Dynamic Uncertainty Handling for Coherent Decision Making in Nuclear Emergency Response

Manuele Leonelli and James Q. Smith

Department of Statistics, The University of Warwick, Coventry, UK, CV47AL, m.leonelli@warwick.ac.uk

INTRODUCTION

Decision making in nuclear emergency management takes place in an uncertain dynamically changing environment, where numerous interdependent events and processes need to be taken into account. Decision makers (DMs) need also to incorporate the opinions of several relevant stakeholders regarding competing objectives. In order to support DMs, decision support systems (DSSs) have been developed to describe the overall problem under study and provide guidelines and explanations about the available countermeasures. One of such systems for nuclear emergency management is RODOS (see e.g. Bartzis et al. 2000), which was the result of the collaboration of several European countries after the Chernobyl accident.

A typical DSS for nuclear emergency response, such as RODOS, often consists of several component DSSs, or *modules*, which using either deterministic or probabilistic methodologies, guide the forecasting of the relevant quantities associated to each of these objectives. The earliest works on modules for nuclear emergency management concerned the initial stages of the accident. For example, French et al. (1995) described the likely development of the profile of the source term by developing a probabilistic belief network. In Thykier-Nielsen et al. (1999) a puff dispersion model was developed to represent the spread of the contamination from the source term. In more recent years probabilistic models have been created to describe also later stages of the emergency: De and Faria (2011) and Richter et al. (2002) introduced further probabilistic models for the deposition and the food chain modules, respectively.

However each of these component DSSs addresses only a subset of the overall problem. In addition each of these is usually separately informed by an appropriate panel of experts, who together will deliver the required judgments to evaluate the efficacy of potential countermeasures. For a DM to be properly supported in nuclear emergency management, these separate judgments need to be incorporated into a unique and balanced description of the problem as a whole. In a Bayesian framework this corresponds to the construction of a single Bayesian probability model, informed by the judgments of separate panels of experts.

So how can a system coherently integrating the different modules into a single entity, identifying and managing the relationships existing between them, be built? One of the conditions that can ensure a coherent description of the problem is the appropriate propagation of the involved uncertainties through the input/output relationships between the modules. For example, a typical source term will provide forecasts of the amount of contaminants that a chimney will spread into the environment. A subsequent dispersion module will use such forecasts to study the future dispersion of radiation in air, which will depend on the level of contamination from that chimney. It is well known that without considering the levels of uncertainty associated to such forecasts, DMs might be lead to choose a dramatically inappropriate sequence of acts (see e.g. Leonelli and Smith 2013a). A second requirement for a coherent aggregation of the panels' judgments is the choice of an overall utility function within a particular family, capturing some form of preferential independence and represented by some factorization of a DM's utility function.

When in addition it is sufficient for the panels of experts, after receiving appropriate information from other panels, to individually deliver their own judgments about the module they oversee, we will call the system *distributed* or *modular*. There are many advantages associated to the use of a distributed system. First, in nuclear emergency management, experts are usually experts about a subset of the problem, as the physical functionality of the plant, the dispersal of the contamination, the medical implications of the intake of radiations, and so on. Modularity ensures that each panel of experts within a specific domain needs only to focus on their area of jurisdiction delivering beliefs about their area of expertise. It is well known that experts' judgments are better calibrated if they have received feedback on similar judgments in the past (Cooke 1991). Second, a potential DM is credible when explaining the rationale behind her actions since, because of the modularity of the system, she is using the best possible information from informed experts. Third, if the system needs to be changed in the light of unexpected developments or implications - and this often happens when facing a crisis - under suitable conditions the management of these new developments can be addressed locally by the relevant expert panel(s). They simply adapt their individual forecasts and the inputs in the light of this new scenario and these can then be folded into the system. Finally, a modular system is fast since it allows the collective to quickly identify its optimal policies as functions of statistics provided by the individual component panels and each panel can locally and independently update its beliefs in the light of new evidences.

Of course for many domains such consistency is not critical and it is admissible and sometimes most valued for the system to express a wide range of views and policies that are not necessarily consistent with each other. However in many domains it is, including nuclear emergency management, as we extensively discussed in Leonelli and Smith (2013a).

Although it has been stated that it is possible to develop a Bayesian methodology to deal with these kinds of problems (see French 1997), it has so far been considered infeasible to actually implement it. However, due to the recent computational and theoretical advances, it is now possible to build coherent distributed Bayesian DSSs in this multi-expert setting. In Leonelli and Smith (2013) we develop a structure that integrates these component modules into a unique Bayesian model. The structure can then form the methodological framework around which an **integrating decision support system** (IDSS) will be used to inform its users. In that paper we introduced a set of assumptions that enables the construction of a unique Bayesian model and developed message passing algorithms, both in the non-dynamic and dynamic case, to compute (optimal) expected utility scores, which use only the delivered panels' judgments and do not break the modularity of the system. These scores updating in real time will then be available to inform the DMs about the relative efficacy of any measure they might entertain.

