Smaller	Rare	Extinct
0.80-1.00	Almost certain	M	H	H	E	E
0.60-0.79	Likely	M	M	H	H	E
0.40-0.59	Coin toss	L	M	M	H	E
0.20-0.39	Unlikely	L	M	M	M	H
0.00-0.19	Rare	L	L	M	M	H

KEY:
L = low risk
M = moderate risk
H = high risk
E = extreme risk

d = GCM minimum
m = Ensemble mean
u = GCM maximum

The results of influence runs and scenario modeling in Bayesian networks can be then put into a decision risk-management framework shown here. "Risk" is essentially the probability of an outcome *coupled* with the condition or utility of that outcome. For example, here, extreme risk results from very high probability of a species such as the polar bear becoming rare or extinct.

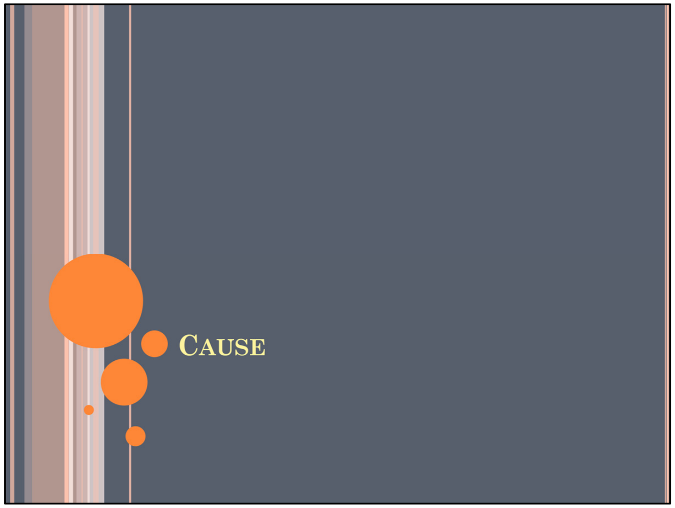

Now let's address the third of the modeling considerations here: cause.

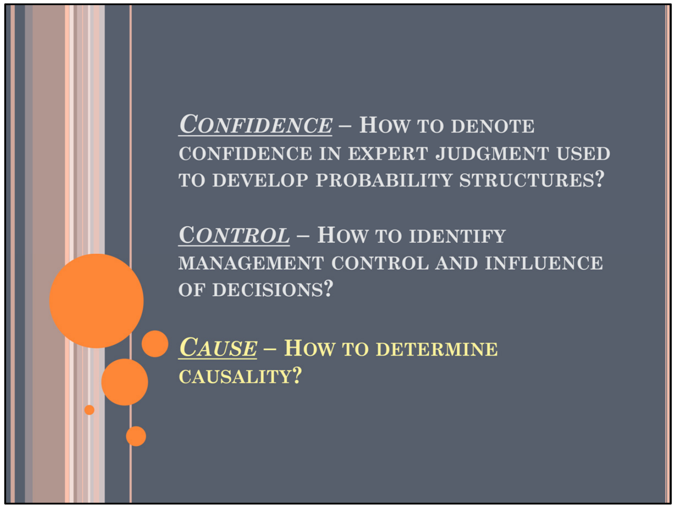

Let's ponder how Bayesian network models aid in determining *causality.*

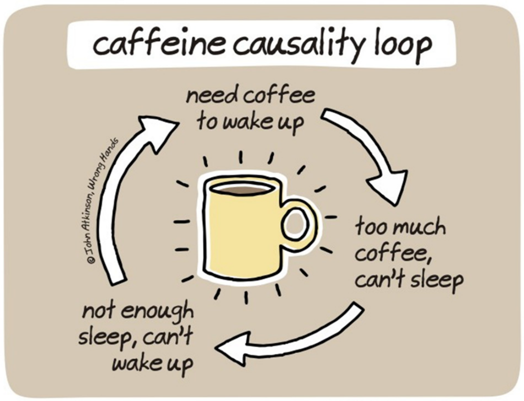

First, here is the causal loop in which I find myself often stuck … but this aside …

Key Points About Causality

- **No model – SEM, path regression, probability networks, Bayesian networks – can reveal causality.**

… it is a truism that no model can reveal *causality.*

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

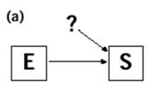

Causal mechanisms can be unexpectedly complicated. For instance, consider that we want to understand how some environmental condition "E" affects some species outcome "S." There will always be some unexplained variation "?" in this simple depiction.

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

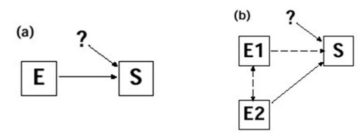

In a more realistic depiction, the true *causal* environmental variable E2 could be hidden from our research but could be *correlated* with the variable we do observe, E1.

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

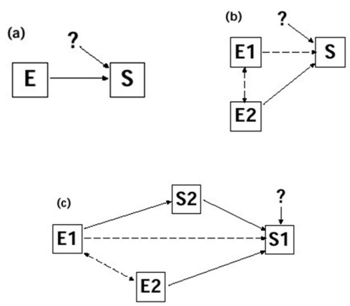

Even more complicated could be another species S2 that somehow affects our target observed species S1, such as by S2 being a key predator on S1, or some other ecological relationship.

