Annals of Actuarial Science, Vol. 6, part 2, pp. 344–380. © Institute and Faculty of Actuaries 2012
doi:10.1017/S1748499512000048
First published online 22 May 2012

Computational intelligence with applications to general insurance: a review

II. Dealing with uncertain knowledge

Pietro Parodi*
Willis Global Solutions (Consulting Group), London

Abstract

This paper argues that most of the problems that actuaries have to deal with in the context of non-life insurance can be usefully cast in the framework of computational intelligence (a.k.a. *artificial* intelligence), the discipline that studies the design of agents which exhibit intelligent behaviour. Finding an adequate framework for actuarial problems has more than a simply theoretical interest: it also allows a technological transfer from the computational intelligence discipline to general insurance, wherever techniques have been developed for problems which are common to both contexts. This has already happened in the past (neural networks, clustering, data mining have all found applications to general insurance) but not in a systematic way. One of the objectives of this paper will therefore be to introduce some useful techniques such as sparsity-based regularisation and dynamic decision networks that are not yet known to the wider actuarial community.

Whilst in the first part of this paper we dealt mainly with data-driven loss modelling under the assumption that all the data were accurate and fully relevant to the exercise, in this second part of the paper we explore how to deal with uncertain knowledge, whether this uncertainty comes from the fact that the data are not fully reliable (e.g. they are estimates) or from the fact that the knowledge is "soft" (e.g. expert beliefs) or not fully relevant (e.g. market information on a given risk). Most importantly, we will deal with the problem of making pricing, reserving and capital decisions under uncertainty. It will be concluded that a Bayesian framework is the most adequate for dealing with uncertainty, and we will present a number of computational intelligence techniques to do this in practice.

Keywords

Intelligent agents; Risk agents; Regularisation; Markov decision processes; Game theory; Bayesian networks; Dynamic decision networks; Fuzzy set theory; Expert systems; Non-monotonic reasoning; Kalman filtering

1. Introduction

In the first part of this paper on the application of computational intelligence techniques in general insurance we have focused on the role of statistical learning in risk costing. The underlying

*Correspondence to: Pietro Parodi, Willis Global Solutions (Consulting Group), Willis Ltd, 51 Lime Street, London EC3 M 7DQ. E-mail: pietro.parodi@willis.com

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

assumption was that data were reliable and were the only source of information (data-driven modelling).

In actuarial practice, however, data are often themselves uncertain: e.g. losses of an insurance company may be estimates, or they need to be revalued (with uncertain claims inflation) to current values, or they may refer to a period in which the risk profile of the insurer was different. Furthermore, data may be insufficient and may need to be supplemented with market information, which is not fully relevant to the insurer, or with "soft" knowledge (e.g. expert knowledge by underwriters).

In this second part of this paper we are therefore going to address the problem of dealing with uncertain knowledge and with the problem of integrating soft/expert knowledge. Section 2 ("Dealing with uncertainty") compares different methodologies (especially **fuzzy set theory**, an extension of set theory in which elements can have partial membership to a set, and **Bayesian networks**, a method for describing a joint probability distribution succinctly in terms of the dependencies between the variables) for dealing with uncertain knowledge. Section 3 ("Temporal reasoning about risk") analyses methodologies to deal with one specific issue – that of the environment changing with time, which makes knowledge accumulated about risk rapidly obsolete. Kalman filtering (an extension of linear regression which allows the regression parameters to vary with time), which has been applied in the past to reserving, will be shown to be a special case of a more general technique, that of **dynamic Bayesian networks**, which are an extension of Bayesian networks. Section 4 ("Making decisions in an uncertain environment") finally moves beyond risk costing (which is not concerned with wider commercial considerations) to address the crucial issue about pricing, reserving and capital modelling: that of making risk management decisions (e.g. how much to charge for a policy, what amount of reserves/capital to hold, etc.) in an uncertain environment. It will be shown that decision making can be analysed in terms of **partially observable Markov decision processes** (POMDP's), which are a formalism for describing sequential decision problems and can be implemented as **dynamic decision networks** (another extension of Bayesian networks).

Section 5 ("Modelling collective behaviour") provides a foray into the thorniest issue of all, that in which the decision-making process needs to take into account the individual strategies of other risk agents (e.g. other insurers) and to react to them. The natural framework for this is provided by **multi-agent systems** theory and by **game theory**. The interaction between risk agents can be modelled as a game of partial information, in which each player only has a limited knowledge of the strategy of its competitors, and implemented by stochastic simulation as in Axelrod (1984). It should not be concealed, however, that at the present stage this is only a speculative solution, and that the problem remains difficult: it is not completely understood theoretically and any stochastic implementation is likely to lead to a strongly sub-optimal solution.

Section 6 ("Conclusions") draws together the overall conclusions of this investigation. The first one is that computational intelligence provides a natural framework for understanding risk. The second is that several computational intelligence methodologies which are little known to the actuarial community (e.g. regularised regression, spectral clustering, partially observable Markov decision processes) can be usefully transferred to the actuarial context. The third is that most if not all of these techniques have an obvious Bayesian interpretation, suggesting that Bayesian theory is a powerful unifying framework for computational intelligence and therefore, indirectly, for non-life actuarial practice. The limitations of the computational intelligence approach are also looked at, calling into question its ability to model the lack of stability of the risk environment, and its

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

effectiveness in taking the human element into account. Finally, the limitations of the investigation itself are mentioned.

2. Dealing with Uncertain and Soft Knowledge

We have so far assumed that data were just data – hard facts that we can count on, and on which we can build a model; we have also assumed that the data was the only information that mattered to the problem. In reality, data are themselves uncertain (e.g., only the estimated value of a claim is often available, rather than the final settled amount), and there is plenty of "soft" knowledge that we cannot ignore.

Computational intelligence has developed many methods to incorporate uncertain and soft knowledge. The most important ones are fuzzy set theory and Bayesian analysis. We are going to take a close look at both these methods in the section and compare their usefulness to actuaries.

Other methods, such as the certainty factors model (initially developed for the Mycin medical diagnosis system and subsequently used in many other "expert systems" in the 1970s and the 1980s) and the Dempster-Shafer belief functions exist, but they are currently viewed as less important. For a brief discussion of these methods, the interested reader is referred to Section 14.7 of the book by Russell & Norvig (2003).

2.1 Fuzzy set theory

The motivation behind fuzzy set theory (Zadeh, 1965) is the attempt to capture and handle analytically the notion of an object (e.g. a number, or a logical value) whose value is not sharply defined: e.g. "a bad winter", "a large loss", "a risky account".

At the core of fuzzy set theory is the concept of *fuzzy membership*, a generalisation of the concept of membership to a set: in the "vanilla" version of set theory, an element x of a space Ω either belongs to a set A or it does not: the value of the membership function $\mu_A(x)$ is either 1 or 0. In fuzzy set theory, $\mu_A(x)$ can be any number between 0 and 1, e.g., we can say that losses of £10, £100k, £10m have membership 0, 0.5, 1 (respectively) to the set of large losses.

Fuzzy set operations (union, intersection…) can also be defined as a natural extension of the operations for standard set theory: e.g., the membership function of the union of two sets is the maximum of the membership functions of the two sets: $\mu_{A\cup B}(x) = \max\{\mu_A(x), \mu_B(x)\}$.

Fuzzy numbers (as opposed to *crisp* numbers) can also be defined based on membership functions. In informal terms, a fuzzy number is a fuzzy subset of the real numbers $\mathbb{R}$ whose membership function is centred on a given real number. A *fuzzy arithmetic* can then be created based on these fuzzy numbers. The essential tool for this is Zadeh's extension principle, which states that if $*$ is a binary operation (e.g. sum, product…) and A, B are two fuzzy numbers, then the membership function for $A(*)B$ is

$$\mu_{A(*)B}(z) = \sup_{x,y}\{\min[\mu_A(x), \mu_B(y)] | x * y = z\} \tag{1}$$

In the simple case of the sum $A \oplus B$ of two triangular fuzzy numbers A and B, things work as in Figure 1. Much of the arithmetic of fuzzy numbers bears an obvious resemblance to interval arithmetic.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

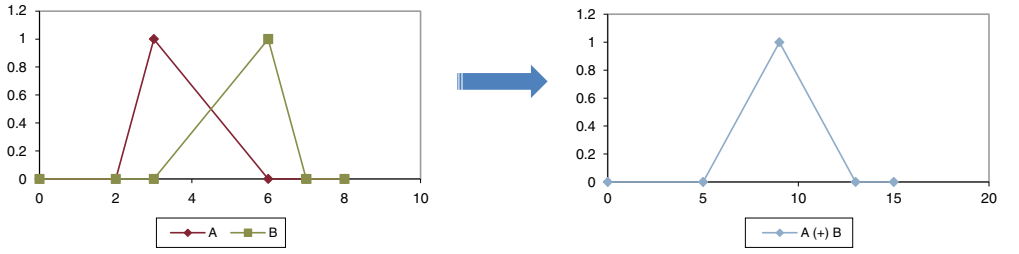


Figure 1. Sum of two fuzzy numbers A, B.

Moving a step forward, we can use Zadeh's extension principle to define *crisp functions of fuzzy numbers:*

$$\mu_{f(A)}(z) = \sup_{x \in \mathbb{R}} \{\mu_A(x) | z = f(x)\} \tag{2}$$

Although the extension principle is simple when applied to monotonic functions and triangular fuzzy numbers, it becomes quickly unwieldy when applied to non-monotonic complex functions and to operations involving fuzzy numbers with complex membership functions. It is perhaps for this reason that (as far as the author is aware) a complete fuzzy calculus has never been developed.

A simple example of fuzzy numbers in non-life insurance is provided by the individual loss estimates for claims that have been reported but have not yet been settled. This will be a fuzzy number with lower limit equal to the amount paid that far and the higher limit equal to the policy limit, or a narrower range if the claim manager is confident about the possible outcome. The highest value of the membership function corresponds to the best estimate.

Fuzzy set theory has been explored at some length in the actuarial literature. The best known example is perhaps the paper by Cummins & Derrig (1997), "Fuzzy financial pricing of property-liability insurance", which attempts to "fuzzify" the Myers-Cohn method for pricing general insurance contracts, according to which the price of a policy is set so that "the net present value of premiums is equal to the present value of the losses, expenses and federal income taxes incurred by the company as a result of issuing a particular policy". The reader interested in a comprehensive overview of the applications of fuzzy set theory to insurance is referred to the paper by Shapiro (2004).

2.2 The Bayesian approach

In the Bayesian approach to modelling uncertainty, both uncertain data and soft/expert knowledge are also modelled as probability distributions. This approach is based on the Bayes theorem: $\Pr(\theta \mid \underline{Z}) = \Pr(\underline{Z} \mid \theta)\ \Pr(\theta)/\Pr(\underline{Z})$, which allows to calculate the posterior probability $\Pr(\theta \mid \underline{Z})$ of θ (e.g., a parameter) conditional to $\underline{Z}$ (e.g., the data) as the probability $\Pr(\underline{Z} \mid \theta)$ of $\underline{Z}$ conditional to θ times the prior probability $\Pr(\theta)$ of θ.

Formally, this is a straightforward consequence of the axioms of probability theory when using conditional probability, but the use of $\Pr(\theta)$ to represent our initial beliefs on θ – rather than an objective state of things – has long been controversial (see for example Popper, 1959). Despite this controversy, the Bayesian approach has gained more and more acceptance in recent decades,

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

undoubtedly aided by the availability of larger computing resources that have made it possible to calculate posterior probabilities in complex situations.

Actuaries are familiar with the Bayesian formalism. For example, credibility theory, a standard tool of general insurance, can be framed in a Bayesian context: the information coming from the data is combined with that coming from an external source (which represents our prior knowledge), and the relative weight is determined by the balance of the variance of our data and the variance of the prior distribution.

In all but the simplest cases, there are many stages between the available data and the final results of a risk costing exercise. Prior knowledge and uncertainty are found at all stages: it is therefore important to be able to determine how uncertainty and beliefs are propagated across all stages of the methodology. Bayesian networks provide a way to do this. The most complete reference to Bayesian networks is probably the book by Pearl (1988). A good, shorter (and simpler) introduction can however be found in the book by Russell & Norvig (2003), arguably the best introduction to artificial intelligence available.

Bayesian networks provide a way of representing a joint probability distribution, highlighting the dependency relationships between the variables. Formally they are **directed, acyclic graphs**, i.e. collections of nodes and arcs where each arc has a direction and there are no cycles:

- Each node represents a random variable and has an associated quantitative probability information.
- Two nodes A and B are connected by an arc from A (the parent node) to B (the descendant) if B is dependent on A.