In this paper for simplicity we will focus our attention on a particular graphical dynamic Bayesian model, the **multi-regression dynamic model** (MDM) (Queen and Smith 1993) and show how this can be applied in practice in a simple situation. We chose a dynamic model here because in crisis management it is critical to have the flexibility to allow for algorithmic updates of the relevant probabilities. This is because throughout the crisis new information is continuously gathered and DMs need to make decisions each time new information becomes available. The MDM is a particularly convenient dynamic version of a Bayesian Network (BN). It is defined on the inputs and outputs of the various component probabilistic models in the system. This has the extremely useful property that the propagation of probabilistic uncertainties through the input/output relationships can be performed in a modular way and in closed form across the different components of the system.

The paper is structured as follows. The next section introduces the MDM and reviews its main properties. We then briefly describe the properties that allow the construction of an IDSS and the algorithm for the identification of optimal expected utility scores within the MDM framework. We proceed to illustrate the operations of the IDSS using the MDM recurrences to demonstrate a simple practical implementation of the methodology. We conclude with a discussion.

THE MULTI-REGRESSION DYNAMIC MODEL

Before formally introducing the statistical model we will use in this paper, we need to set up some notation and introduce some terminology concerning graphs. Let $\{\boldsymbol{Y}_t\}_{t=1,\dots,T}$ be an N-dimensional time series and suppose that $\{\boldsymbol{Y}_t\}$ is partitioned into n vector time series of dimension $r_1,\dots,r_n$ with $\sum_{i=1}^{n} r_i = N$, so that $\boldsymbol{Y}_t^{\mathrm{T}} = (\boldsymbol{Y}_t(1)^{\mathrm{T}},\dots,\boldsymbol{Y}_t(n)^{\mathrm{T}})$. Let $\boldsymbol{Y}^t = (\boldsymbol{Y}_1,\dots,\boldsymbol{Y}_t)^{\mathrm{T}}$, $\boldsymbol{Y}^t(i) = (\boldsymbol{Y}_1(i),\dots,\boldsymbol{Y}_t(i))^{\mathrm{T}}$ and $\boldsymbol{y}_t, \boldsymbol{y}_t(i)$ be the realizations of $\boldsymbol{Y}_t$ and $\boldsymbol{Y}_t(i)$ respectively.

A directed graph $\mathcal{G} = (V(\mathcal{G}), E(\mathcal{G}))$ is a pair of sets satisfying $E(\mathcal{G}) \subset V(\mathcal{G})^2$, such that the elements of $E(\mathcal{G})$ consists of ordered pairs of elements of $V(\mathcal{G})$. The elements of $V(\mathcal{G})$ are called *vertices*, while the elements of $E(\mathcal{G})$ are called *edges*. Figure 1 shows a directed graph with vertices $\{\boldsymbol{Y}_t(1),\dots,\boldsymbol{Y}_t(4)\}$, in which the edges are depicted by arrows connecting two vertices. A vertex $\boldsymbol{Y}_t(i) \in V(\mathcal{G})$ is a *parent* of a vertex $\boldsymbol{Y}_t(j) \in V(\mathcal{G})$, and $\boldsymbol{Y}_t(j)$ is a *child* of $\boldsymbol{Y}_t(i)$ if $\{\boldsymbol{Y}_t(i), \boldsymbol{Y}_t(j)\} \in E(\mathcal{G})$. A vertex with no children is called *leaf*, while a vertex with no parents is called *root*. The set of all parents of a vertex $\boldsymbol{Y}_t(i)$ is the *parent set*, denoted with $pa(\boldsymbol{Y}_t(i))$. We further define the *family* of a vertex $\boldsymbol{Y}_t(i)$ to be the set consisting of $\boldsymbol{Y}_t(i)$ and all its parents, denoted as $Fa(\boldsymbol{Y}_t(i))$. Further recall that a clique of a graph is a maximal subset of its vertices such that every two vertices of the subset are connected by an edge. For example consider the graph in Figure 1, which has two cliques, $C_1 = \{\boldsymbol{Y}_t(1), \boldsymbol{Y}_t(4)\}$ and C_2 which includes vertices $\{\boldsymbol{Y}_t(1), \boldsymbol{Y}_t(2), \boldsymbol{Y}_t(3)\}$. A *decomposable* graph is one where all the parents of a given vertex are joined by an edge. In a decomposable graph there is a unique vertex $\boldsymbol{Y}_t(i)$ for each vertex $\boldsymbol{Y}_t(j)$ for which $pa(\boldsymbol{Y}_t(j)) \subseteq \{pa(\boldsymbol{Y}_t(i)), \boldsymbol{Y}_t(i)\}$. We call this vertex $\boldsymbol{Y}_t(i)$ the *father* of $\boldsymbol{Y}_t(j)$, while $\boldsymbol{Y}_t(j)$ is the *son* of $\boldsymbol{Y}_t(i)$. Note that a vertex can have more than one son, while it cannot have more than one father. Referring to the network in Figure 1, we can deduce, for example, that the father of $\boldsymbol{Y}_t(3)$ is $\boldsymbol{Y}_t(2)$ and that $\boldsymbol{Y}_t(1)$ has two sons, $\boldsymbol{Y}_t(2)$ and $\boldsymbol{Y}_t(4)$. Recall that a directed path is a sequence of edges for which the child vertex associated to that edge (excluding the last edge of the sequence) coincides with the parent vertex of the subsequent one. A directed cycle is a path in which the parent vertex associated to the first edge of the