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

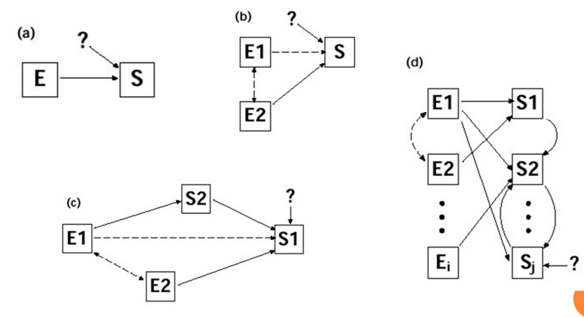

It can get complicated quickly in real-world ecological *communities* consisting of many environmental factors and many species. Also, notice how we are starting to see a repeating pattern in this expanded causal web.

Key Ecological Functions of Organisms

Here's a real-world example … what I call "key ecological functions" of wildlife organisms. Key ecological functions refer to the active ecological "roles" that organisms play in their ecological communities.

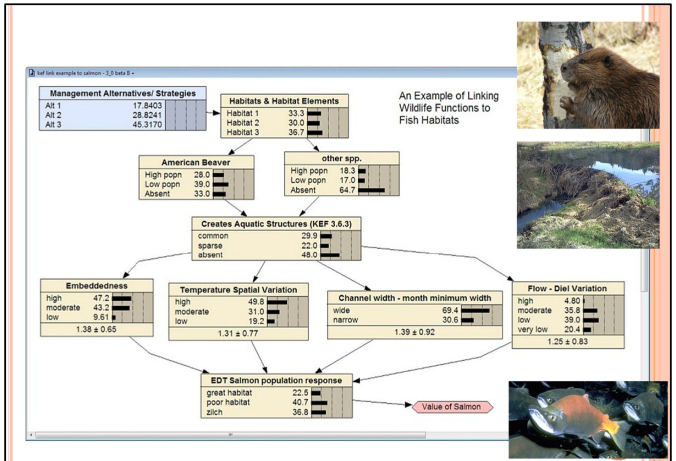

Photo Credits: Dennis Garrison, Gunnison National Forest; Bruce G. Marcot; US Fish and Wildlife Service

In this example is the beaver (among other species), who fells trees to dam streams that alters characteristics of the water and that creates pond habitats used by a wide variety of other species, here shown as salmon.

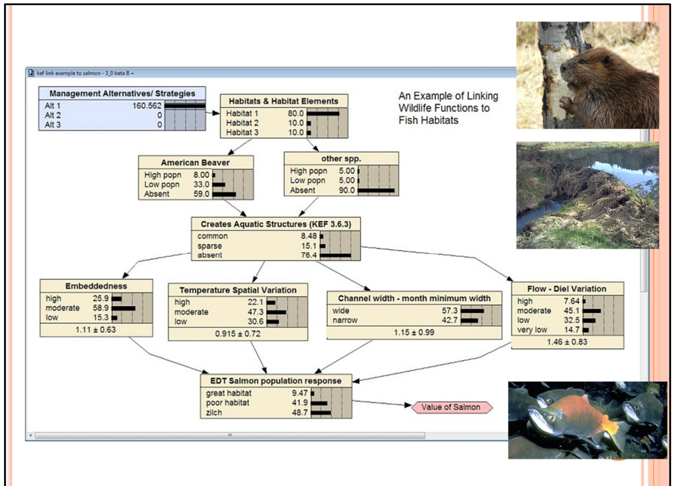

Photo Credits: Dennis Garrison, Gunnison National Forest; Bruce G. Marcot; US Fish and Wildlife Service

In this model, different management alternatives or strategies (in the blue box) can have different affects on habitat elements in turn affecting presence of beavers and other dam-creating species (e.g., nutria) … ultimately affecting salmon population response. The point here is that *causality* in this system is not a direct link from the management activity (or its initial affect on habitats) to salmon response, but instead travels through other species, their key ecological functions, and system responses.

Key Points About Causality

- **Causal mechanisms can be insidiously complex:**
  - **Environmental condition → Species response, ? = residuals**

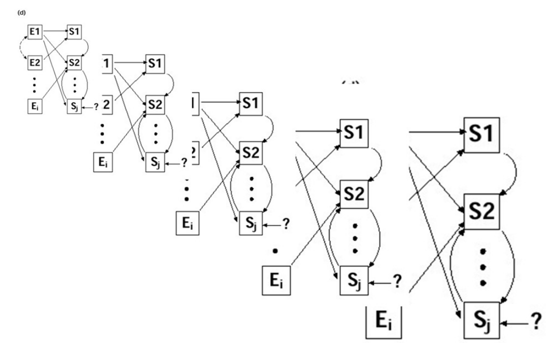

There seems to be a self-similarity of patterns …

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

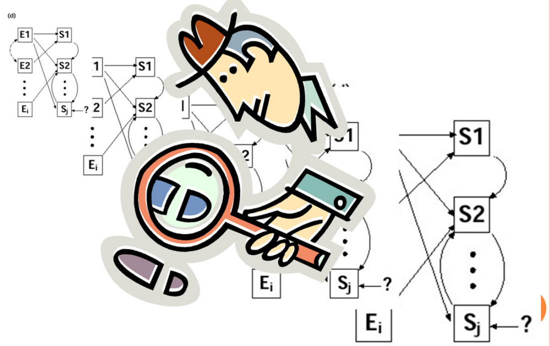

… the deeper you probe into ecological systems …

Key Points About Causality

- Causal mechanisms can be insidiously complex:
  - Environmental condition → Species response, ? = residuals

