Confidence Intervals for Bayesian nets

Margaret Donald
& Kerrie Mengersen

ABNMS2012
November 29, 2012

UNSW AUSTRALIA

1. Why bother?/limitations
2. Queries
3. Case Studies
4. Methods
5. Results
6. Conclusions
7. References

Context/Limitations

1. discrete BNs
2. BNs with binary nodes
3. generalise to finite populations

The distribution of a query is ?

1. asymptotically normal (Van Allen et al,2001, 2008)
2. possibly a mixture of beta distributions (Kleiter, 1996)
3. in some contexts, a mixture of beta distributions, in others, a mixture of gamma distributions (Hooper, 2008)

Queries

In a BN with four nodes, A, B, C, D one might wish to know $p(A = 1|B = 0, C = 0, D = 0)$ or more generally $p(h = H|e = E)$.
These probability estimates from BNs need confidence/credible intervals to

- help assess the validity of the BN
- draw useful conclusions

Case Studies

- Updated Diamond BN (Van Allen et al, 2008)
- Diarrhoea BN (Donald et al, 2009)

Case Studies: Diamond BN updated with experience

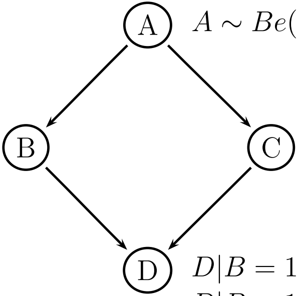

$$A \sim Be(4,2)$$

$$B|A=1 \sim Be(3,2)$$
$$B|A=0 \sim Be(1,2)$$

$$C|A=1 \sim Be(2,4)$$
$$C|A=0 \sim Be(1,2)$$

$$D|B=1,C=1 \sim Be(1,1)$$
$$D|B=1,C=0 \sim Be(2,2)$$
$$D|B=0,C=1 \sim Be(2,1)$$
$$D|B=0,C=0 \sim Be(1,2)$$

Diarrhoea BN of Donald et al. (2009)

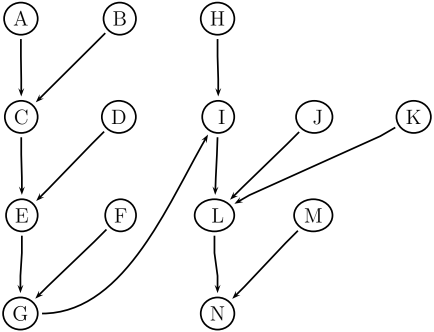

Beta distributions for Probability of 0.5

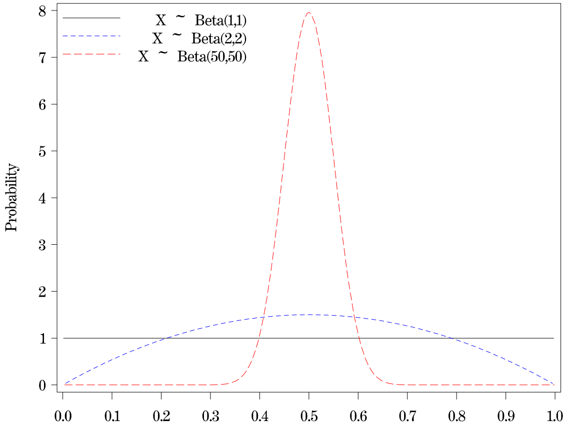

Figure: Some Beta distributions for p=0.5

Some settings for beta priors in the Diarrhoea BN

Node & Value	*p*	α	β
A=1	0.01	6.751	668.37
B=1	0.01	0.599	59.25
D=1	0.99	59.2524	0.59851
F=1	0.80	48.372	12.093
H=1	0.90	30.217	3.35744
J=1	0.90	30.217	3.35744
K=1	0.50	47.52	47.52

Finding α and β in the Beta distribution

For a beta distribution
$E(X) \sim \frac{\alpha}{\alpha+\beta}$
$Var(X) \sim \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$
Having elicited p and its 95% confidence interval,

1. Approximate the variance from $(Q_{.975} - Q_{.025})$ as $1.96 \times s$ where $s^2 = Var(X)$.
2. Equate moments.

This gives, e.g.,
$.98(.95, .99) \Rightarrow Beta(183.494, 3.75)$

Queries: Diamond BN

Query	*(h=H*	|	*e=E)*
1	A=1		
2	A=1	|	B=1
3	A=1	|	B=1, C=1
4	B=1, C=1	|	A=1
5	A=1	|	D=1
6	D=1	|	A=1
7	B=0	|	C=0

Queries: Diarrhoea BN

Query	*(H=h*	|	*E=e)*
1	N=1	|	L=1,M=1
2	N=1	|	L=1,M=2
3	N=1	|	L=1,M=3
4	N=1	|	G=0,M=1
5	N=1	|	G=0,M=2
6	N=1	|	G=0,M=3
7	N=1	|	G=0,H=0,M=1
8	N=1	|	G=0,H=0,M=2
9	N=1	|	G=0,H=0,M=3
10	E=0	|	N=1
11	E=0	|	N=1
12	E=0,H=0	|	N=1

Methods

A BN with p nodes representing a finite population of size n may be thought of as being an $n \times p$ matrix.