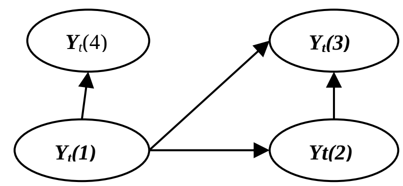

Fig. 1. Example of a DAG.

sequence and the child associated to the last edge coincide. A *directed acyclic graph* (DAG) is a directed graph containing no directed cycles. Note that the graph in Figure 1 is a DAG. A graph in which for each pair of vertices there exists a directed path which joins them is called *connected.*

Definition of the model

The MDM is a multivariate graphical Bayesian model, which uses conditional independence to break a complex multivariate time series model into simpler regression dynamic linear models (DLMs) (West and Harrison 1997) of smaller dimension. Each time slice of the MDM is described by a DAG. In this DAG each child has as parents those contemporaneous variables on which it is regressed. Note that the graphical representation that depicts a BN is also a DAG in which the vertices coincide with random variables/vectors and edges represent statistical dependence.

The state space parameter associated to the time slice DAGs at time t is $\boldsymbol{\theta}_t^{\mathrm{T}} = (\boldsymbol{\theta}_t(1)^{\mathrm{T}}, \dots, \boldsymbol{\theta}_t(n)^{\mathrm{T}})$, where $\boldsymbol{\theta}_t(i)$ is the parameter vector associated to $\boldsymbol{Y}_t(i)$. In the MDM these time slice DAGs exhibit the same topology at each time point. The individual DAGs are then connected to each other, through an algorithmic recursion over the overall parameter vector. In particular the dynamic processes over the observed component series are expressed indirectly through a dynamic expression on its dynamic parameters. Note that the BN is a special case of the MDM in which the parameter vector does not change through time. In the literature the MDM has been applied to several different domains, as for example, brand markets (Queen 1994), vehicles' traffic control (Anacleto et al. 2013) and brain connectivity (Costa et al. 2013) where it has outperformed other available competing models.

The MDM is defined through n observation equations, a system equation and a prior distribution. Specifically,

- Observation equations, $i = 1, \dots n$:

$$\boldsymbol{Y}_t(i) = F_t(i)^{\mathrm{T}}\boldsymbol{\theta}_t(i) + \boldsymbol{v}_t(i), \quad \boldsymbol{v}_t(i) \sim \big(0, \Sigma_t(i)\big) \qquad (1)$$

- System equation:

$$\boldsymbol{\theta}_t = G_t\boldsymbol{\theta}_{t-1} + \boldsymbol{w}_t, \qquad \boldsymbol{w}_t \sim (0, W_t) \qquad (2)$$

- Initial information:

$$\boldsymbol{\theta}_0 | I_0 \sim (\boldsymbol{m}_0, C_0) \qquad (3)$$

Here $F_t(i)^{\mathrm{T}}$ is a $r_i \times s_i$ matrix of regressors, known functions of $pa\big(\boldsymbol{y}^t(i)\big) = \big\{pa\big(\boldsymbol{y}_1(i)\big), \dots, pa\big(\boldsymbol{y}_t(i)\big)\big\}$ and $\boldsymbol{y}^{t-1}(i)$, but not $\boldsymbol{y}^t(i+1), \dots, \boldsymbol{y}^t(n)$ or $\boldsymbol{y}_t(i)$. The s-dimensional vector $\boldsymbol{\theta}_t^{\mathrm{T}} = (\boldsymbol{\theta}_t^{\mathrm{T}}(1), \dots, \boldsymbol{\theta}_t^{\mathrm{T}}(n))$ has elements $\boldsymbol{\theta}_t^{\mathrm{T}}(i)$ of dimension s_i, such that $\sum_{i=1}^n s_i = s$. The observation errors $\boldsymbol{v}_t(i)$ have dimension r_i and can be either assumed known or unknown. The $r_i \times r_i$ matrix $\Sigma_t(i)$ is the observational covariance matrix for $\boldsymbol{Y}_t(i)$. $G_t = \mathrm{blockdiag}\big(G_t(1), \dots, G_t(n)\big)$ and $W_t = \mathrm{blockdiag}\big(W_t(1), \dots, W_t(n)\big)$ are $s \times s$ matrices, where $G_t(i)$ and $W_t(i)$ are, respectively, the $s_i \times s_i$ state evolution matrix and state evolution covariance matrix for $\boldsymbol{\theta}_t(i)$, $i = 1, \dots, n$, which may be functions of $pa(\boldsymbol{y}^{t-1}(i))$ and $\boldsymbol{y}^{t-1}(i)$, but not $\boldsymbol{y}^t(i+1), \dots, \boldsymbol{y}^t(n)$ or $\boldsymbol{y}_t(i)$. The s-dimensional vector $\boldsymbol{w}_t^{\mathrm{T}} = (\boldsymbol{w}_t(1)^{\mathrm{T}}, \dots, \boldsymbol{w}_t(n)^{\mathrm{T}})$ has elements $\boldsymbol{w}_t(i)$, $i = 1, \dots, n$, of dimension s_i called system error vectors. The s-dimensional vector $\boldsymbol{m}_0$ and the $s \times s$ matrix $C_0 = blockdiag(C_0(1), \dots, C_0(n))$ are the first two moments of the distribution of $\boldsymbol{\theta}_0 | I_0$, where I_0 represents the initial information available. The errors $\boldsymbol{v}_t(1), \dots, \boldsymbol{v}_t(n)$ and $\boldsymbol{w}_t(1), \dots, \boldsymbol{w}_t(n)$ are independent of each other and through time.

So for example, assume that the time slice DAG of an MDM is described by the graph in Figure 1. Then for example $F_t(3)$ can be a function of $\boldsymbol{y}^t(1), \boldsymbol{y}^t(2)$ and $\boldsymbol{y}^{t-1}(3)$, but not of $\boldsymbol{y}^t(4)$. Conversely, $F_t(4)$ can be a function of $\boldsymbol{y}^t(1)$ and $\boldsymbol{y}^{t-1}(4)$, but not of $\boldsymbol{y}^t(2)$ and $\boldsymbol{y}^t(3)$.

The main property of MDMs, which is particularly useful in this domain, is that, unlike rival models such as the 2 time slice BN (Murphy, 2002), if the parameter set is a priori independent, which is guaranteed by the blogk-diagonality of the matrix C_0, then it remains independent at each time point given the current information. This result is a dynamic version of the so-called prior to posterior global independence property for static BNs (Cowell et al. 1999), which considers the parameter vectors associated to different vertices to be independent of each other. This natural independence property, together with assumptions that allow the valid construction of an IDSS, will guarantee that it is possible to factor the process so that it can be updated in a modular fashion by the different individual panels. The other main feature of MDMs is the constant topology of the time slice DAGs, which, by construction, never changes through time. This provides us with an enduring framework and will allow the same algorithm for the computation of expected utility scores to continue to be valid as an incident unfolds and new information is added to the system.

Because of the independence of the parameter vectors, when the system is idle, that is when no decisions are/need to be taken, the joint forecast density can be written in product form as

$$f(\boldsymbol{y}_t \mid \boldsymbol{y}^{t-1}) = \prod_{i=1}^n g\left(pa\big(\boldsymbol{y}^t(i)\big), \boldsymbol{y}^{t-1}(i), \boldsymbol{\theta}_t(i)\right), \qquad (4)$$

where

$$g(pa\big(\boldsymbol{y}^t(i)\big), \boldsymbol{y}^{t-1}(i), \boldsymbol{\theta}_t(i)) = \int_{\boldsymbol{\Theta}_t(i)} f\left(\boldsymbol{y}_t(i) \,\middle|\, pa\big(\boldsymbol{y}_t(i)\big), \boldsymbol{y}^{t-1}(i), \boldsymbol{\theta}_t(i)\right) f\big(\boldsymbol{\theta}_t(i) \mid pa\big(\boldsymbol{y}^{t-1}(i)\big), \boldsymbol{y}^{t-1}(i)\big) d\,\boldsymbol{\theta}_t(i), \qquad (5)$$

and $\boldsymbol{\Theta}_t(i)$ is the state space of $\boldsymbol{\theta}_t(i)$.

