A Bayesian Network for a simple example of Drug Economics Decision Making

Norman Fenton (updated 16 March 2014)

Suppose that a relatively cheap drug (drug A) has been used for many years to treat patients with disease X. The drug is considered quite successful since data reveals that 85% of patients using it have a 'good outcome' which means they survive for at least 2 years. The drug is also quite cheap, costing on average $100 for a prolonged course (see Fig 1 (a)). The overall "financial benefit" of the drug (which assumes a 'good outcome' is worth $5000 and is defined as this figure minus the cost) has a mean of $4156.

There is an alternative drug (drug B) that a number of specialists in disease X strongly recommend. However, the data reveals that only 65% of patients using drug B survive for at least 2 years (Fig. 1(b)). Moreover, the average cost of a prolonged course is $500. The overall "financial benefit" of the drug has a mean of just $2777.

(a) Results for Drug A The mean financial benefit is $4156	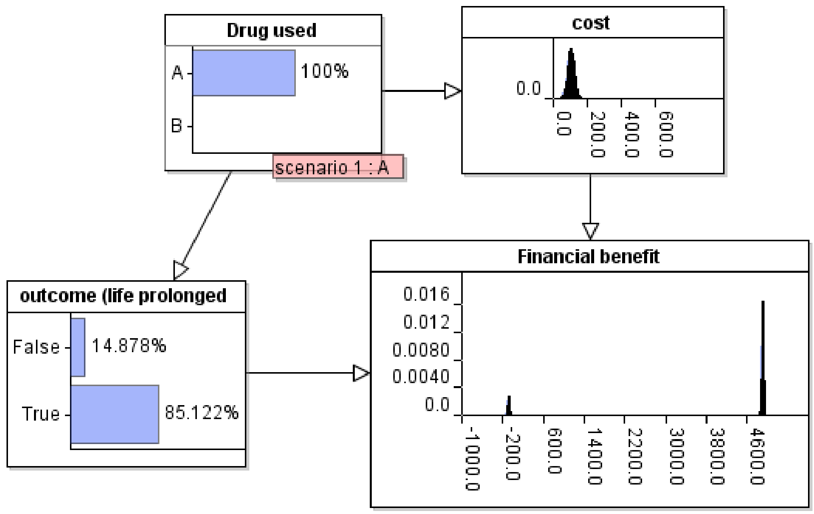
(b) Results for Drug B The mean financial benefit is $2777	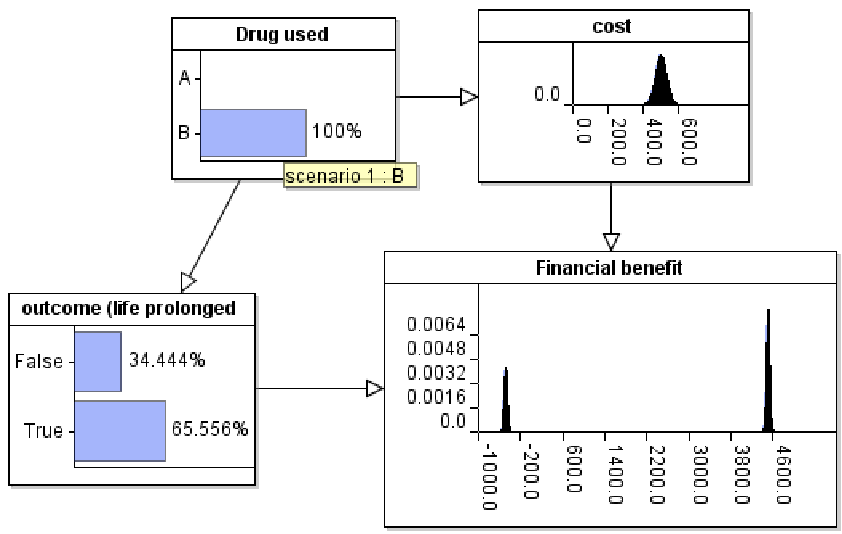

Figure 1 Basic results for drug effectiveness

On seeing the data the Health Authority recommends a ban against the use of drug B. Is this a rational decision?

The answer is no in this case because the data is hiding a very important variable: the ***actual patient condition*** with respect to disease X. Obviously patients who have more serious cases of the disease are more likely to die within two years irrespective of the treatment they receive, while patients who have only a minor condition will survive irrespective. And, crucially, most patients in the database only have the minor condition. The more complete model is shown in Fig 2, with the prior marginal probabilities.