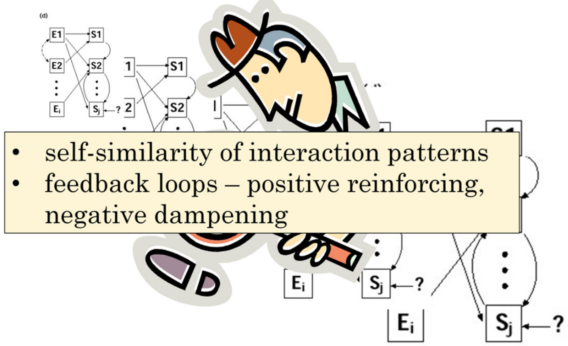

- self-similarity of interaction patterns
- feedback loops – positive reinforcing, negative dampening

… patterns of which result from various kinds of feedback loops and causal connections and influences …

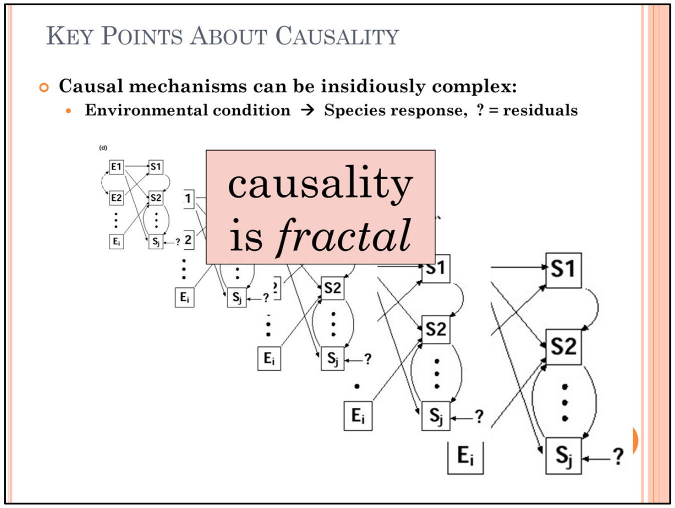

… which had led me to postulate that *causality is fractal* …

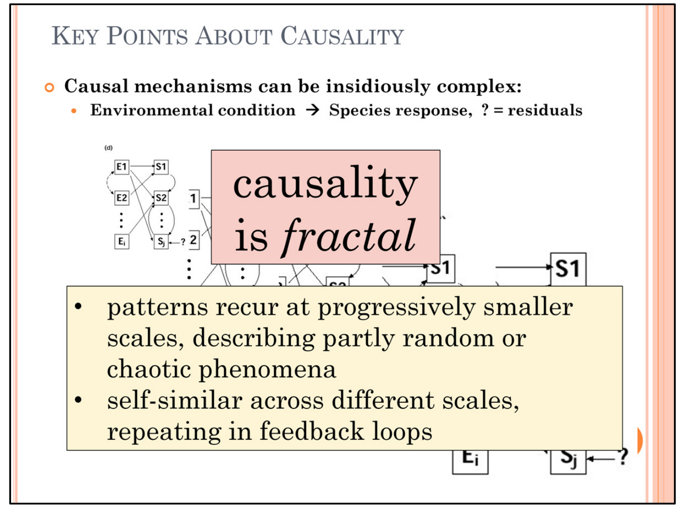

… in that such patterns recur at various scale and appear more or less self-similar.

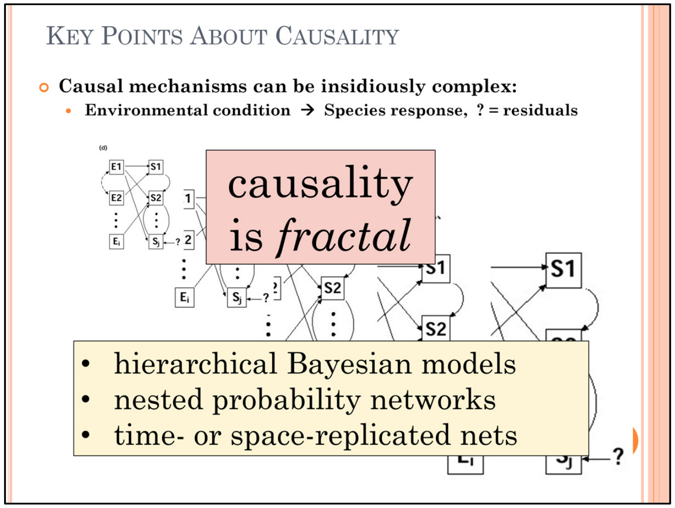

Such fractal patterns could be modeled various ways in Bayesian networks.

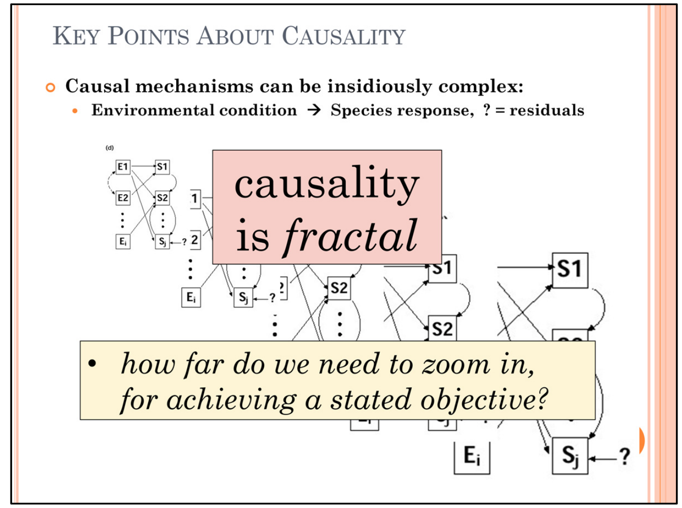

One key question is, to achieve a stated objective, particularly for management control, how much detail is needed in this self-similar, nested causal web structure? How far do we need to zoom in?