When Gaussian errors and conjugate priors are used, the MDM as defined through Equations (1)-(3) enjoys closed form recurrences. This is because a conjugate analysis is possible. Although each component conditional process is a standard distribution, the overall composite distribution is nevertheless often highly non-standard (see Queen and Smith 1993). However its predictive densities can be derived in closed form and follow multivariate T-densities, but many of whose parameters are non-linear functions of parent observations.

Of course in practice some of the modules might not exhibit this level of complexity. In practice it is often assumed that the observational covariances are constant through time and that the matrix of regressors is a linear function of its arguments. In this case the model has been called linear multi-regression dynamic model (LMDM) (Queen and Smith 1993). Such model might even degenerate into the standard BN, which can be expressed as a simple non-stochastic special case of an MDM.

In the following section we will show how the MDM can be used as an overarching structure for an IDSS and then introduce a specific version of the assumptions that we need in this case to allow the construction of an IDSS for MDMs. For a more general treatment of these see Leonelli and Smith (2013). We will then show how expected utility scores for different policies can be computed as a result of the probability factorization in Equations (4)-(5).

INTEGRATING DECISION SUPPORT SYSTEMS

The definition of an IDSS needs three main classes of input which need to be common knowledge between all panels and potential users. These are the specification of the decision problem, the agreed overarching probabilistic structure and the agreed utility factorization. In this section we will next outline these necessary features in more details for each of these aspects of the problem. Let $\boldsymbol{D} = (\boldsymbol{D_0}, \dots, \boldsymbol{D}_{T-1})$ be a vector of possibly multivariate decisions and $\boldsymbol{d} = (\boldsymbol{d}_0, \dots, \boldsymbol{d}_{T-1})$ be a particular instantiation of this vector. Write $\boldsymbol{D}^t = (\boldsymbol{D}_0, \dots, \boldsymbol{D}_t)$ and $\boldsymbol{d}^t = (\boldsymbol{d}_0, \dots, \boldsymbol{d}_t)$. Let UT denote the index set of the attributes of the problem, where UT $\subseteq \{1, \dots, n\}$. We will henceforth call the group consisting of representatives of each panel of experts, potential DMs, relevant stakeholders, *the collective*. This group of people will be responsible for the global structure of the IDSS, while the panels will have local jurisdiction over the module they oversee. The common knowledge will be assumed to hold across this collective. Since the collective must act coherently as a single entity, it will behave like a single decision maker. Henceforth for the purposes of this paper we will refer to this virtual decision maker as the *SupraBayesian* (SB).

The agreed decision space

We begin by briefly introducing and discussing certain structural assumptions which, if made, allow us to define a mathematical framework around which modular decision making can be defined. The first structural assumption we introduce describes the nature of the decision problem the collective is planning to address.

Assumption 1. (Structural consensus)
The collective agrees to observe the values of the series $\{\boldsymbol{Y}_t\}_{t=1,\dots,T}$ or commit to a decision $\boldsymbol{D}_i$, $i = 0, \dots, T-1$, according to the order of the following time sequence:

$$(\boldsymbol{D}_0, \boldsymbol{Y}_1, \boldsymbol{D}_1, \dots, \boldsymbol{D}_{T-1}, \boldsymbol{Y}_T).$$

In addition we will require that

$$f(\boldsymbol{y}_t \mid \boldsymbol{d}, \boldsymbol{y}^{t-1}) = f(\boldsymbol{y}_t \mid \boldsymbol{d}^{t-1}, \boldsymbol{y}^{t-1}),$$

for $t = 1, \dots, T$.

Note that the first part of this assumption demands a certain consistency of time ordering but otherwise simply defines our notation. The second simply means that the distribution of a random variable can only be affected by a decision if that decision is taken before it. In the medical literature, such decisional structure has been called a *dynamic treatment strategy* (Murphy 2003, Dawid and Didelez 2010) and is the usual starting point in this domain. We will also use it for our purposes here.

The probabilistic overarching structure

The second agreement the collective needs to find concerns the qualitative form of the overarching probability model for the time series $\{\boldsymbol{Y}_t\}_{t=1,\dots,T}$. There are many possible choices for this structure. As noted before, in this paper we will assume that such model is the MDM.

Assumption 2. (Probabilistic consensus)
The collective agrees to model $\{\boldsymbol{Y}_t\}_{t=1,\dots,T}$ using a MDM in which the time slice DAG describing the MDM is connected, decomposable and is valid whatever decision $\boldsymbol{d} \in \boldsymbol{D}$ is taken.