The cornerstone for the construction of a Bayesian network is the so-called **chain rule** which gives the joint probability distribution as a function of the conditional probabilities:

$$\Pr(X_1, X_2 \dots X_n) = \Pr(X_n|X_1, \dots X_{n-1})\Pr(X_{n-1}|X_1, \dots X_{n-2}) \dots \Pr(X_2|X_1)\Pr(X_1) \tag{3}$$

To specify a Bayesian network, we need to identify which dependencies in Equation (3) (which is fully general) can be dispensed with, in order to make the network as compact as possible. Therefore we want to replace Equation (3) with the sparser

$$\Pr(X_1, X_2 \dots X_n) = \prod_{i=1}^{n} \Pr(X_i|\text{parents}(X_i)) \tag{4}$$

In Equation (4), parents(X_i) is the set of variables on which X_i actually depends directly. This suggests another view of Bayesian networks (alternative to that of a joint probability distribution) as **a graph encoding a collection of conditional independence statements**. This encoding is called a "topological semantics" and can be expressed by one of the following equivalent specifications:

- a node is conditionally independent of its non-descendants given its parents;
- a node is conditionally independent of all the other nodes in the network, given its Markov blanket (= parents, children, children's parents)

A conditional probability table (CPT) must be associated with each random variable: this provides the quantitative information of the dependency of a node on its parents.

An example of Bayesian network is provided in Figure 4.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

2.3 A practical general insurance example: Determining the parameters of a severity distribution based on estimated losses and using expert knowledge

The differences between different methods to address uncertainty are better understood through a simple example. Suppose we are pricing some type of liability business and that there is a consensus that the proper model for doing so is a single parameter Pareto distribution with cumulative distribution $F(x) = 1-(x_0/x)^{\alpha}$.

The problem is calculating the parameter α of the Pareto distribution, based on:

- loss *estimates* $\{\hat{X}_1, \ldots \hat{X}_n\}$, which are known with uncertainty, and on
- prior knowledge/beliefs about the parameter (e.g. a range of possible values for α)

In the case of crisp data and no prior knowledge available, one can use MLE for a point estimate of the parameter α: $\alpha = n/(\sum_{j=1}^{n} \ln x_j - n \ln x_0)$, and calculate the standard error through the Fisher information matrix (a single value, in this case) or the bootstrap.

When data uncertainty is present and expert knowledge is available, we need a mechanism that gives more weight to data points that are more certain (e.g. settled or almost settled claims) and that strikes a balance between what the data say and what the prior knowledge says. Let us see how this problem is solved using three different methodologies: a heuristic approach, fuzzy set theory, Bayesian networks.

For illustration purposes, we consider an actual numerical example. We have generated 50 losses from a (single-parameter) Pareto distribution with threshold $x_0 = £1\text{m}$ and $\alpha = 2.5$. For each loss we have generated a random outstanding ratio (the amount yet to be paid) and we have applied distortions to each data point based on this outstanding ratio, reflecting the uncertainty on what the ultimate settlement will be, obtaining 31 losses in excess of £1m that are affected by uncertainty. We have then "forgotten" about how these figures were derived and we have tried to infer what the true value of α is. We have also assumed that prior knowledge/underwriting guidelines tell us that α should be in the interval [2,5].

Heuristic approach. Before looking at how fuzzy set theory and Bayesian analysis can help us, let us take a look at how this problem can be solved by using a heuristic approach, as actuaries will often make use of heuristic, sensible rules to deal with uncertainty.

Table 1. A sample of 31 losses randomly drawn from a single-parameter Pareto, which will be used to illustrate the different approaches to uncertainty. The uncertainty on each loss is expressed as a coefficient of variation (squared). For this particular data set, the maximum likelihood estimate of the parameter α is 1.94 (therefore lower than the 'real' value $\alpha = 2.5$). The standard error on the parameter can be calculated by standard MLE tools and turns out to be $\alpha = 0.35$.

Estimated loss	CV^2	Estimated loss	CV^2	Estimated loss	CV^2	Estimated loss	CV^2
5,603,857	0.270	1,934,597	0.148	1,463,245	0.051	1,193,121	0.218
4,556,176	0.393	1,882,209	0.231	1,440,697	0.180	1,191,055	0.220
3,763,176	0.230	1,778,076	0.168	1,440,557	0.149	1,190,398	0.113
2,673,192	0.238	1,741,700	0.089	1,412,091	0.053	1,164,280	0.129
2,256,042	0.196	1,669,716	0.172	1,357,288	0.062	1,158,949	0.092
2,062,769	0.062	1,570,758	0.250	1,252,113	−0.095	1,103,162	0.087
2,016,037	0.266	1,521,100	0.289	1,235,926	0.190	1,041,558	0.065
1,958,076	0.104	1,499,431	0.133	1,194,463	0.114		

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

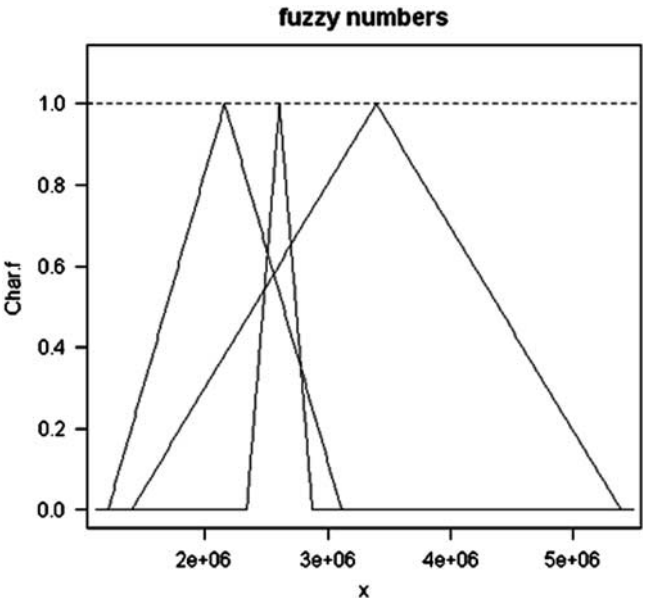


Figure 2. Some of the losses in Table 1, represented as triangular fuzzy numbers. The calculations and graphics of this section have been performed with the R package "fuzzyOP" by Semagul *et al.* (2009).

One simple approach is as follows.

(i) Use MLE to find α_{MLE} and $\mathrm{se}(\alpha_{\mathrm{MLE}})$ and parametric bootstrap to find the standard deviation $\sigma_{\mathrm{BS}}(\alpha)$ of different estimates of α for variations due to data uncertainty. Basically, to calculate $\sigma_{\mathrm{BS}}(\alpha)$ one needs to calculate α_{MLE} many times for different samples based on the initial data set. Unlike non-parametric bootstrap, each sample includes exactly the same losses but the value of each loss is drawn at random from a distribution centred around the best estimate and with a standard deviation proportional to the uncertainty on the loss.

(ii) Use a credibility approach to combine the above with underwriter's opinion: e.g., $\alpha^* = Z\alpha_{\mathrm{MLE}} + (1-Z)\alpha_{\mathrm{UW}}$ (subject to α^* being between 2 and 5). In this expression, α_{UW} is the default value proposed by the underwriter, and $Z = \frac{\sigma^2_{\mathrm{BS}} + \mathrm{se}^2(\alpha_{\mathrm{MLE}})}{\sigma^2_{\mathrm{BS}} + \mathrm{se}^2(\alpha_{\mathrm{MLE}}) + \mathrm{Var}(\alpha_{\mathrm{UW}})}$. The quantity $\mathrm{Var}(\alpha_{\mathrm{UW}})$ is the variance of the underwriter's estimate.

In the case of our numerical example, we have found $\alpha_{\mathrm{MLE}} = 1.94$, $\alpha_{\mathrm{UV}} = 3.5$, $\mathrm{se}(\alpha_{\mathrm{MLE}}) = 0.46$, $\sigma_{\mathrm{BS}}(\alpha) = 0.20$. If we assume that $\mathrm{Var}(\alpha_{\mathrm{UW}}) = 0.25$, this yields $\alpha^* = 2.72$ as our estimate of the Pareto parameter, roughly midway between our MLE estimate and the underwriter's prior guideline.

Despite the fact that this method does not have a solid theoretical foundation and many alternative solutions are possible, some of the methods that reinsurers use in practice are not too dissimilar from that above.

Using fuzzy set theory. Each non-settled loss X_i can be represented as a triangular fuzzy number $(a_1^{(i)}, a_2^{(i)}, a_3^{(i)})$, with width $a_3^{(i)} - a_1^{(i)}$ larger for losses that have a large outstanding percentage. Figure 2 shows examples of some of the 31 actual losses.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

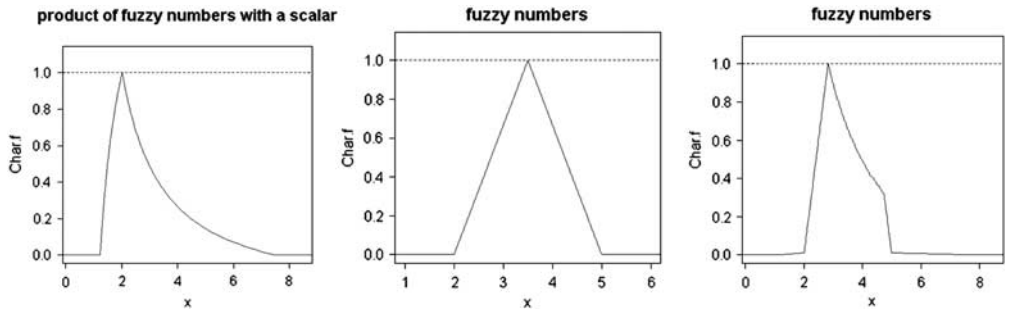


Figure 3. Left: fuzzy number resulting from the calculation of Equation $\alpha^* = n/\left(\sum_{j=1}^{n} \ln x_j - \ln x_0\right)$. Centre: fuzzy number representing the prior knowledge on α (or the underwriting guidelines). Right: fuzzy number representing the final estimate of α, obtained as the intersection of the membership functions at the left and the centre.

Fuzzy arithmetic can be used to produce an MLE-like estimate of the parameters[1]. E.g., for a Pareto distribution, we can use the equation $\alpha^* = n/(\sum_{j=1}^{n} \ln x_j - \ln x_0)$. As the logarithm is a monotonic function, it is easy to obtain the membership function of $\ln x_j$ for all losses x_j using Equation (1). The membership function of $\sum_{j=1}^{n} \ln x_j$ can then be obtained through Equation (2). Since the denominator is positive by definition (all x_j's are greater than x_0)), $n/(\sum_{j=1}^{n} \ln x_j - \ln x_0)$ is a well-defined fuzzy number with finite support. The width of this support (and actually the overall shape of the membership function) is related to the uncertainty on α. Notice that this does not necessarily decrease as $n \to \infty$ (see discussion in Section 2.5).

The overall result is a fuzzy number $\hat{\alpha}$ with membership function $\mu_{\hat{\alpha}}(x)$, whose shape can be seen in Figure 3 (left).

Prior knowledge can be incorporated by representing the parameter as another fuzzy number with membership function $\mu_{\alpha'}(x)$. In this experiment we have modelled the prior knowledge as a triangular number between 2 and 5, centered at 3.5, as in Figure 3 (centre).

The final result $\tilde{\alpha}$ can be obtained as the fuzzy intersection of the two estimates:

$$\mu_{\tilde{\alpha}}(x) = \mu_{\alpha' \cap \hat{\alpha}}(x) \tag{5}$$

and is shown in Figure 3 (right).

The choice of the intersection seems to make sense as it requires that both constraints (those coming from data and that coming from the soft knowledge) be satisfied at the same time.

One computational complication of the procedure is the fact that the membership function becomes more and more complex. Also, it is not clear how to extend this procedure to a case when we have a more complex distribution for which the ML solution is not analytical, e.g. with a GPD.

[1] It may seem odd to use the ML estimate (a probabilistic estimate) in the context of fuzzy set theory. However, the choice of the estimator is not the point here. Other estimators can be used as well. The point is that we need some estimator, and that this estimator will be calculated by performing some algebraic manipulations of fuzzy numbers, and therefore will be a fuzzy number itself.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

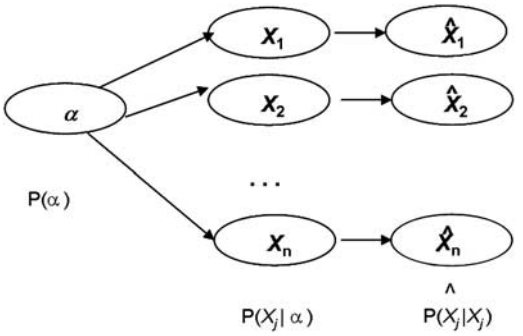


Figure 4. A Bayesian network representing a severity distribution with some prior knowledge on the parameters (α) and uncertainty on the loss amounts.

Using Bayesian networks. In a probability context, each non-settled loss can be represented by a probability distribution centred on the best estimate of the incurred amount. The dispersion of the distribution will be related to the uncertainty regarding the loss, and this in turn depends on a number of factors, such as the outstanding ratio (outstanding amount divided by incurred amount), the size of the loss, the opinion of the claims manager on the possible variability, the age of the loss, etc.