1. Gold standard: Generate m replicates of the $n \times p$ BN. Within each replicate find $p(h = H|e = E)$. Summarise over the m replicates to find a 95% CI of the MC distribution. (We used WinBUGS to do this.)
2. Generate a single replicate of the $n \times p$ BN. Find $p(h = H|e = E)$ for this population. Calculate exact binomial CIs.
3. Use the variance and probability found from method 1, to calculate normal approximations to the CIs.
4. Find the implied BN probabilities for $p(h = H|e = E)$ and $p(e = E)$. Calculate the expected population satisfying $e = E$ as $n \times p(e = E)$ rounded, and use the exact binomial CIs of
   - Clopper-Pearson (1934) (Method 4), or,
   - the Bayes-Laplace of Tuyl et al(2008) (Method 5).

Diamond BN: Query 4, population 25

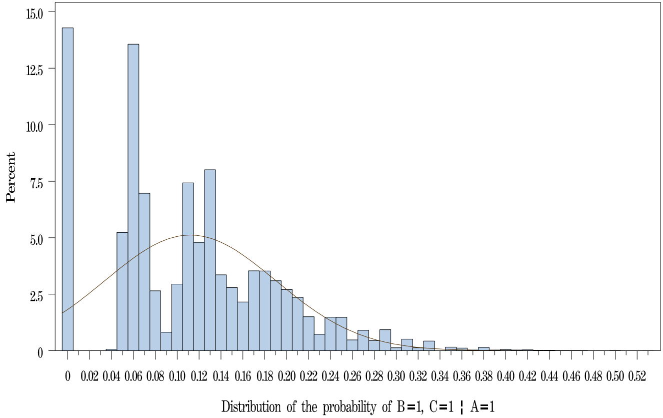

Figure: Query 4: Distribution of p(B=1,C=1|A=1).

Diamond BN: Query 4, population 2000

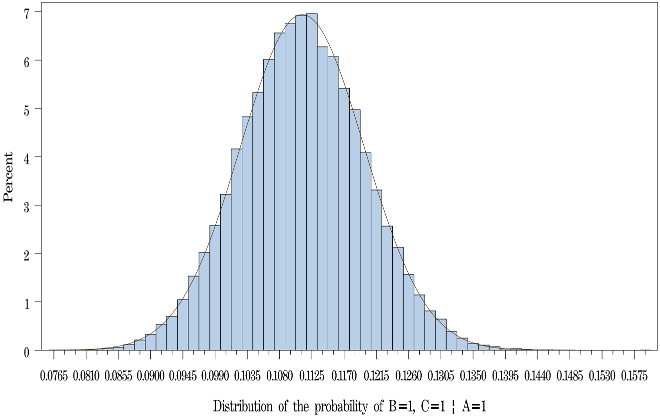

Figure: Query 4: Distribution of p(B=1,C=1|A=1).

Diarrhoea BN: $p(N = 1|G = 0)$, population 50000

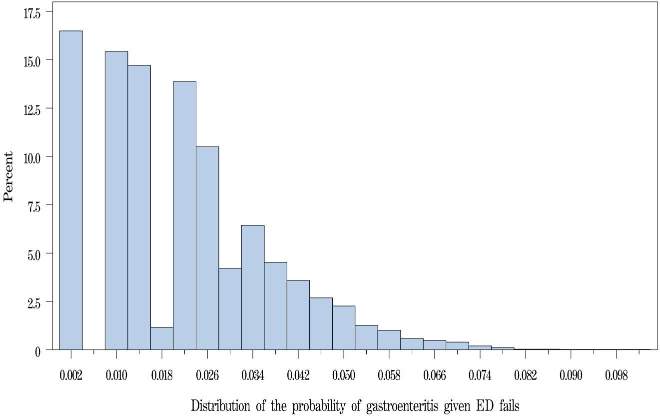

Figure: Distribution of the probability of being infected with gastroenteritis when the endpoint distribution fails.

Diamond BN: population 100

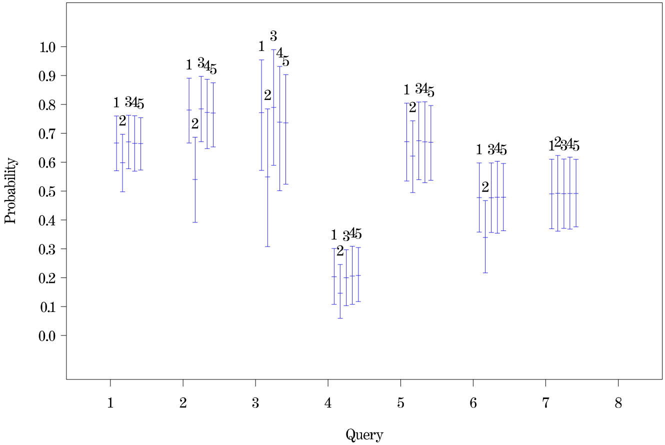

Figure: Diamond BN queries & confidence limits.