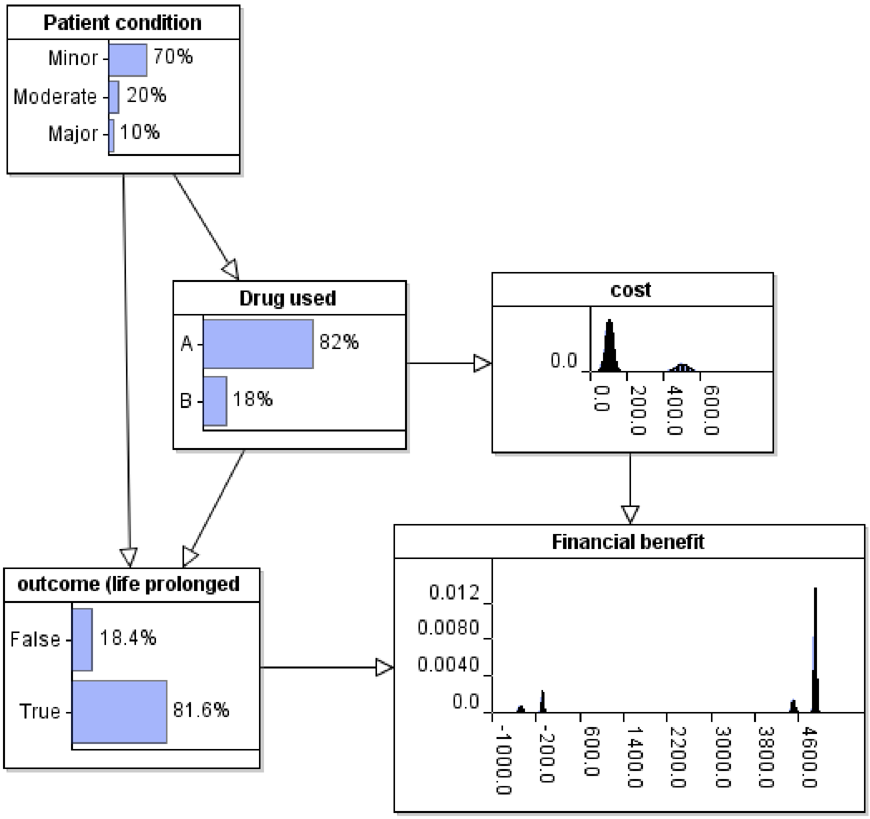

Figure 2 Model with missing variable 'Patient condition' shown

The model (which you can run in AgenaRisk[1] by selecting the model in Examples\Basic\Drug Economics) can be used to 'simulate' a randomized controlled trial by fixing the patient condition before observing the effects of the two drugs. Most crucially, when we do this for patients classified as 'major' we get the very different results shown in Fig 3.

[1] In the AgenaRisk model the node 'patient condition' is a hidden node which you can choose to display or not by selecting the appropriate menu button

(a) Results for Drug A
Only 10% positive outcome.
The mean financial benefit is $400