Prior knowledge on the parameters can be incorporated by a prior distribution on the parameters themselves.

In the case where there is prior knowledge on the parameters but no data uncertainty, an estimate of the parameters θ can be obtained by maximising the posterior probability of θ given the data, $\Pr(\vartheta|x_1, \ldots x_n) \propto \Pr(x_1, \ldots x_n|\vartheta)\Pr(\vartheta)$. If there is no prior knowledge, i.e. $\Pr(\vartheta)$ is the non-informative prior, the term $\Pr(\vartheta)$ can be dropped and the problem reduces to standard maximum likelihood estimation.

Now consider the case where there is data uncertainty as well as prior knowledge: in this case the input is given by the loss estimates $(\hat{x}_1, \ldots \hat{x}_n)$. For each estimate $\hat{x}_j$ we may have some additional information z_j, including, e.g., the outstanding ratio, the loss date, the type of injury, etc.). (For the sake of notational clarity the dependency on these variables is hidden in the equations below.)

This problem lends itself quite naturally to the formalism of Bayesian networks. This makes use of the following standard equation:

$$\Pr(\vartheta|e) = \beta \sum_{y} \Pr(\vartheta, e, y) \tag{6}$$

where $\beta = 1/\Pr(e)$, ϑ represents the *query variables* (in our case: the parameter(s) of the severity distribution), e is a realisation of E, representing the *evidence variables* (in our case: the estimated loss amounts), y is a realisation of Y, the *hidden variables* (in our case: the true loss amounts). Note that all of these quantities should be interpreted as vectorial random variables.

Equation (6) relates the posterior distribution $\Pr(\vartheta|e)$, which is the quantity we want to maximise, with the joint distribution $\Pr(\vartheta, e, y)$, which can be represented as a Bayesian network (see Section 2.2). There are many ways to specify a Bayesian network: one which usually leads to a compact representation is the causal model, where nodes representing causes are added first and nodes representing symptoms are added later. In our example, a simple causal representation is as shown in Figure 4.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

Some structural simplifications have been made in the model of Figure 4, by assuming that:

1. the variables $\hat{x}_j, j = 1 \ldots n$ (estimated losses) depend on α only through the variables $x_j, j = 1 \ldots n$ (true losses);
2. $\hat{x}_j$ does not depend on x_k if $k \neq j$;
3. all x_k are independent of one another given α (a common assumption).

It is always critical to decide which variables can be considered to be independent of the others. In most real-world cases, there is always an argument for objecting to the claim that two given variables are independent. Bayesian networks help introduce some discipline in this: every added link increases the complexity of the network and therefore makes inference harder, so new links should be added only if the growth in complexity is justified by the gain in accuracy. In practice, this means that one often gets rid of tenuous links (see discussion on compactness and node ordering in Russell & Norvig, 2003). In our example, it should be noted that Assumption 2 above is only a rough approximation.

The Bayesian network of Figure 4 corresponds to the following representation of the joint distribution:

$$\Pr(\alpha, x_1, \ldots x_n, \hat{x}_1, \ldots \hat{x}_n) = \Pr(\alpha) \prod_{j=1}^{n} \Pr(x_j|\alpha) \Pr(\hat{x}_j|x_j)$$

After some algebraic manipulations and the application of Bayes' rule, we obtain

$$\Pr(\alpha|\hat{x}_1, \ldots \hat{x}_n) = \Pr(\alpha) \prod_{j=1}^{n} \int \frac{\Pr(x|\hat{x}_j) \Pr(x|\vartheta)}{\int \Pr(x|\vartheta') \Pr(\vartheta') d\vartheta'} d\vartheta \tag{7}$$

Despite the complex form of this equation, its right-hand side is a function of the parameter(s) α, which we can maximise by numerical methods. In our case α is a single parameter (rather than a vector of parameters), the exponent of a Pareto distribution, and the posterior probability can therefore be maximised very efficiently. The three functions we need to know to calculate the integrals in Equation (7) are:

- the conditional probability $\Pr(x|\alpha)$, which for a Pareto distribution is $\Pr(x|\alpha) = \alpha \frac{x_0^\alpha}{x^\alpha}$;
- the prior probability $\Pr(\alpha)$. This should reflect our prior knowledge on the parameters. Here we have assumed a Beta distribution $\Pr(\alpha) = c(\alpha-2)^{a-1}(5-\alpha)^{b-1}$, and $a = b = 6$. This constrains α to be between 2 and 5 and sets the mean at 3.5. The value of a will drive the width of the Beta distribution (the large a is, the narrower the distribution);
- the conditional distribution $\Pr(x|\hat{x})$, or rather – to use the complete notation – $\Pr(x|\hat{x}, z)$. This is the IBNER distribution, and will normally be informed by the actual reserving experience. For simplicity, let us assume here that the IBNER distribution is a Gaussian distribution: $\Pr(x|\hat{x}, z) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\hat{x})^2}{2\sigma^2}\right)$, with $\sigma = CV(\hat{x}, z)\hat{x}$, where the coefficient of variation is estimated based on the actual data. This approach is usually called **empirical Bayesian approach,** as the prior distribution is actually estimated based on some previous, or external, experience, rather than made up from scratch.

By substituting these probabilities into Equation (7), we obtain the posterior distribution of α which we can then maximise – a simple one-dimensional numerical optimisation problem. The (numerically obtained) posterior distribution – with and without data uncertainty – is shown in Figure 5.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

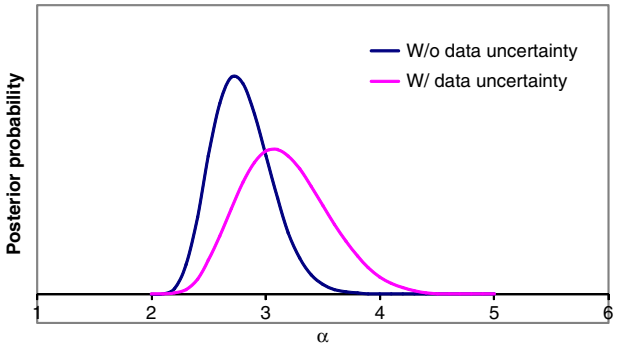


Figure 5. The posterior distribution for α for our example dataset, with and without data uncertainty. Notice that the distribution with data uncertainty is wider and is also closer to the underwriter's guideline $\alpha = 3.5$, probably because data uncertainty generally increases the uncertainty on the parameter calculated from the data and therefore puts more weight on the prior.

2.4 Another practical general insurance example: Model uncertainty

Model uncertainty is a particularly thorny issue as the problem of quantifying model uncertainty is as difficult as that of determining what the correct model is. There is a dumbed-down version of this problem, however, which seems to be more approachable: this is the case where there is only a finite number of competing models available, and we want to select the one that is better at predicting the results on an independent sample. In the example of Section 2.3, we might be undecided as to whether to model our losses with a Pareto (fat-tail) or an exponential (thin-tail) distribution, and we want our data to guide us in choosing between these two alternatives only. The general idea is to assign to each model a prior probability and calculate the posterior probability of the model given the data. This is explored for example in Bishop (2007).

The main idea is to compare a set of L models $\mathcal{M}_i$ where $i = 1, \ldots L$. Each model is a probability distribution, e.g. the probability distribution of losses. Each model has a prior probability $\Pr(\mathcal{M}_i)$. Given a training set $D = \{X_1, \ldots X_n\}$, we are interested in evaluating the posterior distribution:

$$\Pr(\mathcal{M}_i|D) \propto \Pr(\mathcal{M}_i)\Pr(D|\mathcal{M}_i) \tag{8}$$

In the equation above, the term $\Pr(\mathcal{M}_i)$ expresses a prior preference for different models, whereas the term $\Pr(D|\mathcal{M}_i)$ expresses the preference shown by the data themselves for different models, and is called the *model evidence*. Equation (8) can be used to select the best among a number of competing models.

In our case, we might have a preference for Pareto, because (for example) for 70% of our large clients the Pareto distribution turned out to be more successful in fitting our data. Therefore we might assign a 70% prior probability to the Pareto model and a 30% probability to the Exponential model. Again, this can be considered an empirical Bayesian approach. The posterior probability of the Pareto distribution is now given by:

$$\Pr(\mathcal{M}_{\text{Pareto}}|D) \propto \Pr(\mathcal{M}_{\text{Pareto}})\Pr(D|\mathcal{M}_{\text{Pareto}}) \tag{9}$$

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

and if this is larger than 0.5 we may choose a Pareto model (as opposed to an exponential) for modelling our data.

2.5 Comparing the fuzzy set theory and the Bayesian approaches to uncertainty

Several authors have compared the probabilistic and the fuzzy logic approach to uncertainty. One well-known example is the lively debate around the paper by Elkan (1994), which was about alleged inconsistencies in fuzzy logic and whether other paradigms are more adequate to replicate human reasoning (see, e.g., de Silva & Attikiouzel, 1994). The debate has spread to the actuarial community (Cummins & Derrig, 1997).

We here list some of the most contentious issues, and we summarise these points and the points that emerged in the previous sections, in Table 2.

Table 2. A comparison of the fuzzy set theory and the Bayesian approaches to uncertainty.

	Fuzzy set theory (FST)	Bayesian analysis (BA)
Key concept	Membership degree, i.e. the extent to which a certain object belongs to the set of objects having a certain property (e.g. large losses)	Probability, i.e. the size of the set of objects that have a given property as a proportion of all objects within a certain class
Arbitrariness	It is difficult to justify why a given membership function (e.g. triangular) is used to represent a number	There is some arbitrariness in the choice of the prior probabilities, although some of that arbitrariness may be removed by an empirical Bayesian approach
Dealing with uncertainty	FST is more suitable for representing linguistic vagueness rather than quantitative uncertainty. However, it can represent uncertainty *formally* by using fuzzy numbers	BA can represent uncertain knowledge by means of prior probability distributions, which can be updated as more information is made available
Parameter uncertainty	FST cannot be used to model parameter uncertainty as the law of large numbers does not hold	Parameter uncertainty can be easily dealt with by BA: e.g. the error on a maximum likelihood estimator is given by Fisher's information matrix
Model uncertainty	There is no obvious way to deal with model uncertainty within FST	Model uncertainty can be addressed by BA by assigning prior probabilities to models, when there is a finite set of models from which to choose. It is an open (and possibly unsolvable) problem in general
Data uncertainty	FST can formally deal with data uncertainty by representing data points as fuzzy numbers	Data uncertainty can be addressed by BA by modelling data points as probability distributions
Ease of use	Calculations on fuzzy variables are rather unwieldy	Calculations using probabilities are straightforward, and there is a wealth of methods for calculating (posterior) probabilities, e.g. Monte Carlo methods
Modelling complex situations	It is difficult to propagate uncertainties in complex situations, e.g. when a chain of calculations is involved, or where there are loops	Uncertainties can be propagated using techniques such as Bayesian networks

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The wrong type of uncertainty? The most common criticism of fuzzy set theory is probably that it is nothing but an unwieldy version of probability theory, and that everything that can be dealt with by fuzzy set theory can also be dealt with by Bayesian probability. When it comes to representing numbers, membership functions look very much like distribution probabilities indeed – so why not just use those? The classical response to this is that fuzzy set theory captures fuzzy membership to a set, whereas probability distributions aim at calculating the likelihood that an element belongs to a set in a crisp way. The conceptual difference is important: when we say that a loss is "large", we are using a fuzzy concept. No uncertainty is involved here – the loss may be exactly £130,000 and may have already been settled – but whether this loss belongs to the set of large losses is indeed fuzzy (and context-dependent). The author's view on this debate is that although there is no question that fuzzy set theory and probability theory are two conceptually distinct disciplines, each modelling different aspects of uncertainty, the interest of fuzzy set theory seems to be mostly logical and linguistic. The type of uncertainty we are concerned with in insurance is not of a linguistic kind: we may not be able to define crisply what a "large" loss is, but there is hardly a business case for attempting it! We are concerned with uncertainties on the amounts, not with semantic uncertainty. In all the applications of fuzzy set theory in the actuarial literature known to the author it is this second type of uncertainty which is addressed – this is the case, for example, in Cummins & Derrig (1997), where fuzzy logic is used to quantify the uncertainty on the risk-adjusted discount rate for losses, on the losses themselves, on the premiums, etc. In all these cases, fuzzy logic provides a formalism which produces reasonable results, but there is no reason why a richer probabilistic framework could not be used in its stead.

Is fuzzy set theory unwieldy? One of the main reasons of the recent success of Bayesian analysis lies in the fact that many tools have been developed for the calculation of posterior distributions: these include Markov Chain Monte Carlo techniques, Gibbs sampling, Metropolis-Hastings (see, e.g., Mackay (2003) or Roberts & Casella (2004) for a review of these methods), and all the techniques applicable to Bayesian networks to take advantage of their sparsity. On the other hand, anything beyond simple arithmetic is painful with fuzzy set theory. Until a proper calculus is developed, only elementary applications to general insurance will be possible.