It is worth pointing out at this stage that, although we assume that the probabilistic model is an MDM, there are several other dynamic models that can be used as a qualitative overarching structure, as showed in Leonelli and Smith (2013). Indeed, the methodology works in a more general setting. The MDM simply represents an important special case of the permitted form of the overarching model. In this paper we later use this form to demonstrate some of the explicit recurrences that can be derived for the IDSS methodology.

We envisage here that the collective will be facilitated into discovering a common structure during Decision Conferences, where a facilitator guides the collective's discussion and exploration of the implications of the possible models (French et al. 2009). During these meetings they will also discuss any other necessary qualitative probabilistic assumption that might need to be made, or further refine the model specification. These often include the above mentioned local and global independence of the parameter set. We note that these assumptions are well documented and are standard ones made in most applications of graphical model technologies. These consequences can also be expressed and explored in common language.

In practice, Assumption 2 states that the collective need to agree on whether or not one time series affects another. This then will define the qualitative input/output relationships of series of observations in the IDSS.

The utility overarching structure

Just as for the probability structure, the collective needs to find an agreement about the utility factorization which depicts the preferential independences they are ready to assume. In order for the algorithm we give below to work this utility factorization needs to lie within some specified family. In Leonelli and Smith (2013) we formally define this family of factorizations, which we called *dynamic compatible*. For the purposes of this paper, we will only present the intuition behind this factorization and an application to the DAG in Figure 1.

Denote with $\boldsymbol{Y}_{\mathrm{UT}}^{T}$ the vector containing all the vectors up to time T whose index lie in UT. Let also $\boldsymbol{Y}_{\mathrm{t,UT}}$ be a subvector of $\boldsymbol{Y}_{\mathrm{UT}}^{T}$, which includes only those vectors with time index equal to t, and $\boldsymbol{y}_{\mathrm{t,UT}}$ one of its instantiation. We assume here that a set of assumptions implying the existence of a utility function u over $\boldsymbol{Y}_{\mathrm{UT}}^{T}$ is appropriate (see for example French and Insua 2000). Recall that a multilinear utility factorization (Keeney 1993) for $\boldsymbol{Y}_{\mathrm{t,UT}}$ can be written as

$$u(\boldsymbol{y}_{\mathrm{t,UT}}) = \sum_{I\in\mathcal{P}(\mathrm{UT})} k_I \prod_{i\in I} u(\boldsymbol{y}_t(i)),$$

where $\mathcal{P}$ represents the power set and k_I is a criterion weight (see Keeney 1993). This factorization is appropriate when each attribute is *utility independent* of the others. Now a dynamic compatible factorization is a simplification of a multilinear one within each time slice, in which the interactions between terms are dropped if they belong to different cliques in the time slice DAG. In addition it consists of a linear combination of these factorizations through time. Thus a simple value independent utility factorization is always in the family of compatible utilities whatever the graph. However the compatible class is a much more expressive one and gives more flexibility in making a problem to an elicited preference structure.

To illustrate this new class of utility factorizations, consider the graph in Figure 1. Assume also that UT$=\{2,3,4\}$. Then the most general compatible utility factorization can be written as

$$u\big(\boldsymbol{y}^T(2),\boldsymbol{y}^T(3),\boldsymbol{y}^T(4)\big) = \sum_{t=1}^{\mathrm{T}} k_{t,2}u\big(\boldsymbol{y}_t(2)\big) + k_{t,3}u\big(\boldsymbol{y}_t(3)\big) + k_{t,4}u\big(\boldsymbol{y}_t(4)\big) + k_{t,2,3}u\big(\boldsymbol{y}_t(2)\big)u\big(\boldsymbol{y}_t(3)\big)$$

Note that the above factorization is a particular subclass of the multilinear one within each time slice, in which all the interaction terms between time series with index in $\{2,3\}$ and $\{4\}$ are dropped. Compatible models in general only allow sets whose elements belong to the same family defined by the decomposable DAG of the collectively agreed MDM. Note that these families of graphically indexed utilities are rather different from others suggested in the literature (Gonzales and Perny 2005, Abbas 2010) and are designed for our particular distributive coherence property we develop here.

We are now ready to make the following assumption.

<u>Assumption 3. (Preferential consensus)</u>
The collective is able to identify an agreed dynamic compatible utility decomposition over $\mathbf{Y}_{\mathrm{UT}}^{T}$ and to elicit the associated common criterion weights.

Again, as for the probabilistic counterpart these issues can be discussed during joint decision conferences across the collective. Specifically this is where we might expect the values of the criterion weights k_I within the compatible utility function to be decided.