Diamond BN: population 2000

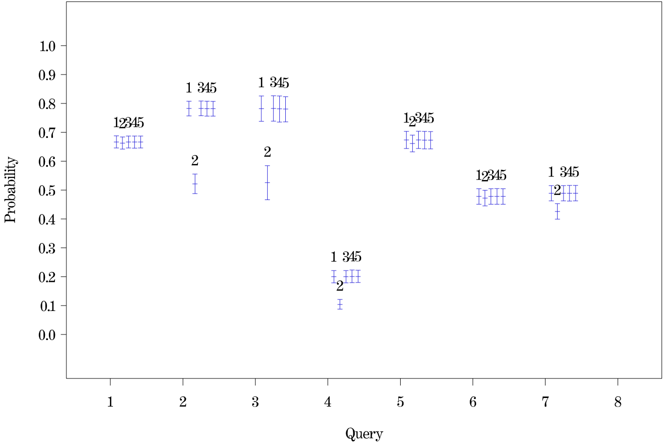

Figure: Diamond BN queries & confidence limits.

Diarrhoea BN: population 50000

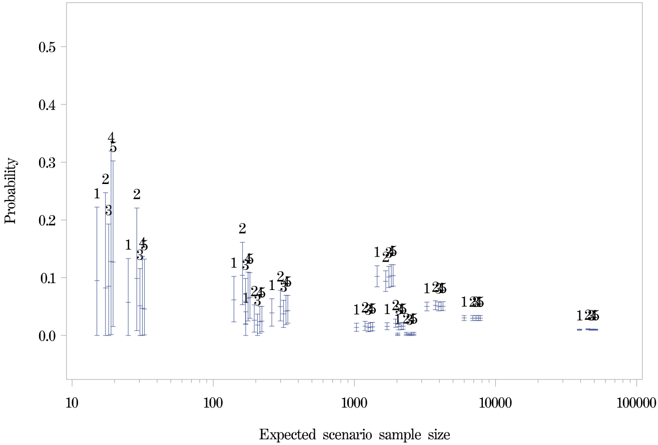

Figure: Diarrhoea BN queries & confidence limits.

Summary

- Distributions of queries
  - Are asymptotically normal, but
  - it is not always clear when this behaviour occurs.
- Finding CIs
  - 'the gold standard' for CIs (method 1) can be difficult to apply when using WinBUGS. It is all too easy to mismatch the beta distributions to the correct condition.
  - Using exact probabilities and expected scenario sizes together with 'exact' binomial CIs gives results comparable to the gold standard. But these can be difficult to calculate. They should be used where they are directly available from the BN
  - When $p(h = H|e = E$ and/or $p(e = E)$ are difficult to find, a single simulation of the BN with the desired population size can give useful idea of the uncertainty associated with the query.

Postscript

Confidence intervals can be produced using BN software, by

1. Adding "experience" which matches the sum of α and β, for each condition in the CPT tables.
2. Producing a set of simulated cases of size $m \times n$, thereby producing an $(m \times n) \times p$ matrix.
3. Calculate the query value for each $n \times p$ matrix,
4. and summarise the m query results.

(You probably need to do this in a single simulation, in order to ensure the randomness of successive population draws.)

Selected References

Clopper, C. and Pearson, E.: 1934, The use of confidence or fiducial limits illustrated in the case of the binomial, *Biometrika* **26**(4), 404–413.

Donald, M., Cook, A. and Mengersen, K.: 2009, Bayesian network for risk of diarrhoea associated with the use of recycled water, *Risk Analysis* **29**(12), 1672–1685.

Hooper, P. M.: 2008, Exact distribution theory for belief net responses, *Bayesian Analysis* **3**(3), 615–624.

Kleiter, G.: 1996, Propagating imprecise probabilities in Bayesian networks, *Artificial Intelligence* **88**(1-2), 143–161.

Tuyl, F., Gerlach, R. and Mengersen, K.: 2008, A comparison of Bayes-Laplace, Jeffreys, and other priors, *The American Statistician* **62**(1), 40–44.

Van Allen, T., Greiner, R. and Hooper, P.: 2001, Bayesian error-bars for Belief Net inference, *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI-01)*, Citeseer, Seattle.

Van Allen, T., Singh, A., Greiner, R. and Hooper, P.: 2008, Quantifying the uncertainty of a Belief Net response: Bayesian error-bars for Belief Net inference, *Artificial Intelligence* **172**, 483–513.

Thank you for listening. Any questions?