Bayesian Networks

- Acyclic
- Directed
- Markovian

Remember that Bayesian networks are acyclic digraphs … that are essentially Markovian in structure, that is, each node depends only its immediate parent nodes.

Bayesian Networks

- Acyclic
- Directed
- Markovian

How to represent non-Markovian processes?

So one question is, how to represent *non*-Markovian processes in a Bayesian network model? These could be important for depicting the complexity of causal influences that could be latent (hidden) or could be operating across time steps, geographic scales, or nested functions such as shown in the beaver example.

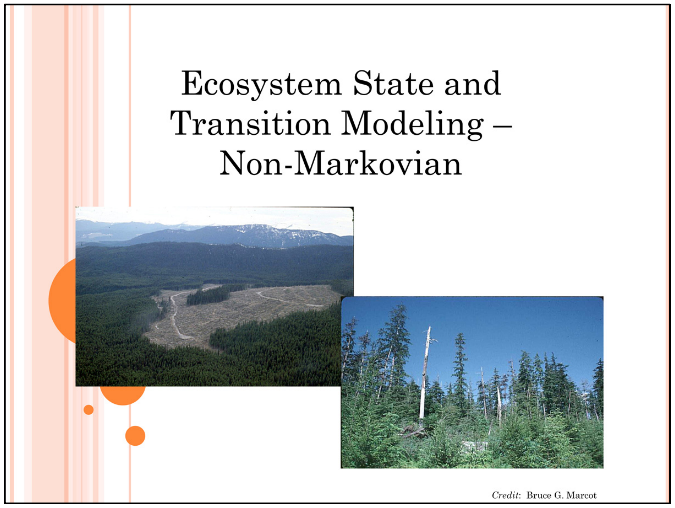

Here’s a simple “cheat” I developed for depicting non-Markovian processes in a Bayes net …

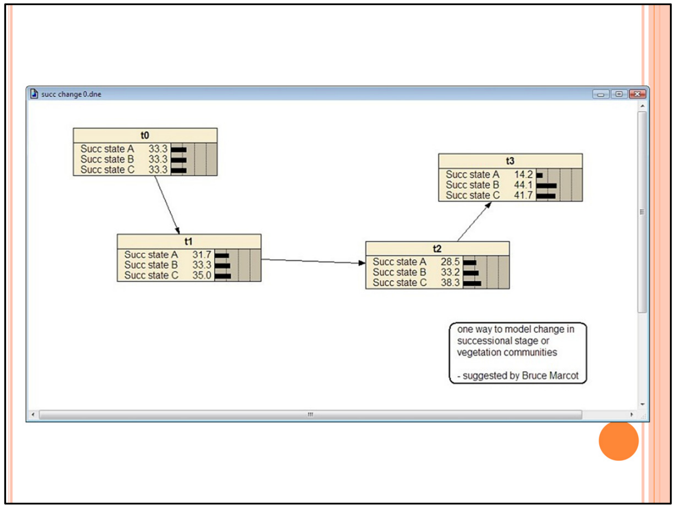

The project at hand was to project future successional states of a vegetation community based on transition probabilities across time periods from the present (t0) into the future (t1, t2, t3). In this model, the transitions occur independently and successively across the time periods in a classic Markovian structure where conditions in each time period is influenced ONLY by conditions in the immediately preceding time period.

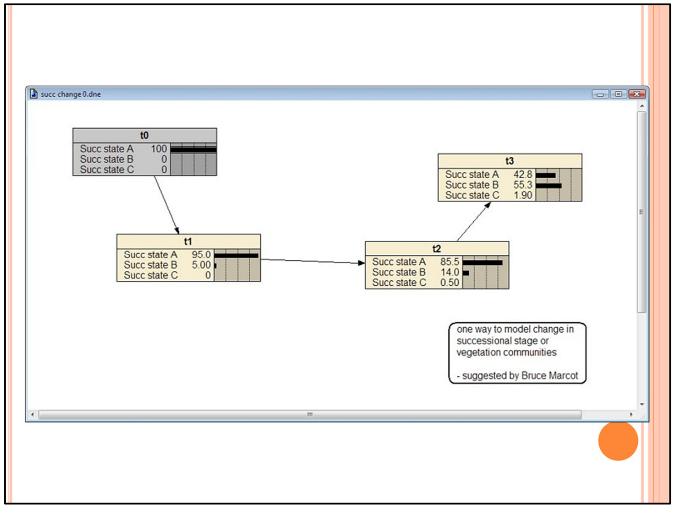

For instance, in the present time period t0, let's say that the vegetation community is 100% in successional state A. This model projects its transitions into future states as shown here.