Thus a dynamic compatible utility factorization can be written as a linear combination of terms, whose arguments belong to the same family in the time slice DAG. Specifically, because of the uniqueness of the father in the decomposable MDM jointly defined by the collective, the above factorization can be written as

$$u\big(\boldsymbol{y}^T(2),\boldsymbol{y}^T(3),\boldsymbol{y}^T(4)\big) = \sum_{t=1}^{\mathrm{T}} u_{t,2}(Fa(\boldsymbol{y}_t(2))) + u_{t,3}(Fa(\boldsymbol{y}_t(3))) + u_{t,4}(Fa(\boldsymbol{y}_t(4)))\ ,$$

where

$$u_{t,3}\Big(Fa\big(\boldsymbol{y}_t(3)\big)\Big) = \quad k_{t,3}u\big(\boldsymbol{y}_t(3)\big) + k_{t,2,3}u\big(\boldsymbol{y}_t(2)\big)u\big(\boldsymbol{y}_t(3)\big),$$

and $\quad u_{t,i}\Big(Fa\big(\boldsymbol{y}_t(i)\big)\Big) = k_{t,i}u\big(\boldsymbol{y}_t(i)\big), \quad$ for $\quad i = 2,4.$ Specifically each term $u_{t,i}(\cdot)$ includes $k_{t,i}u\big(\boldsymbol{y}_t(i)\big)$ and any interaction term between $u\big(\boldsymbol{y}_t(i)\big)$ and utilities over attributes with a lower index within the same time slice. These terms will be central in the algorithm for the computation of the expected utility scores. Henceforth we call them *t/i-th utility functions*.

The local panels' structure

The previous assumptions specified the overarching structure of the IDSS. Under this construction, the panels of experts need to agree on their own local specifications. Each of these needs to be consistent with the overarching structure. Specifically each panel needs to provide a unique, possibly conditional, probability distribution and a marginal utility function. These two requirements are specified in the following assumption.

Assumption 4. (Local panels' consensus)
Every expert within a panel G_i agrees on a probabilistic model for the associated parameters and sample distribution of their particular component DSS. This will enable the collective to calculate and potentially deliver the output $\boldsymbol{Y}_t(i)$, $i = 1, \dots, n$, $t = 1, \dots, T$, as function of its inputs. In addition, every expert in G_i agrees on a marginal utility function over $\boldsymbol{Y}_t(i)$, if $i \in$ UT, $t = 1, \dots, T$.

We note here that the panels can come to this agreement in a variety of ways. As before, this consensus can be the result of decision conferences. Alternatively, the agreement can sometimes be found remotely, following for example a *Delphi Protocol* (see e.g. French et al. 2009). This agreement might even consist of planning to use a particular piece of software on agreed expert inputs, for example a particular probabilistic emulator (O'Hagan 2006). A panel might also agree to follow certain pooling axioms (see e.g. French 2011, Faria and Smith 1997, Wisse et al. 2008).

In practice, within the MDM framework, each panel will have to choose values (or prior distributions) for $\Sigma_t(i)$, $G_t(i)$, $W_t(i)$, $\boldsymbol{m}_0(i)$ and $C_0(i)$, for each possible combination of previous decisions.

Under the four assumptions introduced above, the probabilistic part of the IDSS can be represented, for each available policy, as a product of conditional distributions like in Equations (4)-(5), where now in addition decisions are included. Specifically,

$$f(\boldsymbol{y}^T|\boldsymbol{d}) = \prod_{t=1}^{T}\prod_{i=1}^{n} h\big(\boldsymbol{y}^t(i), pa\big(\boldsymbol{y}^t(i)\big), \boldsymbol{\theta}_t(i), \boldsymbol{d}^{t-1}\big),$$

where,

$$h\big(\boldsymbol{y}^t(i), pa\big(\boldsymbol{y}^t(i)\big), \boldsymbol{\theta}_t(i), \boldsymbol{d}^{t-1}\big) = \int_{\boldsymbol{\Theta}_t(i)} f\big(\boldsymbol{y}_t(i)\big|pa\big(\boldsymbol{y}_t(i)\big), \boldsymbol{y}^{t-1}(i), \boldsymbol{\theta}_t(i), \boldsymbol{d}^{t-1}\big)$$
$$f\big(\boldsymbol{\theta}_t(i)\big|pa\big(\boldsymbol{y}^{t-1}(i)\big), \boldsymbol{y}^{t-1}(i), \boldsymbol{d}^{t-1}\big) d\,\boldsymbol{\theta}_t(i).$$

It will be this product structure which will drive our results.

