The Bayesian Art of Managing Risks

Armin Haas, Jette Krause

3 April 2007

Decision makers are paid for making decisions. Very often, the tricky part of their job is that they have to decide under uncertainty. They are confronted with several possible futures and have to come up with a judgement of how to assess the situation.

Basically, there are a couple of standard procedures of how to deal with this problem. One is the wait-and-see approach: Just postpone the decision until more evidence decreases the uncertainty. Unfortunately, very often this strategy cannot be applied since the decision is due.

Another one is to consult experts, be it a lawyer, a business consultant, a political adviser, or a scientist. The wellknown problem of this approach is that very often the advice from different experts is contradictory. We know the situation when consulting, say, some physicians about a severe medical problem: three doctors, three opinions.

Still another strategy is to pick an alternative at random.

As social scientists, we are basically interested in two questions:

1. How do decision makers behave in such situations?

2. How should they behave?

The first question addresses the positive aspect taking the empirical perspective. The second question is normative and asks for the optimal way of dealing with such a situation.

Our answer is that decision makers can very often be described as performing a specific way of handling uncertainty in their decision processes. This way we call Bayesian Risk Management. In the following, we will sketch the main ideas of BRM and argue why decision makers are wise applying it.

East or West?

In October 2005, hurricane Wilma ravaged the peninsula of Yucatan, Mexico. In some places, all the tourist facilities got destroyed. The owner of a ravaged hotel now faces a severe choice: Should he rebuild his hotel at the very place or should he move? In fact, some hoteliers are contemplating about moving to the pacific coast of Mexico.

In this respect, the main problem for the hotel owners are hurricanes. How will they develop in intensity and frequency? Any hotel owner has to tackle a double stage problem: On the first level, he must come up with an idea of how and how often hurricanes will damage his facilities and cause business interruptions. On the second level, he must form expectations about how he will be able to get insurance coverage for these damages. For performing the latter, he does not only have to project his own expectations, but the expectations of the insurance and reinsurance industry. The worst thing he could encounter is a situation in which the insurance industry, at some instance, would stop giving cover for a facility that to a considerable extent had not yet been depreciated.

A natural first step in tackling this tough challenge seems to be consulting experts for tropical storms. In late April 2006, these experts gathered in Monterey, California. As a guest, our hotel owner would have been rather puzzled: The experts did neither agree on basic theories nor on what to expect. Some even challenged the meteorological measurement and thus the time series of hurricane activities. In particular, the experts did disagree on whether the pattern of hurricane intensity in the Caribbean oscillates or follows a trend. If an oscillation prevails, the question would be about its frequency and amplitude. For the practical purpose of our hotel owner, this is about figuring out at which point of the oscillation he is right now. If a trend prevails, the question is about the dynamics of the process. Here, the hotel owner needs to have an estimate of the growth rate that governs the trend. In fact, the decision problem may be more complicated as the alternatives of oscillation and trend may not be exclusive but be overlaid.

The Basic Principle of BRM

Let us consider a decision maker who has to deal with a figure that is essential for the decision to take. This figure could be, for example, the annual intensity or frequency of hurricanes. When asking a scientist, a

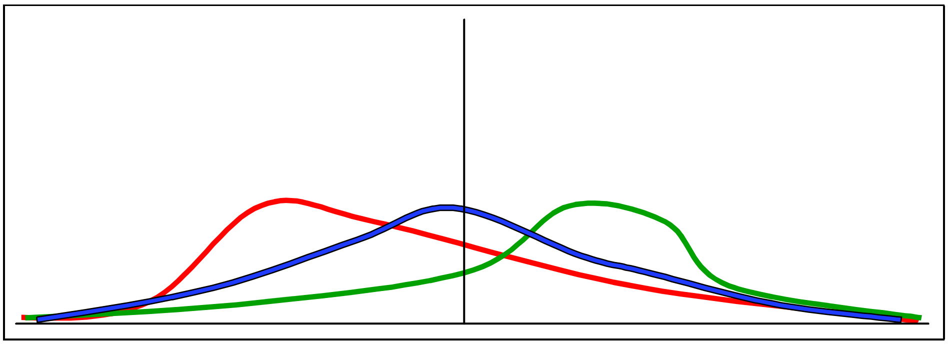

Figure 1: Three possible density functions for an index value.

reasonable expertise she could get is a density function over the range of possible values for the intensity or frequency. For the following, let us deal with an abstract index value. This index can be thought of as being an index of intensity, frequency, or the like. Suppose that, after asking three experts, our decision maker is confronted with the three density functions sketched in figure 1. In Baysian notion, these density functions depict so-called *first order* probabilities. A reasonable way to proceed is to weigh the alternative density functions with a measure that reflects her preference of the specific functions. One way of doing this is attaching to each density function a so-called *second order* probability. Figure 2 shows an example of second order probabilities attached to the density functions.

The procedure just sketched seems to be rather ad hoc, and it is. This caveat is important but not crucial. Crucial is that the decision maker is willing and able to learn. In our example, this means that the decision

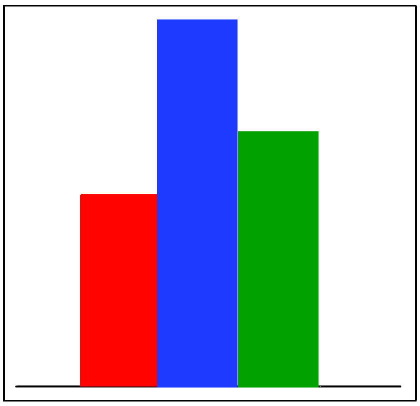

Figure 2: Second order probabilities of the three possible density functions of figure 1.