Parameter uncertainty. Another issue with fuzzy set theory is that it deals poorly with parameter uncertainty. As observed for example in Nikolaidis *et al.* (1999), the law of large numbers does not hold for fuzzy numbers. The average of n fuzzy numbers with the same membership function has *exactly* the same membership function of the component numbers, regardless of n. On the contrary, in a probabilistic setting the average of n independent numbers will have a narrower distribution than its components, and this narrowing can be quantified by the central limit theorem. This leads to the apparent paradox that adding more information does not lead to a reduction in uncertainty in fuzzy set theory! The sum of fuzzy numbers can therefore be likened to the sum of perfectly correlated variables.

Model uncertainty. There is no obvious way by which fuzzy set theory can deal with model uncertainty. In Cummins & Derrig (1997), one of the stated reasons to use fuzzy set theory was actually that there is uncertainty as to whether capital asset pricing model is the correct model to produce a risk-adjusted discount rate for losses, r_L (see Cummins & Derrig (1997), Section 2.3). However, this fact is only used to justify an additional uncertainty on the result of this model – r_L itself. In other words, this model uncertainty is translated into a sort of parameter uncertainty. On the contrary, Bayesian theory incorporates model uncertainty by using the prior probability on models, and gives a recipe on how to choose between a number of competing models. It is a limited solution but it is further-reaching than anything one can do with fuzzy set theory.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

Other issues. As discussed, e.g., in Bender (1996), it is tricky to use fuzzy set theory when multiple levels of operations and decisions must be made. Furthermore, some of the basic definitions of fuzzy set theory, such as the definition of intersection, lead to counterintuitive results and may need *ad hoc* amendments for specific applications (Lemaire, 1990).

3 Temporal Reasoning about Risk

We have so far implicitly assumed that the environment is stationary, and that the only uncertainties lie in the limited amount of data available (leading to parameter and model uncertainty) and to the inaccuracy in the data themselves (data uncertainty). We now introduce the dimension of time, which contributes to uncertainty in at least two ways:

i. Evidence about risk unfolds gradually: in some cases, such as for long-tail liability classes, it may take decades to know what the final losses for a given underwriting year were (e.g. asbestosis claims in the 1990s);
ii. Risk itself is ever changing, which makes the evidence collected in the past gradually obsolete. This is a thornier issue, as many of the techniques for interpreting data rely on an assumption that the environment is static.

One obvious illustration of these issues is experience rating. To calculate the pure loss cost for a policy, claims data are collected over a number of years, and a frequency/severity model is fitted to the historical data. As new evidence becomes available, we are able to refine our estimates of the model parameters, while at the same time we may be forced to revise the parameters themselves or even the models as new trends (e.g., a drop in claims frequency due to improved risk control mechanisms) become evident. Another obvious example is reserving – the projection of the ultimate reserve is periodically updated in the light of new evidence, with occasional overhauls to take into account radical changes in the environment (e.g. the introduction of new claimant-friendly legislation for liability claims).

Several approaches have been used in computational intelligence to model time dependence (Sections 3.2, 3.3, and 3.4). Before describing these approaches, we need to give some definitions and explain some generalities on temporal models.

3.1 Probabilistic temporal models

Temporal aspects are usually addressed in computational intelligence by using a probabilistic temporal model (also known as a state-space model).

A probabilistic temporal model requires three ingredients:

i. *A transition model*. This is the conditional distribution $\Pr(X_t|X_{t-1})$, and fully describes time evolution in the case of first-order Markov processes.
ii. *A sensor model (aka observation model)*. This is the probability $\Pr(E_t|X_t)$, describing how the evidence is affected by the state of the world. We assume that the evidence only depends on the current state (*Markov property of evidence*), i.e. $\Pr(E_t|X_{0:t}, E_{0:t-1}) = \Pr(E_t|X_t)$. The state of the world is supposed to be "hidden", that is, measurable only through the observation model.
iii. *A prior probability*. This is the probability $\Pr(X_0)$ over the states at time zero.

Notice that the transition and the observation model themselves must often be learned from observations.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The probabilistic model above immediately sheds some light on the notion of time dependence, and it allows us to define different levels of stationarity in terms of the transition model (note that the time t in the following relations is discrete):

- a *static system* is one where the state of the system does not change: $X_t = X_{t-1}$;
- a *stationary system* is one where $\Pr(X_t|X_{t-1})$ does not itself depend on t, that is, $\Pr\left(X_{t+k}|X_{t+k-1}\right) = \Pr(X_t|X_{t-1}) = F$ for all k;
- a *non-stationary system* is one where $\Pr(X_t|X_{t-1})$ depends on t, that is, $\Pr(X_t|X_{t-1}) = F_t$, where the distribution F_t is often not known. In this case, if the time dependence is "slow", it may still be possible to model the system as locally linear (or constant) with slowly changing parameters;
- the worst case of all is of course the one where events at different times are completely independent of each other: $\Pr(X_t|X_{t-1}) = \Pr(X_t)$. In this case prediction is, by definition, impossible. A conceptually different but practically similar case is that in which there exists some dependency but the distribution F_t changes so quickly that the past experience becomes almost immediately irrelevant for prediction purposes.

Many results of probabilistic temporal models assume that the system is *stationary* and has the *Markov* property of evidence. However, they can usually be extended at least to the case where the system is quasi-stationary (i.e., it changes slowly).

The ingredients of the probabilistic temporal model can be combined to specify the whole joint distribution:

$$\Pr(X_0, X_1, \ldots X_t; E_1, \ldots E_t) = \Pr(X_0) \prod_{j=1}^{t} \Pr\left(X_j|X_{j-1}\right) \Pr\left(E_j|X_j\right) \tag{10}$$

The temporal model above allows to solve several different inference problems, including (list borrowed from Russell & Norvig, 2003):

- The *filtering problem*: i.e., calculating the belief state (the posterior distribution $\Pr(X_t|e_{1:t})$ of the current state given all the evidence to date); for example, X_t may represent the number of losses expected in year t given all claims reported up to t (a non-trivial problem when IBNR claims are a significant portion of the total claims);
- The *prediction problem*: i.e., calculating the posterior distribution of a future state given the evidence until now, $\Pr\left(X_{t+k}|e_{1:t}\right)$. A typical example is experience rating – the calculation of the future distribution of claims given historical losses. Note that prediction can be seen as filtering without the benefit of new evidence.
- The *smoothing problem:* i.e., calculating the posterior distribution of a past state given all evidence up to the present, $\Pr(X_k|e_{1:t})$, with $k<t$. A classical example is reserving, where we need to keep updating our estimate of losses in year k for many years after that – until we are fully sure that no more claims can be reported for that period.

The solution to any of these problems usually comes in the form of a recursive equation, with a repeated application of the Bayes' rule and of the Markov property of evidence. As an example, the recursive equation for filtering can be written as:

$$\Pr(X_{t+1}|e_{1:t}) = \alpha \Pr(e_{t+1}|X_{t+1}) \sum_{x_t} \Pr(X_{t+1}|x_t) \Pr(x_t|e_{1:t}) \tag{11}$$

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The recursive equation for prediction is instead:

$$\Pr(X_{t+k+1}|e_{1:t}) = \alpha \Pr(e_{t+1}|X_{t+1}) \sum_{x_{t+k}} \Pr(X_{t+k+1}|x_{t+k}) \Pr(x_{t+k}|e_{1:t}) \tag{12}$$

so that prediction can be seen as filtering without the benefit of new evidence.

We now look at different examples of temporal probabilistic models: hidden Markov models, Kalman filtering and dynamic Bayesian networks.

3.2 Hidden Markov models (HMMs)

A hidden Markov model (HMM) is a temporal probabilistic model, in which the state of the process is described by a single discrete random variable (possibly, a vector variable). This representation allows "a simple and elegant matrix implementation" of the generic model (Russell & Norvig, 2003). In the equations below it is assumed that the system is Markov and stationary. As usual, the results can easily be extended to the quasi-stationary case.

The transition model is represented as

$$T_{ij} = \Pr(x_t = j|x_{t-1} = i)$$

The sensor model is represented as the diagonal matrix whose diagonal elements are given by the values of $\Pr(e_t|x_t = i)$.

$$O_t = \text{diag}\,(\Pr(e_t|x_t = i))$$

The updating equation is

$$f_{1:t+1} = \alpha O_{t+1} T^T f_{1:t}$$

where T^T is the transpose of T and $f_{1:t} = \Pr(x_t|e_{1:t})$ is the so-called "forward message".

3.3 Kalman filtering

Kalman filtering is a type of regression analysis enhanced with a mechanism for updating the regression parameters, which are assumed to evolve in time. A classical application of Kalman filtering is radar tracking: once an object is identified at time t, its position at time $t + \Delta$ is determined by combining the noisy observation at time $t + \Delta$ with an assumption on where the object is expected to be found at time $t + \Delta$ based on its position at time t.

Although various extensions of the Kalman filter can be found in the literature, the standard model of Kalman filtering requires the system to have the *Markov property of evidence* and to be *stationary*: furthermore, it is assumed that the current state follows a multivariate Gaussian distribution. The **transition model** for the Kalman filter is therefore $\Pr(x_{t+1}|x_t) = \mathbb{N}(Fx_t, \Sigma_x)(x_{t+1})$, where F and Σ_x are matrices describing the linear transition model and transition noise covariance. The **sensor model** is $\Pr(z_t|x_t) = \mathbb{N}(Hx_t, \Sigma_z)(z_t)$, where H and Σ_z are matrices describing the linear sensor model and sensor noise covariance.

The update equation for the mean μ_t of x_t is now:

$$\mu_{t+1} = F\mu_t + K_{t+1}(z_{t+1} - HF\mu_t) \tag{13}$$

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

and the update equation for the covariance is:

$$\Sigma_{t+1} = (I - K_{t+1})(F\Sigma_t F^T + \Sigma_x) \tag{14}$$

where K_{t+1} (the Kalman gain) is given by:

$$K_{t+1} = (F\Sigma_t F^T + \Sigma_x)H^T(H(F\Sigma_t F^T + \Sigma_x)H^T + \Sigma_z)^{-1} \tag{15}$$

Equations (13), (14) and (15) are horrible to behold, but a few moments' look at them will remind the actuary of credibility theory: the Kalman gain is the credibility given to the new observation z_{t+1} (or rather, the difference between this and the predicted observation $HF\mu_t$), relative to the predicted state at time $t+1$, $F\mu_t$.

As mentioned at the beginning of this section, the standard model, which is Markov, stationary and with a Gaussian transition and sensor model, has been extended in several directions.

- The **extended Kalman filter** (**EKF**) incorporates non-linearities by only requiring that the model is locally linear. In practice, this means that the system is only locally stationary (Zarchan & Musoff, 2005; Ribeiro, 2004). This works well for smooth, well-behaved systems – that is, systems that do not have abrupt changes;
- The **switching Kalman filter**, which attempts to model strongly non-linear behaviour such as obstacle avoidance by using multiple Kalman filters among which it is possible to switch (Ghahramani & Hinton, 1998; Hamilton, 1989).

Both cases above are special cases of dynamic Bayesian networks, that provide a far more general solution to the problems addressed by the Kalman filter.

3.4 Dynamic Bayesian networks (DBNs)

Dynamic Bayesian networks are Bayesian networks representing temporal models. They are completely general in the sense that they do not add any further assumptions to those of general temporal models.

The joint probability distribution of Equation (10) already provides a full specification of a DBN, as it clarifies the dependency structure of the network. This can be rendered graphically as in Figure 6, both in the rolled-up (recursive) version and in the unrolled version. The term "unrolling" refers to the development of the recursion.

The reason why the recursive formulation is important is that it allows to store the whole probabilistic model in constant space (the network does not grow with time) and update the network in constant time. The recursive formulation is equivalent to the unrolled one only if the system is stationary.

The Kalman filter is a special case of a DBN. This is illustrated in Figure 7, where a classical Kalman filter is shown as a rolled-up DBN.

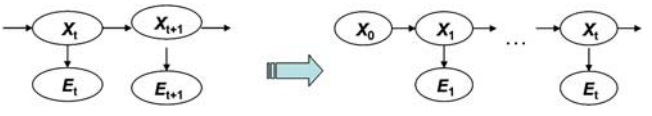


Figure 6. A dynamic Bayesian network, in recursive (left) and unrolled (right) form.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

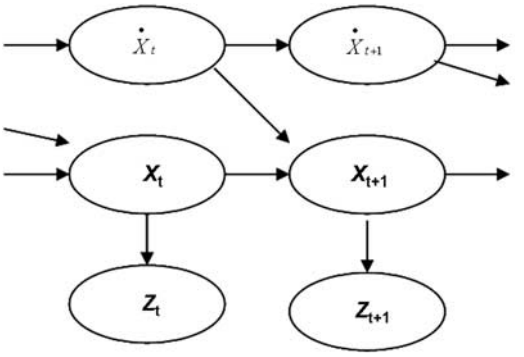


Figure 7. Representation of a classical Kalman filter where both the position (x_t) and the speed (x_t) of an object are tracked as a dynamic Bayesian network are shown.