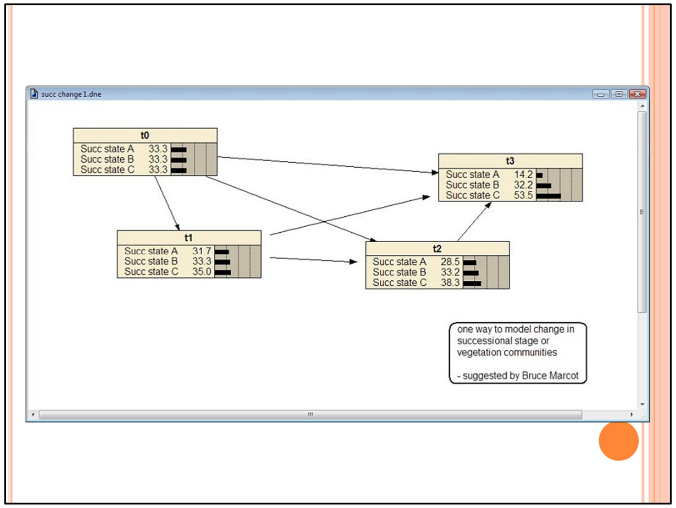

So to "cheat" the Markovian process, one can link successive future time periods to all the previous time periods.

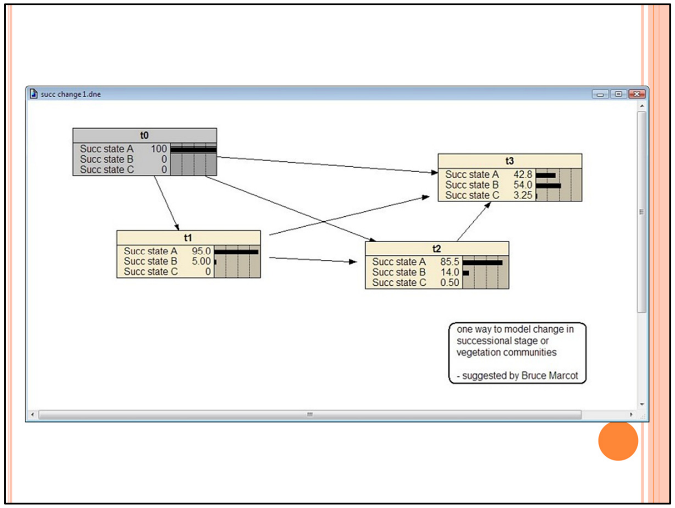

In this way, future vegetation conditions can be influences by conditions in the immediately recent time period AND by those in previous periods as well. Why is this useful?

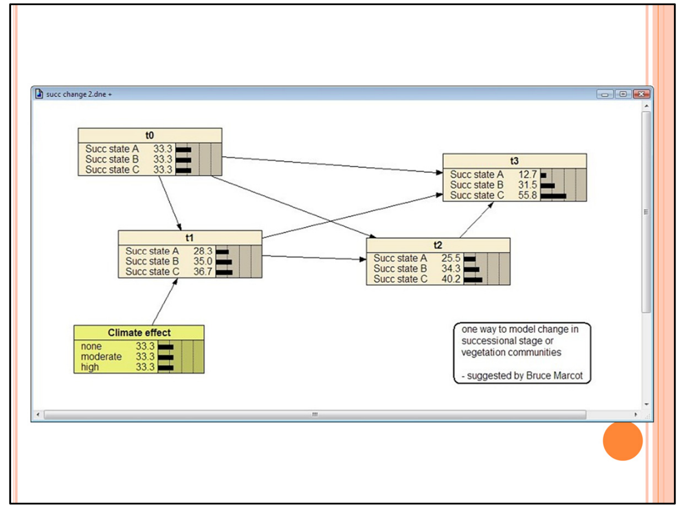

Well, consider if we introduce some stressor such as climate effect …

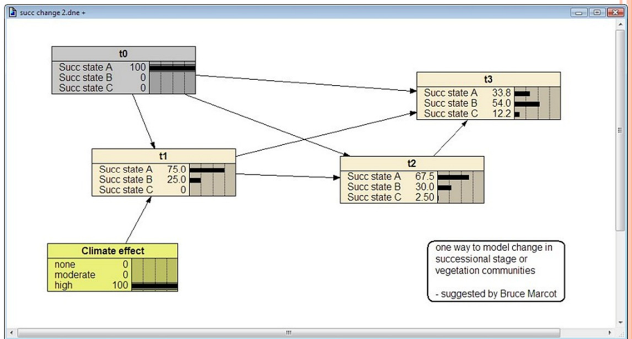

… that can change the course of vegetation succession at a particular future time period. Its influence might play out more strongly over time in such a (pseudo) "non-Markovian" structure as this one.

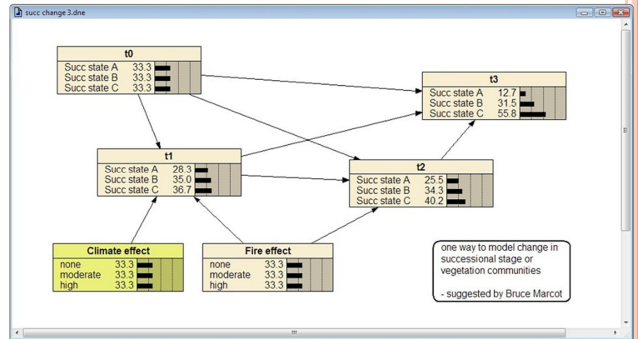

One can also introduce other stressors that could have different and additional influences, as well.