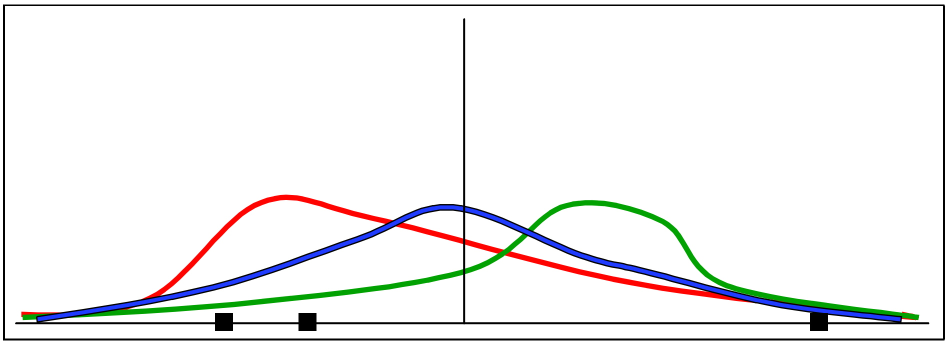

Figure 3: The three density functions and three realisations of an index value.

maker updates her second order probabilities. Let us assume that after three years have passed, the decision maker knows the annual index values of these three years as indicated in figure 3. A reasonable conclusion is to raise the second order probabilities of those densities for which the realised index values are more likely. The two index values on the left are much more likely for the red density function than for the green one, whereas all density functions assign similar probabilities to the index value on the right. Therefore, it seems reasonable that the red density function will gain probability and the green one will lose. Figure 4 displays the updated second order probabilities.

This simple example also shows that is is less obvious how the second order probability of the blue density function should change, if at all. What is needed is a clear rule how the weights given to the density functions

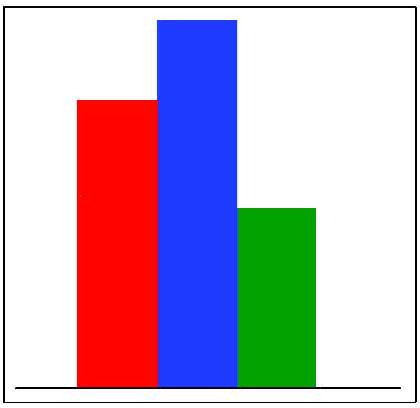

Figure 4: Updated second order probabilities of the three possible density functions.

on consideration, i.e. the second order probabilities, should reasonably be updated. DeFinetti's Theorem states such a rule.

DeFinetti's Theorem

In order to give an idea of DeFinetti's theorem, we consider an index that can have several discrete values. Formally, each value can be described as a discrete state s out of a discrete state space S (the set of all possible states, i.e. values). As an example for such states, think of the number of hurricanes of category four in a hurricane season. A specific expectation of the future can then be described by the first order probability a decision maker attaches to each possible state s. In the continous example of the previous paragraphs, the three densities were examples of such first order probabilities (in a continous state space with an indefinite number of states – we now focus on discrete states in order to ease notation). When confronted with several expectations, i.e. several (discrete) first order probability distributions, the decision maker attaches a second order probability to each distribution.

Now suppose that out of all possible states, a specific state s has realised. In our hurricane example, such a state could be that in a hurricane season 3 hurricanes of category 4 have occured.

A possible updating rule for the second order probability $p_{2,t+1}$ of a specific first order probability distribution j reads as follows:

$$p_{2,t+1}(j) = p_{2,t}(j) \cdot \frac{p_j(s)}{\sum_{i \in I} p_{2,t}(i) \cdot p_i(s)}$$

whereas $p_{2,t}(j)$ denotes the second order probability $p_{2,t}(j)$ of the first order probability distribution j in period t, and $p_j(s)$ denotes the first order probability for the state s under the assumption of the first order distribution j. The sum of the denominator is taken over all possible distributions $i \in I$.

DeFinetti has proven that, under rather general assumptions, this Bayesian updating rule converges towards the "true" distribution of the random process under consideration.

References

Agosta J.M., Gardos, T. (2004): Bayes Network 'Smart' Diagnostics, Intel Technology Journal, Santa Clara.

Haas, A., Jaeger, C. (2005): Agents, Bayes, and Climatic Risks – a modular modelling approach, Advances in Geosciences, 4, p. 3–7.

Hasselmann, K. (1998): Conventional and Bayesian Approach to Climate-Change Detection and Attribution, Quarterly Journal of the Royal Meteorological Society, 124(552): p. 2541–2565.

Jaeger, C., Krause, J., Haas, A., Klein, R., Hasselmann, K. (2007): The Fraction of Attributable Risk for the Alpine Heat Wave, to be submitted to Risk Analysis.

Jaeger, C. et al. (2001): Risk, Uncertainty, and Rational Action, London.

Kerr, R.A. (2006): A Tempestuous Birth for Hurricane Climatology, Science, Vol. 312, p. 676-678.

Kreps, D. (1988): Notes on the Theory of Choice, Boulder.

Rouanet, H. et al. (1998): New Ways. In: Statistical Methodology, Bern.