Table 3. Example of a run-off triangle. The amounts in the triangle represent the incremental paid amounts for each origin year and for each development year (0 to 4). Amounts shown are in £1,000.

Origin year, s	Inflation index	Volume index	0	1	2	3	4
1	0.598	1	753.5	648.9	311.7	173.5	71.3
2	0.665	0.899	642.3	648.4	249.7	206.5	
3	0.748	0.858	715.8	661.1	309.4		
4	0.853	0.863	841.6	862.6			
5	1	0.813	968.8				

3.5 Practical general insurance example: Reserving

A pioneering application of probabilistic temporal models to general insurance is that of De Jong & Zehnwirth (1983) in their work on reserving. The general case is quite complex so for illustrative purposes we will focus on the simple example that can be found in Section 5 of their paper.

This is based on the run-off triangle in Table 3 (borrowed from De Jong & Zehnwirth, 1983), that gives the amount paid for each development year and for each origin year, taking into account some measure of inflation (the inflation index) and of exposure (the volume index).

In this example, the observation (evidence) variables z_t are the successive diagonals of the paid triangle, as they represent the evidence collected in the latest calendar year. In the example above,

$z_1 = (753.5)^T$

$z_2 = (642.3, 648.9)^T$

$z_3 = (715.8, 648.4, 311.7)^T$

$z_4 = (841.6, 661.1, 249.6, 173.5)^T$

$z_5 = (968.8, 862.6, 309.4, 206.5, 71.3)^T$

The underlying parameters are in this simple case the ultimate incurred claims x_s for each origin year s.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The **observation equation**, which connects the paid amounts to the ultimate claims x_s is (in this extremely simplified model):

$$z_t(d) = (d + 1)\exp(-d)n_{t-d}\lambda_t x_{t-d} + u_{t-d}$$

where u_{t-d} is a zero-mean random noise; $(d + 1)\exp(-d)x_{t-d}$ is the expected value of payments for origin year $t-d$ and development year d, which is decreasing monotonically as a function of the development year; n_{t-d} is the exposure index for origin year $t-d$; λ_t is the inflation index for calendar year t;

The transition model is given by the simple random walk model:

$$x_s = x_{s-1} + v_s$$

where v_s is again a zero-mean random noise.

We also need an initial state x_0 corresponding to the unobserved origin year $s = 0$.

In De Jong & Zehnwirth (1983), this problem is approached by using extended Kalman filtering, with the structural parameters $(F, H, \ldots)$ varying with time and therefore indicated below as $F(t)$, etc. The list below explains what these parameters look like in this simple case:

- the matrix $F(t)$ is a matrix with zeros everywhere except at elements $(1,1),(2,1),(3,2),\ldots(t,t-1)$ where its value is 1;
- the matrix $\Sigma_x(t)$ is written as $G(t)V(t)G^T(t)$ where $G(t)$ is a $t \times 1$ matrix whose only non-zero element is (1, 1) which can be set to 1, and $V(t)$ was estimated by MLE methods as $V(t) = 626.3$;
- the initial state x_0 corresponding to origin year $s = 0$ and its covariance matrix Σ_0 are back-calculated with MLE methods assuming that a number of diagonals are already given, obtaining the values $x_0 = 1146.3$ and $\Sigma_0 = 2667.3$;
- the covariance matrix $\Sigma_z(t)$ is a diagonal matrix whose diagonal values are the first t values from diag(16283.6, 12204.8, 4509.3, 2755.7, 876.8)

Kalman's update rules (Equations (13), (14), (15)) are then used to produce the projected values of the claims triangle. The estimates for the states xt generated by the Kalman filter are shown in Table 4:

Table 4. The estimates of the states x_t at different times. The quantities shown in italic and smaller case are the standard errors.

	Time				
	1	2	3	4	5
x5					1,147.30
					53.70
x4				1,141.20	1,146.20
				55.90	*48.90*
x3			1,156.40	1,141.10	1,140.60
			57.20	*51.20*	*45.30*
x2		1,157.50	1,157.10	1,140.20	1,138.50
		57.20	*52.30*	*47.30*	*42.80*
x1	1,154.00	1,158.60	1,156.90	1,142.50	1,139.40
	55.40	*52.20*	*48.20*	*44.40*	*41.20*

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

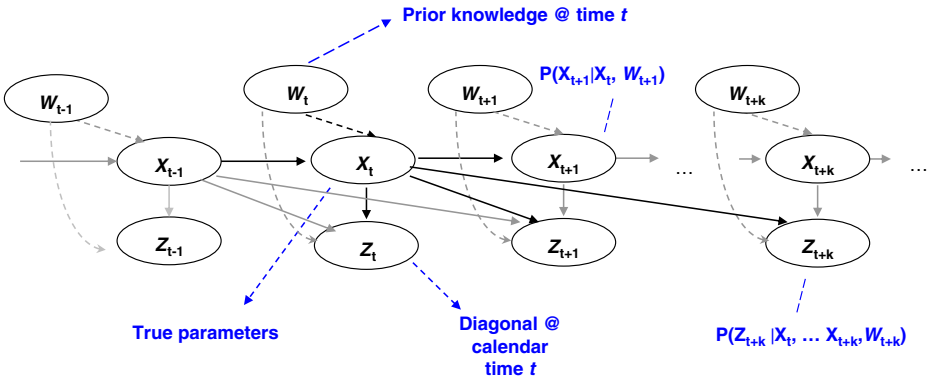


Figure 8. A dynamic Bayesian network for loss reserving. The example in De Jong & Zehnwirth (1983) does not require the variables W_t depicting prior knowledge, which are therefore connected to the rest of the network with dashed lines. All the edges except those to and from X_t are grayed.

Kalman's update rules (Equations (13), (14), (15)) are then used to produce the projected values of the claims triangle. The structural parameters $(F, H \ldots)$ vary slowly with time and therefore this can be regarded as an application of the Extended Kalman Filter (EKF).

The results obtained in De Jong & Zehnwirth (1983) could also be obtained by a dynamic Bayesian network like that in Figure 8. The model can be enhanced by adding nodes which represent (possibly discrete) random variables W_t incorporating prior knowledge on the things that might have an impact on the general level of claims for a given origin year (e.g. a bad winter) or on the amounts paid from a given year on (for example, new legislation on court awards for bodily injury liability). The transition model and the sensor model including these additional variables need not be linear or Gaussian.

Note that the model shown in Figure 8 requires some rearrangement before being cast as a standard transition/sensor model, since X_t affects $Z_t, Z_{t+1}, \ldots Z_{t+k}$ (k is the maximum number of development periods such that the hidden variables X_t are relevant to the evidence variables. In other words, we can assume that $\Pr(Z_{t+j} \mid X_t, \Omega_t) = \Pr(Z_{t+j} | \Omega_t)$ for $j > k$), where Ω_t denotes the collection of all other variables). Note that, as long as k is finite, this process can always be transformed into a first-order Markov process by a suitable redefinition of variables.

3.6 Comparison of different methods for temporal reasoning

In the previous sections we have looked at probabilistic reasoning over time, and at some specific models for dealing with it. We here summarise the advantages and disadvantages of the different methods:

- Kalman filtering crucially depend on the assumption that both the sensor (observation) model and the transition model are linear multivariate Gaussian. Although the model can be extended into different directions (EKF, switching Kalman filter), the Kalman filter is simply not general enough to model all interesting situations in general insurance. In the practical case we looked at (reserving), many of the most relevant uncertainties over time depend on factors that are strongly non-linear and non-Gaussian, such as future changes of legislation, future changes in the reserving

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

guidelines, sudden change in the riskiness of an account, etc. Some (although not all) of these non-linearities can be incorporated in a general probabilistic model, but it is awkward to address them with Kalman filters;

- dynamic Bayesian networks are a more powerful method than Kalman filters:
  - all Kalman filters can be represented as dynamic Bayesian networks or as hidden Markov models, but not vice versa;
  - serious non-linearities can be addressed by dynamic Bayesian networks with both discrete and continuous variables;
  - numerical methods such as Markov Chain Monte Carlo can be used for approximate inference in dynamic Bayesian networks;
- hidden Markov models and dynamic Bayesian networks are equivalent formulations in the discrete case – however, dynamic Bayesian networks are more compact and allow gains from the sparsity of the connections between the random variables.

It should be stressed that *all* temporal models as we have presented them here have some limitations. In all cases, a prior model of the possible future changes is needed. This can be linear, or locally linear, or allow for strong non-linearities, but in all cases the information we have on the transition model is crucial. If $\Pr(X_t|X_{t-1}) = \Pr(X_t)$, or if, less radically, $\Pr(X_t|X_{t-1})$ changes too quickly, prediction is impossible and all we can do is rely on the very latest data. In practice, it is seldom the case that the past is completely irrelevant to the future, but there are times where the rules of the game change quickly and our ability to model future changes is very limited.

4 Making Decisions in an Uncertain Environment

Ultimately, risk professionals need to understand risk because they have to make informed decisions. They are, in the language of computational intelligence, "intelligent agents", because they learn from the environment and they act on it.

Except in trivial cases, making decisions is not a one-off activity but a regular one: for example, an insurance company has to decide on the amount of reserves to put aside for future claims on a regular basis, perhaps on a quarterly basis, and has to decide on price changes perhaps on a monthly basis, if not almost in real time.

In this section we address the problem of a risk agent making *sequential* decisions in an uncertain environment. This is a well-known problem in computational intelligence: "Designing an intelligent agent which can move in an environment making the best decisions – i.e., the decisions which maximise utility" (Russell & Norvig, 2003). Notice how the concept of utility has crept in here in a purely computational intelligence context, before even considering applications to the actuarial context.

It is helpful to have a specific example in mind while reading this section, therefore we will assume that our intelligent agent is an insurance company and that its objective is to maximise the amount of capital in the long run. The insurance company's name is R-Age (formerly Risk Agents Ltd), a small outfit which sells a single product so that its financial fate is completely determined by the success of this product. The company has a very simple rating system – it charges each customer the same fee for the policy! There is little incentive to have a better rating system as R-Age happens

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

to be the only player selling this product in its very niche market, and its customer base is quite homogeneous anyway.

4.1 Individual agents and the environment

Each risk agent is assumed to be autonomous and to incorporate strategies to interact with the environment. The type of environment in which the agent moves is crucial. An environment can be (see Russell & Norvig, 2003; Wooldridge, 2001):

Fully v partially observable. A fully observable environment is simply one in which the agent always knows where it is. In the insurance world we usually have partially observable environments as we do not usually know for sure what our true underlying financial position really is (this is perhaps clearer in the case of long-tail business, where it takes a long time for the financial position of a given underwriting year to become finalised).

Deterministic v stochastic. A deterministic environment is one where the outcome of our actions is known for sure (in terms of the state the agent will be in). This is never the case in insurance: once for example we set the price of one policy we cannot know for sure what the volume sold and the claims to be paid will be.

Static v dynamic. A static environment is simply one that does not change with time, except for the actions of the agent. A dynamic environment has other processes operating in it, and these change the environment in ways that the agent cannot control. Needless to say, the insurance environment is highly dynamic. However, in some cases it may be necessary, or convenient, to assume that the environment does not change in the short term.

Discrete v continuous. A discrete environment is one with a fixed, finite number of states and of actions that can be performed. Again, the type of environments we are interested in are continuous. However, this distinction is not crucial, as we can usually approximate continuous environments with discrete ones.

Finally, note that the term "environment" may be misleading, as in insurance we normally mean it to encompass the context where an insurance company operates: its competitors, the judicial climate, the regulation, etc. However, the meaning in which we use it here is broader, and includes everything that affects the financial results and that is not determined directly by the actions of the agent. Specifically, *it includes the distribution of losses*.

4.2 Markov decision processes

Although there is no such thing as a fully observable environment in general insurance, it is useful to start considering this type of environment, because it prepares the ground for dealing with partially observable environments.

In this section we therefore assume that the environment is fully observable, stochastic, static and discrete. In our example, we are basically assuming that R-Age is so good at pricing that it is able to determine the parameters of the loss distribution with infinite accuracy. It is also so market-wise that it knows the demand curve for that policy with infinite accuracy, so that it knows exactly what the volume sold will be for a given price. However, underwriters at R-Age fall short of omniscience

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

as they cannot predict the actual losses which will be incurred next year – only the distribution from which they are drawn.

This situation can be modelled by a *Markov decision process*, which is defined by:

- An *initial state* S_0
- A Markovian *transition model* $T(s, a, s')$: at each time t, an agent will be in state s and will be able to perform an action a. As a consequence, it will move to state s' with probability $T(s, a, s')$. The transition model is Markovian in the sense that the probability of reaching s' from s is assumed only to depend on s and not on the states that precede s.
- A *reward function* $R(s)$, establishing the reward received by the agent in state s.