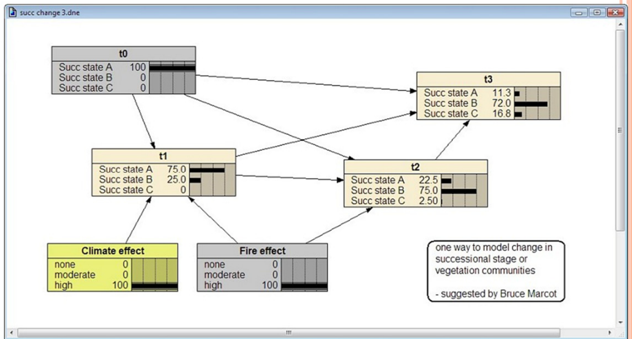

So, although this Bayesian network really is still Markovian (each node is influence only by its directly-linked parents), it is functional non-Markovian in terms of the influence of time periods.

Probability vs. Likelihood

I want to briefly note the difference between probability and likelihood in a Bayesian network context, as this has implications for thinking about causality.

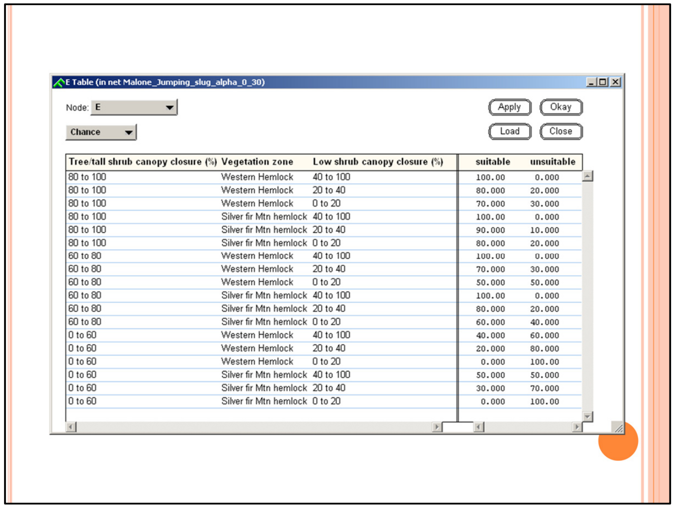

E Table (in net Malone_Jumping_slug_alpha_0_30)

Node: E

Chance

Apply Okay Load Close

Tree/tall shrub canopy closure (%)	Vegetation zone	Low shrub canopy closure (%)	suitable	unsuitable
80 to 100	Western Hemlock	40 to 100	100.00	0.000
80 to 100	Western Hemlock	20 to 40	80.000	20.000
80 to 100	Western Hemlock	0 to 20	70.000	30.000
80 to 100	Silver fir Mtn hemlock	40 to 100	100.00	0.000
80 to 100	Silver fir Mtn hemlock	20 to 40	90.000	10.000
80 to 100	Silver fir Mtn hemlock	0 to 20	80.000	20.000
60 to 80	Western Hemlock	40 to 100	100.00	0.000
60 to 80	Western Hemlock	20 to 40	70.000	30.000
60 to 80	Western Hemlock	0 to 20	50.000	50.000
60 to 80	Silver fir Mtn hemlock	40 to 100	100.00	0.000
60 to 80	Silver fir Mtn hemlock	20 to 40	80.000	20.000
60 to 80	Silver fir Mtn hemlock	0 to 20	60.000	40.000
0 to 60	Western Hemlock	40 to 100	40.000	60.000
0 to 60	Western Hemlock	20 to 40	20.000	80.000
0 to 60	Western Hemlock	0 to 20	0.000	100.00
0 to 60	Silver fir Mtn hemlock	40 to 100	50.000	50.000
0 to 60	Silver fir Mtn hemlock	20 to 40	30.000	70.000
0 to 60	Silver fir Mtn hemlock	0 to 20	0.000	100.00

As we know, a conditional probability table (CPT) in a Bayesian network model essentially lists all possible conditions of the parent node(s), shown on the left, and the outcome state probabilities of the child node, show on the right. Here's an example CPT from some model I built.