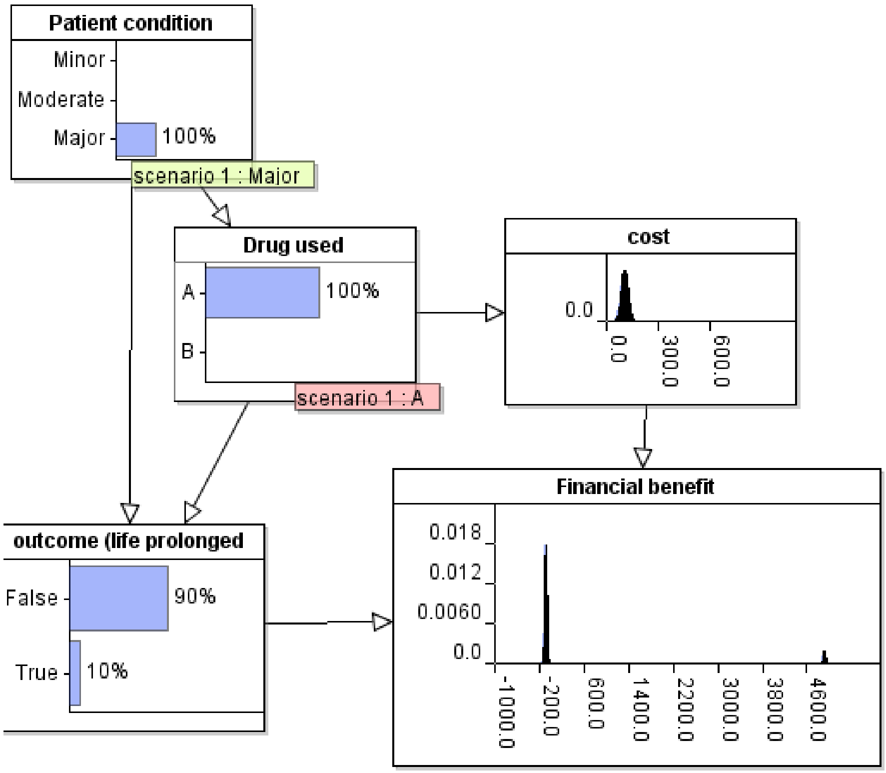

(b) Results for Drug B
30% positive outcome.
The mean financial benefit is $1000

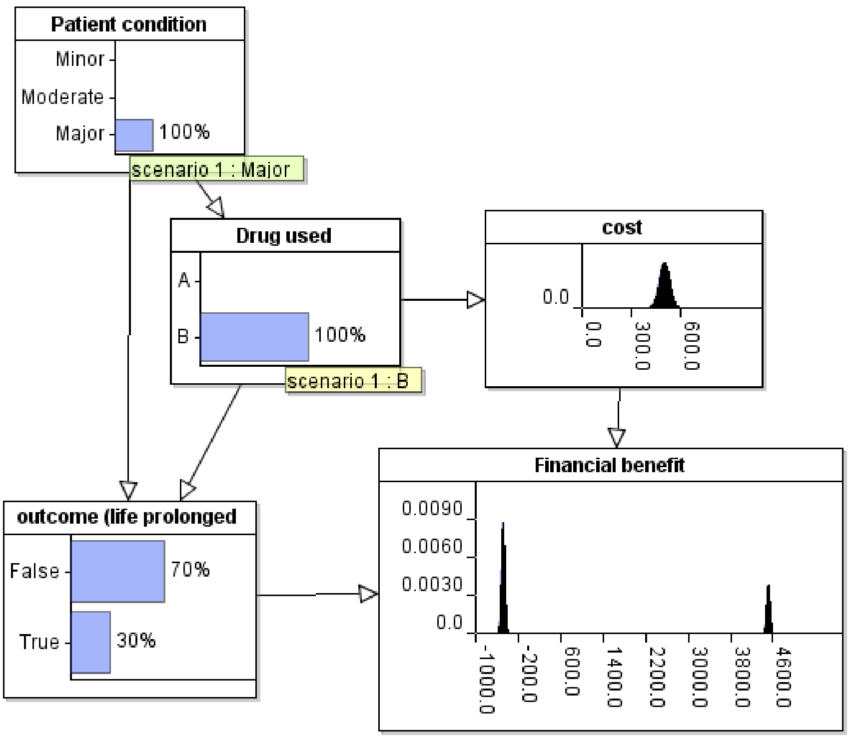

Figure 3 Results with 'Patient condition' major

In this case Drug A results in only 10% positive outcomes, while drug B results in 30%. The mean financial benefit is respectively $400 and $1000. The *same* data as used in the original recommendation therefore leads to a very different recommendation: use drug B when the patient condition is major.

The example might seem to suggest that decisions about drug effectiveness must always be subject to proper randomized controlled trials in which variables like 'condition of patient' can be controlled. However, in reality we will rarely have sufficient patients in a trial to account for all possible variables that need to be controlled. It is far more realistic and useful to use the observational data that is available (as in the above BN), combined with expert judgement about the impact of variables for which there is little or no data to enhance a casual BN.

This example also demonstrates how utility nodes (such as the node 'financial benefit' here) can be incorporated into a tool like AgenaRisk to support decision-making. The NPT for the node is defined as the partitioned expression shown in table 1

Table 1 NPT for node 'financial benefit'

outcome (life prolonged more than 2 years)	False	True
Expressions	Arithmetic(-cost)	Arithmetic(5000-cost)