How is this relevant to R-AGE? The state s in which R-AGE is at time t may be simply defined as a pair $s = (C_t, V_t)$ where C_t is the capital available and V_t is the volume sold at time t. Time is discretised and is in years. The initial state is the initial capital C_0 and the initial volume V_0. Actions are simply decisions on prices: therefore a is another symbol for P_t, the price charged at time t. The transition probability $T(s, a, s')$ is the probability that if the price is set to a the capital will become C_{t+1} (the volume V_{t+1} is deterministic in this simplified setting), and will depend exclusively on the aggregate losses S_{t+1} incurred between t and $t + 1$ (for simplicity, we assume that these include expenses):

$$C_{t+1} = C_t + V_t P_t - S_t \tag{16}$$

If we know the distribution of S_t (e.g., a lognormal distribution) we know the distribution for C_{t+1} given C_t and P_t: that is, the transition probability $T(s, a, s')$.

Equation (16) also allows us to define the reward for each state as:

$$R_t = V_t P_t - S_t \tag{17}$$

which is simply the profit during year t.

Having defined what a Markov decision process is, we need a method to decide how good a *strategy* σ (i.e., a prescription specifying what the agent should do in any state it might reach) is. $\sigma(s)$ is the action prescribed for the agent when this is in state s. We therefore need an ordering on the set of strategies – a *utility function*. Given the stochastic nature of the environment, the quality of a strategy σ is measured by its *expected* utility, where the expectation is calculated over all possible *environment histories* generated by σ. The term "histories" might be a bit misleading here, as it is used differently from normal actuarial practice: by environment history we mean in this context a sequence $[s_1, s_2, \ldots s_n, \ldots]$ of *future* states.

An *optimal strategy* is thus one that achieves the highest expected utility. Interestingly, if the utility function is constructed wisely, strategies which yield high rewards in the short-term but are too risky will be penalised. As a result, there is a familiar balance to be struck between risk and reward, and this arises simply as a consequence of the stochastic nature of the problem: "*the careful balancing of risk and reward is a characteristic of MDPs that does not arise in deterministic search problems; moreover, it is a characteristic of many real-world decision problems*" (Russell & Norvig, 2003).

How can we assign a utility function $U_h([s_0, s_1, \ldots s_n])$ to environment histories $[s_0, s_1, \ldots s_n]$? This depends, first of all, on whether the horizon for decision-making is finite (only rewards accumulated

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

up to year N are relevant) or infinite (the whole future history is relevant). Finite horizon problems are more complicated because they are inherently non-stationary: the optimal decision will depend on how close you are to the deadline. We will assume that

- the horizon is infinite – which simply means in our case that one does not know *a priori* how long one is going to stay in business;
- the agent's preferences between state sequences are stationary, meaning that if the utility of $[s_1, s_2, \ldots s_n, \ldots]$ is greater than that of $[s'_1, s'_2, \ldots s'_n, \ldots]$, then the utility of $[s_0, s_1, s_2, \ldots s_n, \ldots]$ is greater than that of $[s_0, s'_1, s'_2, \ldots s'_n, \ldots]$.

It can be shown that under these conditions there is only one way in which we can assign utilities to sequences:

$$U_h([s_0, s_1, \ldots s_n]) = \sum_{j=0}^{\infty} \gamma^j R(s_j) \tag{18}$$

where γ, the discount factor, is a number between 0 and 1. The choice $\gamma = 1$ is usually not acceptable as it leads to infinite utilities.

Note that without ever mentioning the time value of money and without any other actuarial consideration, but simply assuming an infinite horizon and the stationarity of the agent's preference over state sequences, we have found that the only sensible way in which we can define utility is as the sum of the present value of future rewards!

We are now in a position to define the optimal strategy σ as the one which maximises the expected utility as defined in Equation (18):

$$\sigma^* = \arg\max \mathbb{E}\left(\sum_{t=0}^{\infty} \gamma^t R(s_t)|\sigma\right) \tag{19}$$

In the case of R-AGE, the utility is simply the present value of future profits. As to what discount factor γ to use, that is a usual problem – in a normal situation, one solution is using the return expected by investors for the level of riskiness of this company (in this case, this policy). In the specific case addressed in this section – for which the only real uncertainty is in the randomness of the aggregate losses, whose parameters are assumed to be known with certainty – it would probably be more appropriate to choose a lower discount rate, such as the risk-free discount rate plus a "premium" related to, say, the coefficient of variation of the aggregate loss distribution.

How do we solve Equation (19)? The utilities of states are the unique solutions of the so-called Bellman equation, which relates the utility of one state s to the discounted utility of the state of its neighbours' s':

$$U(s) = R(s) + \gamma \max_a \sum_{s'} T(s, a, s')U(s') \tag{20}$$

The Bellman equations (one for each state) can be solved iteratively, by starting with an arbitrary value $U_0(s)$ of the utilities and then updating the utilities with this equation, called a Bellman update:

$$U_{i+1}(s) \leftarrow R(s) + \gamma \max_a \sum_{s'} T(s, a, s')U_i(s') \tag{21}$$

It can be easily shown that this iteration, which has the form $U_{i+1} \leftarrow BU_i$, is a contraction in the space of the utility vectors, and has therefore a unique fixed point, to which the iteration ultimately converges.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The speed of the convergence is related to the value of γ: as one would expect, the smaller γ is, the quicker the algorithm converges to the solution (see Russell & Norvig (2003), Section 17.2, for a more detailed discussion of the algorithm).

The discussion above applies to the case where the number of states is finite. This is not the case of R-Age, because the premium of a policy can be any number. However, the real world case can be approximated with a scenario involving a finite number of states.

A more important issue is the assumption that the environment be fully observable. As we will see in the next section, this is not a satisfactory assumption, and to address this case we need the more advanced formalism of partially observable MDPs.

4.3 Partially observable Markov decision processes (POMDP)

In practice, R-Age will not be able to determine its state (C_t, V_t) with certainty for $t > 0$: e.g., when writing long-tail business some claims will only be known or settled many years after their occurrence. The estimate of the current capital depends on the reserve estimates, which are uncertain (owing to IBNR, IBNER, etc). As a consequence, the true underlying financial position of the company is not known with certainty. Furthermore, the transition matrix $T(s,a,s')$ will be uncertain, as the loss distribution is built upon a finite set of observations[2].

In the language of intelligent agents we say that the environment is not fully observable and the agent does not know for certain which state it is in. This can be modelled by the so-called partially observable MDPs, or POMDPs ("pom-dee-pees").

A POMDP requires the following ingredients:

- A *transition model* $T(s,a,s')$
- A *reward function* $R(s)$
- An *observation model* $O(s,o)$ that specifies the probability of perceiving the observation o in state s
- A *belief state*, i.e. a mapping between actual states s and the probability $b(s)$ assigned to each of them. The belief state can also be defined as *the posterior distribution over the current state, given all evidence to date*: or, using the notation of our section on temporal models, $\Pr(X_t|e_{1:t})$.

As new observations are made, the belief state is updated, according to this equation, which is simply an application of Bayes' theorem, with a non-standard format.

$$b'(s') = \alpha O(s', o) \sum_s T(s, a, s') b(s) \tag{22}$$

which can be abbreviated as

$$b' = \textsc{Forward}(b, a, o) \tag{23}$$

[2] When the only uncertainty is that on $T(s,a,s')$, because losses are reported quickly, a workaround is to incorporate the uncertainty on the transition matrix in the MDP model. Each row in the transition matrix $T(s,a,s')$ is a sample from a Dirichlet distribution (a multivariate extension of the Beta distribution), which has expectation given by the transition beliefs, sums to 1 but also has uncertainty which shrinks as more observations are made (Roberts, 2011).

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

What makes POMDPs work is that the optimal strategy σ^* depends only on the agent's current belief state and not on the state it is in. This is how a POMDP makes decisions:

POMDP *decision process*

1. Given the current belief state b, execute $\sigma^*(b)$ (the optimal strategy given the belief state)
2. Receive observation O
3. Set the current belief state to $b' = \textsc{Forward}(b, a, o)$ and repeat

Interestingly, the problem of solving a POMDP on a physical state space can be reduced to that of solving an MDP on the corresponding belief state space:

- The probability of reaching b' from b given action a is

$$\begin{aligned}\tau(b,a,b') &= \Pr(b'|a,b) = \\ &= \sum_o \Pr(b'|o,a,b)\Pr(o|a,b) = \\ &= \sum_o \Pr(b'|o,a,b)\sum_{s'} O(s',o)\sum_s T(s,a,s')b(s)\end{aligned} \tag{24}$$

  where $\Pr(b' \mid o,a,b)$ is 1 if $b' = \textsc{Forward}(b,a,o)$, 0 otherwise.

- The reward for belief state b can be defined as:

$$\rho(b) = \sum_s b(s)R(s) \tag{25}$$

  (that is, the reward for a belief state is the expected value of the reward for individual states)

These equations have the same form as the definition for fully observable MDPs (Section 4.2), the only difference being that the state s is replaced by the *belief state* $b(s)$. We have thus reduced the problem of solving POMDPs in the state space to the problem of solving MDPs in the belief space. It can be shown that an optimal strategy $\sigma^*(b)$ for this MDP is an optimal strategy for the original POMDP.

The difficulty, however, is that the space of belief states is continuous and usually high-dimensional. Solutions for this case have been developed as well, but their computational complexity becomes rapidly prohibitive as we increase the number of states: the problem falls in the class of PSPACE-hard problems, which are widely thought to require exponential time to solve (Papadimitriou, 1994). Luckily, approximate methods for solving POMDPs have been developed, making use of the so-called "dynamic decision networks" (DDNs), which we will describe in the next section.

Before going into that, however, let us see how POMDPs apply to R-Age. As we said at the start of this section, the company will not know its financial position for each year with certainty, mainly because of reserve uncertainty. We may assume that the financial position at the initial year (C_0, V_0) will be known, but any successive year will be uncertain. How does the company decide what to do in order to maximise capital in the long run?

The belief state b in any year different from 0 will be determined among the other things by the claims reserves.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

Instead of having a state (C_t, V_t), in the partially observable case we have a belief state b(s) which gives the probability of each state and that we can write for R-AGE as $\Pr(C_t, V_t, \boldsymbol{\mu})$, where $\boldsymbol{\mu}$ is a vector including all the parameters for which there is some uncertainty (e.g., the parameters of a lognormal distribution modelling the total losses. $\boldsymbol{\mu}$ can also be written as $\boldsymbol{\mu_t}$ to allow for parameters varying in time).

The reward for the belief state $b(s)$ is defined as (in the case when states can be discretised)

$$\rho(\Pr(C_t, V_t, \boldsymbol{\mu})) = \sum_{C_t, V_t, \boldsymbol{\mu}} \Pr(C_t, V_t, \boldsymbol{\mu})(V_t P_t - S_t)$$

where the sum operator is to be intended as an integral wherever the variables are continuous.

The optimal strategy for R-AGE does not depend on the current state but on what we know about the current state, exactly as in the general case. The three-step *POMDP decision process* can be rewritten for R-AGE as follows.

1. Given the current estimated distribution $\Pr(C_t, V_t, \mu)$, set the policy premium equal to $P^*(t)$
2. Receive observations on reported losses, paid losses, volumes sold, etc.
3. Revise the reserve estimates, the total loss parameters, etc. thus producing a new estimate $\Pr(C_{t+1}, V_{t+1}, \boldsymbol{\mu})$, and repeat.

4.4 Dynamic decision networks

As mentioned in the previous section, POMDPs are difficult to solve. However, approximate methods are available that provide sub-optimal solutions. Specifically, *dynamic decision networks* provide a comprehensive approach to agent design for partially observable stochastic environments.

In this approach, the agent is represented as a *dynamic decision network* (DDN), i.e. a dynamic Bayesian network (Section 3.4) augmented with decision and utility nodes. **Decision nodes** are nodes where decision makers have a choice of action – e.g., what price R-AGE should charge for its policies. **Utility nodes** are related to the utility of the state (including the rewards received). A filtering algorithm similar to that described in Equation (11) is used to incorporate new evidence, the effect of actions (which is treated as evidence) and to update the belief state, e.g. (in the case of R-AGE) the probability distribution of the capital held. Decisions are made by considering possible courses of action and then choosing the one that leads to the best results (in terms of the utility function).

Summing up the notation:

- X_t is the state variable (a vector, in general) and E_t is the evidence variable, as usual. In the case of R-AGE, $X_t = (C_t, V_t, \boldsymbol{\mu})$ and E_t is the collection of reported claims, paid claims, information on volume sold, etc.
- A_t is the action at time t, e.g. setting the price P_t for the policies
- R_t is the reward at time t, e.g. the profits in year t
- U_t is the utility of the state at time t, i.e. (by definition) the sum of discounted future rewards.

The transition model for this network is $\Pr(X_{t+1}|X_t, A_t)$ ($T(s, a, s')$ in the language of POMDPs) and the observation model is $\Pr(E_t|X_t)$ ($O(s,o)$ in the language of POMDPs).