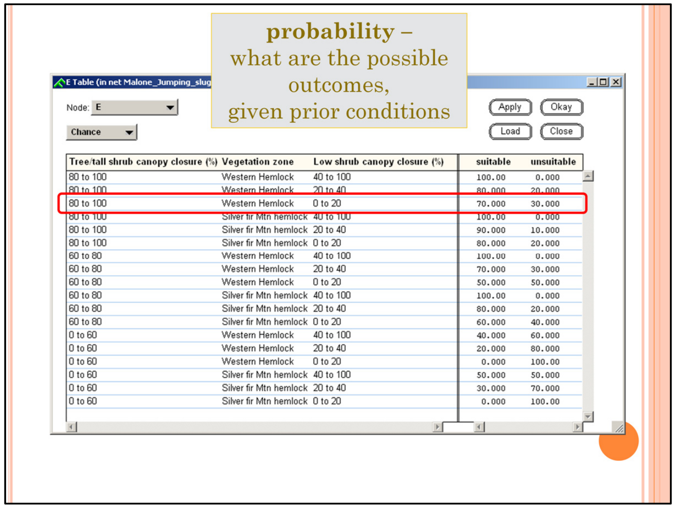

E Table (in net Malone_Jumping_slug

Node: E

Chance

Apply Okay Load Close

Tree/tall shrub canopy closure (%)	Vegetation zone	Low shrub canopy closure (%)	suitable	unsuitable
80 to 100	Western Hemlock	40 to 100	100.00	0.000
80 to 100	Western Hemlock	20 to 40	80.000	20.000
80 to 100	Western Hemlock	0 to 20	70.000	30.000
80 to 100	Silver fir Mtn hemlock	40 to 100	100.00	0.000
80 to 100	Silver fir Mtn hemlock	20 to 40	90.000	10.000
80 to 100	Silver fir Mtn hemlock	0 to 20	80.000	20.000
60 to 80	Western Hemlock	40 to 100	100.00	0.000
60 to 80	Western Hemlock	20 to 40	70.000	30.000
60 to 80	Western Hemlock	0 to 20	50.000	50.000
60 to 80	Silver fir Mtn hemlock	40 to 100	100.00	0.000
60 to 80	Silver fir Mtn hemlock	20 to 40	80.000	20.000
60 to 80	Silver fir Mtn hemlock	0 to 20	60.000	40.000
0 to 60	Western Hemlock	40 to 100	40.000	60.000
0 to 60	Western Hemlock	20 to 40	20.000	80.000
0 to 60	Western Hemlock	0 to 20	0.000	100.00
0 to 60	Silver fir Mtn hemlock	40 to 100	50.000	50.000
0 to 60	Silver fir Mtn hemlock	20 to 40	30.000	70.000
0 to 60	Silver fir Mtn hemlock	0 to 20	0.000	100.00

The highlighted row here says that when tree or tall shrub canopy closure is 80 to 100 %, and the vegetation zone is Western Hemlock, and the low shrub canopy closure is 0 to 20 %, then the site has a 70 % probability of being "suitable" (for some species) and 30% being "unsuitable." Each row in a CPT therefore is a statement of *probability* – that is, what are the possible outcome states (suitable, unsuitable) *given* some set of prior conditions.

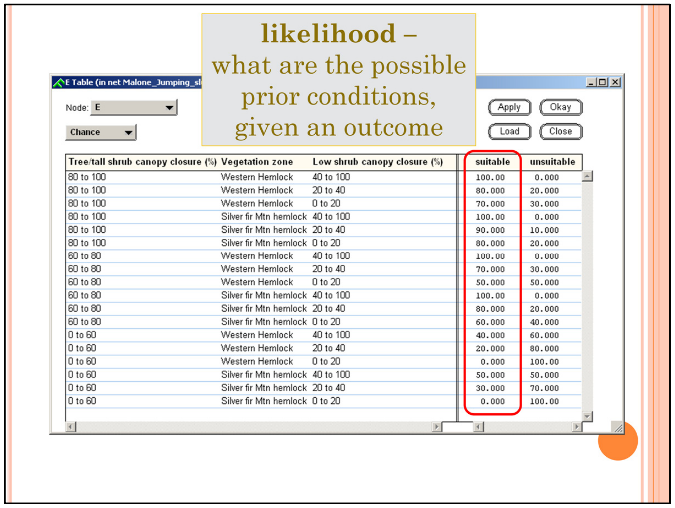

Tree/tall shrub canopy closure (%)	Vegetation zone	Low shrub canopy closure (%)	suitable	unsuitable
80 to 100	Western Hemlock	40 to 100	100.00	0.000
80 to 100	Western Hemlock	20 to 40	80.000	20.000
80 to 100	Western Hemlock	0 to 20	70.000	30.000
80 to 100	Silver fir Mtn hemlock	40 to 100	100.00	0.000
80 to 100	Silver fir Mtn hemlock	20 to 40	90.000	10.000
80 to 100	Silver fir Mtn hemlock	0 to 20	80.000	20.000
60 to 80	Western Hemlock	40 to 100	100.00	0.000
60 to 80	Western Hemlock	20 to 40	70.000	30.000
60 to 80	Western Hemlock	0 to 20	50.000	50.000
60 to 80	Silver fir Mtn hemlock	40 to 100	100.00	0.000
60 to 80	Silver fir Mtn hemlock	20 to 40	80.000	20.000
60 to 80	Silver fir Mtn hemlock	0 to 20	60.000	40.000
0 to 60	Western Hemlock	40 to 100	40.000	60.000
0 to 60	Western Hemlock	20 to 40	20.000	80.000
0 to 60	Western Hemlock	0 to 20	0.000	100.00
0 to 60	Silver fir Mtn hemlock	40 to 100	50.000	50.000
0 to 60	Silver fir Mtn hemlock	20 to 40	30.000	70.000
0 to 60	Silver fir Mtn hemlock	0 to 20	0.000	100.00

Contrast that with a column in the CPT, which represents *likelihood* – that is, what are the possible *prior conditions* that could lead to this specified *outcome state*?