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

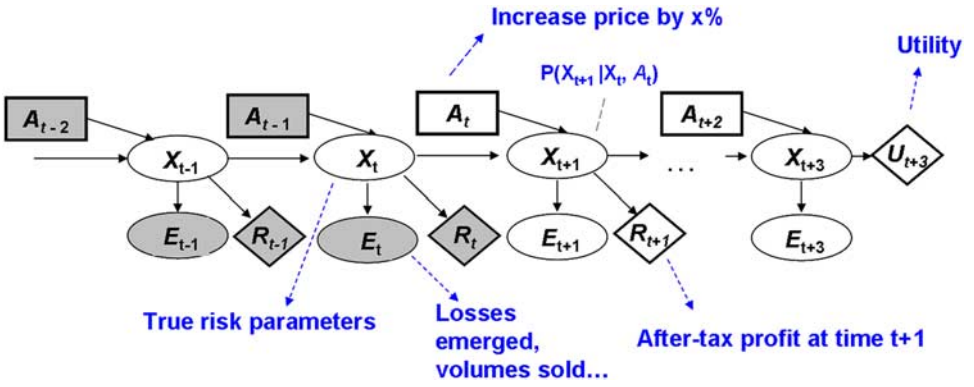


Figure 9. An insurance company pricing decision process seen as a dynamic decision network.

Figure 9 provides the generic structure for a DDN, and its interpretation for our R-AGE example. The network has been projected for three future periods to visualise the way in which a network like this is "solved": since the problem of finding the optimal strategy is intractable, the way one produces a (sub-optimal) strategy is by a look-ahead method:

1. Assume we are at time t
2. Set $k = 1$
3. Explore all possible actions at time $t + k$
4. For a given action, the updated *belief state* (not the actual state!) is uniquely determined by the monitoring algorithm explained above
5. Explore all possible "environment actions", i.e. observations
6. Calculate reward
7. Set $k \leftarrow k + 1$ and go back to Step 2, until desired look-ahead depth is achieved
8. When desired look-ahead depth is achieved, calculate the utility function in approximate form.

Note that we are using an approximate utility in the method above: if we knew the exact utility, we would not need a look-ahead algorithm – we could directly calculate U_t! Also note that looking ahead only a few steps is usually sufficient when the discount factor used is small enough, as the rewards beyond a certain level might be negligible.

Note that since this exhaustive search is possible only for a discrete number of states, the problem needs to be discretised in a sensible way before carrying out the computation (see Russell & Norvig (2003) for a discussion of computational complexity issues). Also, stochastic search over a random sample of actions and observations can sometimes replace exhaustive search, as in dynamic financial analysis exercises.

Overall, POMDPs and especially dynamic decision networks represent a good way of representing risk agents such as an insurance company, at least in situations that are not too complex. Various approximation methods – beyond those illustrated here – can be used to make near-optimal decisions with reasonable computational complexity. The interested reader is referred as usual to Russell & Norvig (2003) for a more thorough treatment of these techniques (see especially Chapter 21 on reinforcement learning).

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

The reader will have noticed analogies with ruin theory, dynamic financial analysis and pricing optimisation. As a matter of fact, all these theories/activities can be put in the context of risk agents making decisions in an uncertain environment:

- POMDPs are an extension of ruin theory in the sense that they also enable us to consider the effect of decisions and other uncertainties (including model and data uncertainty) on the probability of ruin in a Bayesian context;
- DFA is a special case of POMDPs where utility is calculated for a random sample of possible actions and observations – possibly in an interactive fashion;
- pricing optimisation can also be put into the context of POMDPs: in this case the demand modelling is given special emphasis.

4.5 Towards more realistic examples

The R-Age example outlined at the beginning of Section 4 was of course simplistic: an insurer trading in a single class of business, charging all customers the same amount, facing no direct competition, and knowing both the parameters of the loss distribution and the demand curve perfectly.

As we have seen in Section 4.3, POMDPs allow us to lift the assumption about knowing the loss distribution perfectly: the parameters of the loss distribution can be learned gradually as time progresses. If the distribution changes with the time, as it probably will, the changing parameters can still be learned if the changes are not too quick (see Section 3). Analogous considerations can be made about the demand curve, which needs to be learned as we go along.

Other assumptions, such as the fact that R-Age deals only with one product and charges all customers the same amount, have allowed us to simplify the problem but have not changed its nature, and were made only for clarity of treatment. Having multiple products, and different classes of customers, does not change the problem in principle: however, the resulting optimisation problem is far more complex, and the solution might be one in which the pricing of some or all the products is significantly sub-optimal: e.g. cross-subsidisation is a case in point.

A far more crucial assumption, however, was that R-Age was the only player in its niche market. When this is not the case, learning the characteristics of the environment as a whole – which includes not only the loss process and the customers but also the competitors and many other players – becomes an altogether different challenge, which will be addressed in the next section.

5 Modelling Collective Behaviour

We have thus far modelled the environment as an amorphous "blur". To a certain extent, this is inevitable, for example in the context of personal lines pricing, where millions of customers are involved and these cannot be dealt with individually.

However, in many circumstances the granularity of the environment cannot be ignored, and the strategy of a risk agent must take into account the interaction with other individual players: e.g. the pricing strategy of an insurance company must take into account what the other insurance companies do. This is obvious to risk practitioners, and it is also a theoretical necessity: something as basic as the underwriting cycle cannot be explained within the framework of a risk agent interacting with an amorphous environment.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

Modelling the interaction between individual risk agents is the most challenging aspect of developing an understanding of risk, and this section has only very tentative suggestions to make.

The most promising technique for modelling the behaviour of many interacting risk agents is, in the author's opinion, that of multi-agent systems (MAS) with the incorporation of game theory to model the interaction of the different agents. Section 5.1 briefly illustrates what MAS are; Section 5.2 illustrates the role of game theory; Section 5.3 describes a possible approach to multi-agent insurance problems.

For background reading, Russell & Norvig (2003) ("The Intelligent Agent Book", as the subtitle goes) has a full chapter (Chapter 17) devoted to the aspects of intelligent agents that are of more direct interest to us. A book fully devoted to multi-agent systems is Wooldridge (2001), which has a good section on game theory. However, this book is mainly aimed at software engineers and is less concerned with the quantitative aspects of intelligent agents.

5.1 Multi-agent systems (MAS)

Multi-agent systems are collections of intelligent agents, which interact both with the environment and among themselves. They can be used in a cooperative or non-cooperative fashion. For example, agents may cooperate to solve a given problem, by breaking it up into smaller tasks: this is a familiar use of agents in a software engineering context. Alternatively, they can be made to compete with each other. In this latter non-cooperative framework, which is the most interesting to us, there are two main ways in which we can use multi-agent systems:

- *Agent design*: designing agents that deal with other agents with an optimal strategy, so as to maximise their utility.
- *Mechanism design*: designing rules to regulate the interaction between different agents in order to maximise some overall measure of utility.

If we look at insurance companies as agents, agent design is what is required for an insurance company which wants to thrive among its competitors, or that of an individual wanting to buy a household policy. Mechanism design is instead the job of the regulator. When devising new regulations, it is important to understand what the effect of the regulation is going to be, and to test different types of regulations so as to reduce the instability of the system. This effect can only be understood at the collective, and not at the individual, level.

An example which follows this spirit although it ignores the formalism of multi-agent systems is the paper presented at GIRO 2008 by Taylor (Taylor, 2008), which investigates the effect of catastrophes and of regulation on producing cyclical effects such as the underwriting cycle.

5.2 Games of partial information

Game theory plays a role whenever there are agents that have to adopt successful interaction strategies, and can be used for both agent design and mechanism design. The games we are mostly interested in when modelling the behaviour of the insurance markets:

- are repeated games, in which agents meet one another repeatedly and keep memory of the previous encounters;
- involve several different agents making decisions simultaneously;

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

- concern information which is only partially observable, in the sense that each agent has only partial information on what its opponents are doing;
- have stochastic payoffs, at least because of randomness of the loss process;
- are mostly non-cooperative, as the insurance market is competitive.

Games in this type of setting are called *games of partial information*. Popular examples include card games such as poker and bridge, in which each player can only see a subset of the cards, or the modelling of nuclear wars, where the location of the opponents' weapons is unknown. Games of partial information can be solved theoretically by similar techniques to those used in Section 3.4, where a solution maximising utility was outlined by a look-ahead technique. A key difference is that one's belief state is observable whereas that of the opponents are not.

However, game theory has not been widely used for agent design. According to Russell & Norvig (2003), there are two main reasons for this:

- In a Nash equilibrium solution, a player is assuming that the opponents will play the equilibrium strategy, without being able to incorporate any prior belief that the player might have of what the opponent(s) might do. To address this problem, the notion of Bayes-Nash equilibrium has been introduced which explicitly incorporates our prior beliefs about the other players' likely strategies.
- No good way is currently known to combine game theory with POMDP theory in the framework of agent design.

For these reasons, game theory has been more successful in mechanism design, by which one attempts to maximise some global utility, basically attempting to discourage behaviours on the part of individual agents which might decrease the overall utility. This can be obtained by the imposition of rules that incorporate all externalities – that is, effects on global utility not recognised in the interaction between agents – explicitly. This has been successfully used for example in designing auctions (Wooldridge, 2001; Russell & Norvig, 2003). Another example, which is more directly related to general insurance, regards the optimal allocation of quotas of risk among insurers that evaluate risk differently (Fragnelli & Marina, 2003; Ambrosino *et al*., 2006).

In most cases concerning general insurance, an analytical solution of games involving risk agents will of course be impossible, given the large uncertainties and the large number of variables involved. The most promising application of game theory to multi-agent systems in insurance is perhaps through the stochastic modelling of agents, much along the lines of the work done by Axelrod (1984) on the evolution of cooperation. Axelrod devised a tournament among agents (represented as programming code) that would meet repeatedly at random and could choose at each encounter whether to "co-operate" or "defect". The payoff matrix was devised in such a way that mutual cooperation was preferred to mutual defection, whereas the best outcome for an agent would occur when the agent defected while the opponent cooperated. The idea was to have a repeated version of the famous prisoner's dilemma, in which the most rational decision is for both players to defect. The fact that the game is repeated creates a long-term incentive to cooperate (to avoid retaliation), which makes the game more realistic.

The key ideas of the Axelrod experiment are its stochastic nature (agents meet at random) and that rather than trying to design an optimal agent based on theoretical considerations (which proves to be extremely complicated even in simple settings such as this), one can design "tournaments" to test different strategies, and tentatively adopt the strategy which performs better in a number of tournaments.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

5.3 A possible approach to modelling the interaction among risk agents

In the light of the discussions in the previous sections, one possible approach to modelling risk agents is as follows (a more detailed example can be found in Parodi, 2009):

(i) Model each agent (e.g. an insurer, a customer, a reinsurer, etc) as a POMDP;

(ii) model the interaction between the risk agents as a game of partial information, where the state of one agent depends on the previous actions not only of the agent itself but of the other agents as well, and the risk agents have a (partial) knowledge of the other agents' strategies;

(iii) test the goodness of strategies by stochastic simulation, by having different strategies compete in an Axelrod-style experiment (not necessarily by one-to-one encounters, but by synchronised decisions).

This framework is in principle suitable to address both the agent design and the mechanism design problem. Note, however, that the balance between that information which is publicly available and that which is privately held by the risk agents may vary depending on the context, and especially between the agent design context and the mechanism design context. For example, the regulator might have a more in-depth view of the position of the different risk agents. Obviously, the optimal strategy will vary depending on the amount of information available.

One weakness of this approach is that by testing strategies in this way there is no assurance that optimality is achieved. Actually, the uncertainties involved in the process are so large that there are doubts that the solution is even sub-optimal. One must of course be aware that the level of complexity of even the most basic model is so high that there is a danger that we are able to reproduce well-known phenomena (underwriting cycle, etc) in a spurious fashion, simply because the number of parameters is sufficient to imitate those phenomena by appropriate tuning. At the same time, even the most complicated version of the model will pale in comparison with the messy reality of the insurance market and will not be realistic. Why bother, then? There are at least two reasons to attempt this investigation anyway:

- Although stochastic simulation of market behaviour will not reproduce the reality of the insurance market in an exhaustive way, it may bring to light unexpected collective phenomena;
- if we are able to reproduce known phenomena (bubbles, cycles, etc) with a very small number of parameters, and with a broad range of parameter values (that is, robustly), we may indeed have found something which has predictive value and that could be used for further investigation.

The overriding consideration should be to consider these models with much caution, as useful tools for exploratory analysis rather than rigorous means of understanding collective behaviour.

5.4 A practical example from financial modelling

An example which is very much in the spirit of the approach outlined in Section 5.3 can be found in the paper by Darley *et al*. (2001) "Learning, evolution and tick size effects in a simulation of the Nasdaq stock market". As the title suggests, this example comes from the field of financial modelling rather than from that of non-life insurance. However, I believe it is worth including it as it may easily provide a template for an analogous insurance exercise.

The authors provide an agent-based model of the stock market. The model is structured as follows:

i. Market makers and investors are modelled as individual agents.

ii. The market is made of a single security whose true value (or fundamental value) is assumed to fluctuate as the result of a stochastic process.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

iii. Investors do not know the true value, but a noisy version of it.

iv. Each agent behaves according to its individual strategy, e.g. you have investors who make their decision based on the true value they estimate, and those that act based on estimated trends. Market makers can also make use of information not available to the investors, e.g. all the trades they are involved in.

v. The individual strategies are either built in or "learned" based on market data – which allows for the discovery of new strategies.

Agents are then made to compete against one another in the market with a stochastic simulation. Based on this set-up, the authors have obtained a number of interesting results. We cite two of them, the first related to mechanism design and the second to agent design:

- reducing the tick size (i.e. the minimum amount by which stock prices can change) below a certain limit makes it more difficult for participants to discover the "true" value of the security (*price discovery*), if parasitic strategies such as day traders are allowed;
- under a variety of conditions, artificial learning strategies were seen to outperform built-in strategies based on expert knowledge or extracted from data.

As mentioned at the beginning of this section, this work may provide a template for similar works in an insurance context:

- the "price discovery" problem that Darley *et al*. (2001) addresses is not at all dissimilar from the problem of determining the pure cost of an insurance policy;
- as in the stock market, the market participants (insurers and insureds) never see the "true value" (fundamental value) of the policy, but only a noisy version of it, the noise coming from various sources of uncertainty: process, parameter, model and data uncertainty;
- whilst the stock market has constraints such as the tick size, the insurance market has its own constraints – such as capital requirements – on market participants. Stochastic simulation will help discover unintended consequences of these constraints;
- pricing strategies may be tested against one another, exactly as investing strategies in the stock market. Furthermore, it will be possible to compare "expert" strategies which attempt to replicate the behaviour of actual players in the market with artificial learning strategies.

6 Conclusions

The overall conclusion of this investigation is that computational intelligence provides a natural framework for addressing a large range of actuarial problems.

More specifically, these are the highlights of this investigation:

- The problem of risk costing can be seen in the context of statistical learning (see Part I).
- Uncertain and soft knowledge can be dealt with most successfully in a Bayesian context. Although other approaches (fuzzy set theory, non-monotonic logic) are possible and may lead to reasonable solutions, the Bayesian approach seems to the author to be the most suited to the highly numerical insurance context.
- The fact that risk changes over time and that the information on past risk also unfolds gradually can be addressed by a probabilistic temporal model (or state-space model), which allows the time evolution of the parameters and the structure of a model. Problems that can be naturally framed in this context include reserving and pricing in the face of a changing risk profile.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

- Insurers and other players in the insurance market can be represented as *intelligent agents* which move in an uncertain environment ("risk agents"). Risk agents can be modelled as partially observable Markov decision processes (POMDPs), as it is assumed that the state in which the agent is is not known for certain. This reflects the fact that insurance companies do not usually know their current financial position with certainty, especially if their portfolio has long-tail elements.
- The environment cannot always be considered a "blur" – it contains other agents whose behaviour cannot be ignored if we are to make successful decisions. Risk agents, in fact, interact with each other by competing, cooperating, exchanging goods and information, forming what is called a *multi-agent system*. Game theory then provides a framework for modelling the interaction between agents. The problem of optimising utility in a multi-agent context is very complex and any solution is at this stage purely speculative.

Overall, we have shown that most actuarial problems can be framed quite naturally in the computational intelligence context, and that by doing so we have a number of results from computational intelligence that we can tap into.

6.1 It's a Bayesian world

An overwhelming consideration is that most of the tools that appear to be helpful for solving general insurance problems are Bayesian in nature. Supervised learning (part I) can be interpreted in a Bayesian framework – we have the familiar trade-off between the prior knowledge (the penalty function) and the goodness of fit. Data and model uncertainty can be addressed by Bayesian techniques (see Equations (7) and (8)). The temporal element can be addressed by probabilistic temporal models, which in most cases are applications of the Bayes' rule (see for example Equations (11) and (12)). Making decisions in an uncertain environment can be addressed by partially observable Markov decision processes, which are an application of Bayes' theorem (see for example Equation (24), which gives the probability of transition between belief states). Multi-agent systems can also be viewed in a Bayesian context (as collections of POMDPs), and the interaction between different risk agents can be viewed as the quest for a Bayesian-Nash equilibrium. Although the power of the Bayesian paradigm will be no surprise to many actuaries, what was surprising to the author is the wealth of Bayesian applications beyond the standard tools of Bayesian analysis (the basic application of Bayes' theorem and the use of Bayesian networks).

The underlying reason for the ubiquity of the Bayesian approach is that the risk professional's job constantly entails integrating information coming from data with many other sources of information, including data from the market, from other clients, from one's previous experience and gut feel, from analogous cases – something which can be collectively called "prior knowledge". Even in the case of supervised learning, which seems to be exclusively data-driven, learning is guided by the prior belief that models should be as simple as possible, and that therefore complexity should be punished.

6.2 Limitations of computational intelligence as a framework for understanding risk in general insurance

The main limitation of computational intelligence techniques is that they themselves capture the "ecological" element only up to a point:

- There is a hidden assumption in many applications of machine learning that the environment is stable. For example, we make sure that the inferences we draw from our data are sound by putting aside a subset of these data and testing our theory (based on a training set) against this

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

subset (see Section 2.3). However, this may not tell us much of what will happen when we use our theory on data that have not been collected yet – e.g. losses that have yet to happen. This is a crucial difficulty for us in that risk is ever-changing. The probabilistic temporal model helps us deal with the problem of a changing risk, but in turn rely on the law by which the risk changes to be known or at least slowly varying – basically, we must have some reason to believe that the past will be at least of *some* help as a guide to the future!

- Changes in the environment are often introduced by people. Especially in finance, people "work the system" constantly and introduce changes that are difficult to model almost by definition. This again is a perfectly legitimate ecological aspect of risk, and one for which current techniques are inadequate. To solve this problem, we would need what is usually called "strong artificial intelligence" (that is, agents that actually exhibit general intelligence, much like humans). This is, despite some claims to the contrary, a highly speculative branch of computational intelligence, with little to show for in terms of practical achievements. And humans themselves are often not very good at moving in this type of environment.
- Despite the fact that multi-agent systems are a very promising paradigm for collective behaviour, a lot more research is needed before they can actually be used for taking optimal decisions at an individual and at the regulator's level. The sheer model complexity may prevent them *ever* to be used as optimal decision tools: however, they would still be useful in the production of unexpected scenarios.
- A final consideration involves the usual question as to the degree to which the techniques presented here (in both the first and the second paper) can replace experience and judgment. This question is asked every time a new technique is introduced which automates actuarial analysis and underwriting to some degree: but it is particularly meaningful here as we are speaking of techniques drawn from computational intelligence – the very discipline that used to have as its mission to equate or improve on human intelligence. The author's opinion on this is that these techniques cannot replace the actuary's (or the underwriter's) judgment more than blood tests and nuclear magnetic resonance (NMR) devices can replace the doctor's judgment; rather, they give more ammunition for making good decisions, as in the doctor's case. Furthermore, many of these statistical techniques are Bayesian in nature, and rather than replacing experience and judgment they provide an explicit way of *incorporating* that experience and judgment into the decision process, by making them part of the prior probability distribution. Any attempt to go beyond that and build some sort of artificial actuary/underwriter is almost certain to end up like many other "strong artificial intelligence" projects undertaken from the 1960s onwards – in an embarrassing failure.

Acknowledgements

I would like to thank the Profession for giving me the opportunity to carry out a piece of research that was in my mind since I have started working as an actuary, but never had the time to pursue. In particular I would like to thank Julian Lowe, my internal supervisor, for his time and guidance and especially for the kind encouragement throughout the project. I would also like to thank Mark Flower for his painstaking revision of an earlier draft.

I am indebted to Alessandro Verri for giving me the opportunity of spending several weeks at the University of Genoa, and for many mind-opening conversations. My thanks extend to his team, for talking me through their projects and for their precious feedback on this work.

For this second part of the paper, I am especially thankful to Stephen Roberts for advice on how to incorporate the uncertainty on the transition state matrix in a Markov decision process, to Vito Fragnelli for discussion on the use of game theory in insurance, and to Viviana Mascardi for introducing me to multi-agent systems techniques.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

I would also like to thank the following for helpful discussion: Jill Boulton, James Coyle, Phil Ellis, Clinton Freeman, Chris Gingell, Di Kuang, Eamonn McMurrough, Chris Russell, Andy Smyth, Jen Tan Ning, Claire Wilkinson.

References

Ambrosino, D., Fragnelli, V. & Marina, M. (2006). Resolving an insurance allocation problem: A procedural approach. *Social Choice and Welfare*, **26**(3), 625–643.

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Bender, E.A. (1996). *Mathematical Methods in Artificial Intelligence*. IEEE Computer Society.

Bishop, C.M. (2007). *Pattern Recognition and Machine Learning*. Information Science and Statistics.

Cummins, J.D. & Derrig, R.A. (1997). Fuzzy financial pricing of property-liability insurance. *North American Actuarial Journal*, **1**(4), 21–44.

Darley, V., Outkin, A., Plate, T. & Gao, F. (2001). Learning, Evolution and Tick Size Effects in a Simulation of the Nasdaq Stock Market. *Proceedings of the Joint 5th World Multiconference on Systemics, Cybernetics and Informatics (SCI 2001) and the 7th International Conference on Information Systems Analysis and Synthesis (ISAS 2001)*.

De Jong, P. & Zehnwirth, B. (1983). Claims reserving, state-space models and the kalman filter. *Journal of the Institute of Actuaries*, **110**, 157–181.

de Silva, C.J.S. & Attikiouzel, Y. (1994). A better path to duplicating human reasoning. *IEEE Expert: Intelligent Systems and Their Applications*, **9**(4), 14.

Elkan, C. (1994). The paradoxical success of fuzzy logic. *IEEE Expert: Intelligent Systems and Their Applications*, **9**(4), 3–8.

Fragnelli, V. & Marina, M.E. (2003). A fair procedure in insurance. *Insurance: Mathematics and Economics*, **33**(1), 75–85.

Ghahramani, Z. & Hinton, G.E. (1998). Variational learning for switching state-space models. *Neural Computation*, **12**(4), 963–996.

Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, **57**, 357–384.

Lemaire, J. (1990). Fuzzy insurance. *ASTIN Bulletin*, **20**(1), 33–55.

MacKay, D.J.C. (2003). *Information Theory, Inference, and Learning Algorithms (4th Printing)*. Cambridge University Press.

Nikolaidis, E., Cudney, H.H. & Chen, Q. (1999). Bayesian and fuzzy set methods for design under uncertainty. *Virginia Polytechnic Institute and State University*. 13th ASCE Engineering Mechanics Division Conference. Baltimore, Maryland.

Papadimitriou, C. (1994). *Computational complexity*. Addison-Wesley, Reading, Massachusetts.

Parodi, P. (2009). Computational intelligence techniques for general insurance. Research dissertation for the Actuarial Profession. Available at http://www.actuaries.org.uk/research-and-resources/documents/computational-intelligence-techniques-general-insurance

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems : networks of plausible inference*. Morgan Kaufmann Publishers, San Mateo, Calif., 1988.

Popper, K. (1959). The propensity interpretation of probability. *The British Journal for the Philosophy of Science*, **X**(37), 25–42.

Roberts, C.P. & Casella, G. (2004). *Monte Carlo Statistical Methods* (*2nd Edition*). Springer-Verlag, New York.

Russell, S.J. & Norvig, (2003). *Artificial Intelligence: A Modern Approach* (*Second Edition*). Prentice Hall.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press

Ribeiro, I. (2004). Kalman and Extended Kalman Filters: Concept, Derivation and Properties. *Institute for Systems and Robotics, IST*, (http://www.isr.ist.utl.pt/ ~ mir)

Roberts, S. (2011). (http://www.robots.ox.ac.uk/ ~ sjrob). Personal communication.

Semagul, A., Altindas, E., Macit, R., Umar, S. & Unal, H. (2009). fuzzyOP: Fuzzy numbers and the main mathematical operations. R Package (http://cran.r-project.org/web/packages/fuzzyOP/index.html)

Shapiro, A.F. (2004). Fuzzy logic in insurance. *Insurance: Mathematics and Economics*, **35**(2), 399–424.

Taylor, G. (2008). A simple model of insurance market dynamics. *GIRO proceedings*. Available at http://www.actuaries.org.uk/research-and-resources/documents/simple-model-insurance-market-dynamics-slides

Wooldridge, M.J. (2001). *Multi-agent systems : an introduction*. Wiley, Chichester GBA1-Z6596 Michael Wooldridge.

Zadeh, L.A. (1965). Fuzzy sets. *Information and Control*, **8**, 338–353.

Zarchan, P. & Musoff, H. (2005). *Fundamentals of Kalman Filtering: A Practical Approach* (2nd ed). American Institute of Aeronautics and Astronautics.

https://doi.org/10.1017/S1748499512000048 Published online by Cambridge University Press