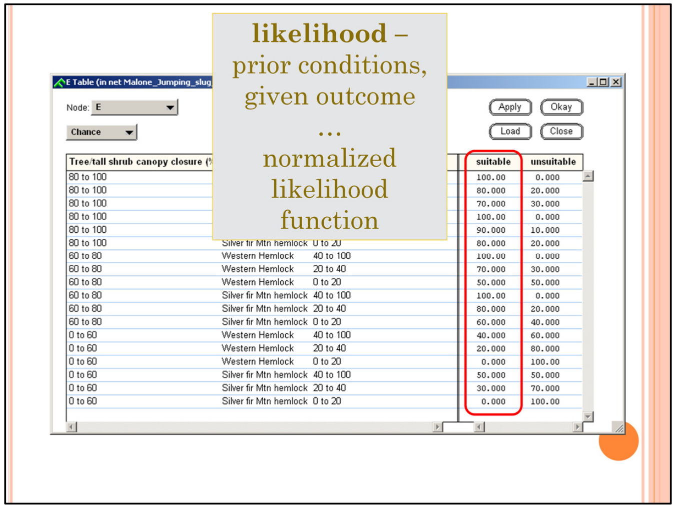

Tree/tall shrub canopy closure (			suitable	unsuitable
80 to 100			100.00	0.000
80 to 100			80.000	20.000
80 to 100			70.000	30.000
80 to 100			100.00	0.000
80 to 100			90.000	10.000
80 to 100	Silver fir Mtn hemlock	0 to 20	80.000	20.000
60 to 80	Western Hemlock	40 to 100	100.00	0.000
60 to 80	Western Hemlock	20 to 40	70.000	30.000
60 to 80	Western Hemlock	0 to 20	50.000	50.000
60 to 80	Silver fir Mtn hemlock	40 to 100	100.00	0.000
60 to 80	Silver fir Mtn hemlock	20 to 40	80.000	20.000
60 to 80	Silver fir Mtn hemlock	0 to 20	60.000	40.000
0 to 60	Western Hemlock	40 to 100	40.000	60.000
0 to 60	Western Hemlock	20 to 40	20.000	80.000
0 to 60	Western Hemlock	0 to 20	0.000	100.00
0 to 60	Silver fir Mtn hemlock	40 to 100	50.000	50.000
0 to 60	Silver fir Mtn hemlock	20 to 40	30.000	70.000
0 to 60	Silver fir Mtn hemlock	0 to 20	0.000	100.00

In fact, if a CPT is parameterized with "pegged" values of 0 and 100, then the column here can be interpreted as a normalized likelihood function. So again, it is important to differentiate between *probability* and *likelihood*. In the context of causality, probability tells you potential causal outcomes from some specified prior condition, whereas likelihood tells you potential causal conditions that could result in some specified outcome.

CONFIDENCE – HOW TO DENOTE CONFIDENCE IN EXPERT JUDGMENT USED TO DEVELOP PROBABILITY STRUCTURES?

CONTROL – HOW TO IDENTIFY MANAGEMENT CONTROL AND INFLUENCE OF DECISIONS?

CAUSE – HOW TO DETERMINE CAUSALITY?

So now we have delved into these three main aspects of Bayesian network modeling – confidence, control, and cause.

NEW AREAS OF RESEARCH

I'll finish by suggesting some new areas of research in Bayesian network modeling.

New Areas of Research

- Develop theory and tools for fractal Bayesian networks & a deeper theory of causality
  - self-similarity
  - feedback loops ++ and – –

First is the need for a deeper theory and tools for modeling fractal causal structures.

New Areas of Research

- Develop theory and tools for fractal Bayesian networks & a deeper theory of causality
  - self-similarity
  - feedback loops ++ and – –
- Merge Bayesian network modeling with structured equation modeling (SEM)

Second is the exciting area of more formally and fully merging Bayesian network modeling with the area of frequentist statistical modeling known as *structural equation modeling* or SEM. SEM is often used to depict causal structures, as well.

New Areas of Research

- Develop theory and tools for fractal Bayesian networks & a deeper theory of causality
  - self-similarity
  - feedback loops ++ and – –
- Merge Bayesian network modeling with structured equation modeling (SEM)

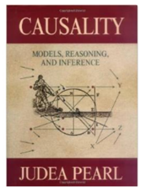

And just to mention that some groundwork on this area has been provided by Pearl and others.

New Areas of Research

- Develop theory and tools for fractal Bayesian networks & a deeper theory of causality
  - self-similarity
  - feedback loops ++ and – –
- Merge Bayesian network modeling with structured equation modeling (SEM)
- Develop a means of explicitly representing confidence in expert-judgment based nets

And the third are of new research could be developing an explicit way to represent confidence in expert-based Bayesian network models. Remember that CPTs depict probabilities (or likelihoods) of how conditions affect outcomes, but CPTs do *not* depict the degree of confidence in those values.

Confidence in BNs

The Journal of Wildlife Management 76(6):1298–1309; 2012; DOI: 10.1002/jwmg.366

Habitat Relations

Using Bayesian Networks to Incorporate Uncertainty in Habitat Suitability Index Models

GEORGE F. WILHERE,[1] *Washington Department of Fish and Wildlife, Habitat Program, 600 Capitol Way N., Olympia, WA 98501-1091, USA*

Available online at www.sciencedirect.com

ScienceDirect

ELSEVIER

Artificial Intelligence 172 (2008) 483–513

Artificial Intelligence

www.elsevier.com/locate/artint

Quantifying the uncertainty of a belief net response: Bayesian error-bars for belief net inference

Tim Van Allen[a], Ajit Singh[b], Russell Greiner[c,*], Peter Hooper[d]

There too has been some groundwork laid for this third area of suggeste research.

New Areas of Research

- Develop theory and tools for fractal Bayesian networks & a deeper theory of causality
  - self-similarity
  - feedback loops ++ and – –
- Merge Bayesian network modeling with structured equation modeling (SEM)
- Develop a means of explicitly representing confidence in expert-judgment based nets

So I leave you with these suggestions for future research areas.

The end! Now get to work.