Now that the specification of an IDSS based on a MDM has been fully described, we can use this framework to compute coherent expected utility scores (see Leonelli and Smith 2013). In order to calculate these scores we derive a message passing algorithm between the panels and the SB that will then lead to their fast and transparent computation from the delivered judgments of the panels and to be placed before the incident team during the actual unfolding of the emergency.

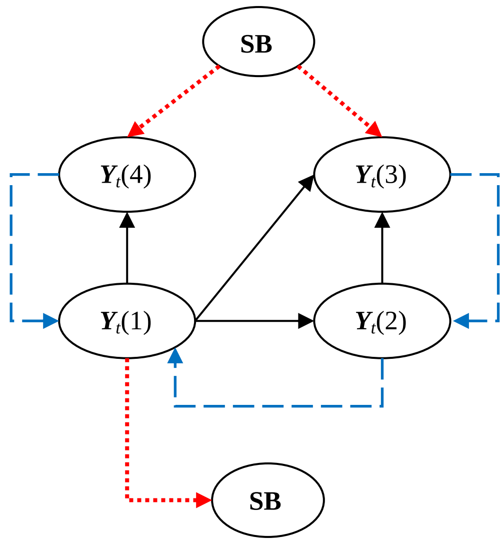

Fig. 2. Message passing algorithm for the time slice DAG of Figure 1.

The computation of the expected utility scores

Assume a potential DM enquires the system about the expected utility scores for the initial decision $\boldsymbol{D}_0$, given that she will behave optimally for the subsequent decisions. The IDSS needs to be able to compute such scores for each $\boldsymbol{d}_0 \in \boldsymbol{D}_0$, using only the judgments provided by the panels. This can be achieved under the assumptions introduced before, using a message passing algorithm between panels and the SB, which is based on standard backward induction. A full description and the general demonstration of this algorithm is given in Leonelli and Smith (2013). However a formal presentation of this algorithm to compute the expected utility scores in this dynamic framework requires the introduction of several supporting quantities and of a rather complex notation, which is beyond the scope of this short paper. So here for brevity we limit ourselves to simply describing how this algorithm works for an MDM based on the small time slice DAG given in Figure 1.

The algorithm is a very familiar type of backward induction on a development that we assume lasts T periods of time. However because of the modularity demand on the system it also has the flavor of a propagation algorithm over a Junction Tree (see e.g. Jensen and Nielsen 2007). The procedure we define is in fact recursive and works in exactly the same way within each of the time slice DAGs. The procedure recurses backwards from the final DAG (time T) back to the initial one ($t = 1$).

More specifically consider therefore the final DAG for $t = T$. The SB communicates the query to the panels overseeing leaf vertices, as described by the red arrows at the top of Figure 2. In this picture red arrows depict communications between the SB and panels, blue dashed arrows depict communications between panels, and the black ones represent statistical dependence in the time slice DAG of the MDM. The panels associated to leaf vertices (in the example G_4 and G_3) separately consider their T/i-th utility function and marginalize it with respect to the conditional distribution of $\boldsymbol{Y}_T(i)$. These panels then perform another marginalization step with respect to the conditional distribution of the associated parameter vector $\boldsymbol{\theta}_T(i)$. After these two steps, Panel G_3 communicates the result of the two marginalization steps to the panel overseeing the father vertex of $\boldsymbol{Y}_T(3)$, that is G_2. Specifically G_3 communicates the function $\bar{u}_{T,3}(\boldsymbol{d}^{T-1}, \boldsymbol{y}^{T-1}(3), \boldsymbol{y}^{T}(2), \boldsymbol{y}^{T}(1))$, defined as

$$\bar{u}_{T,3}(\cdot) = \int_{\boldsymbol{\Theta}_T(3)} f(\boldsymbol{y}_T(3)|\boldsymbol{y}_T(2), \boldsymbol{y}_T(1), \boldsymbol{d}^{T-1}, \boldsymbol{\theta}_T(3), \boldsymbol{y}^{T-1}(3)) \\ f\left(\boldsymbol{\theta}_T(3)\middle|\boldsymbol{d}^{T-1}, \boldsymbol{y}^{T-1}(3), \boldsymbol{y}^{T-1}(2), \boldsymbol{y}^{T-1}(1)\right) d\boldsymbol{\theta}_t(3)$$

Using a similar procedure G_4 communicates its result to G_1. These two message passing operations are described in Figure 2 by the blue dashed lines connecting the relevant vertices. At this point G_2 needs to add its T/i-th utility function to the incoming message from G_3. It can then perform the two marginalization steps over this new quantity. As described in Figure 2 by the dashed blue line from $\boldsymbol{Y}_T(2)$ to $\boldsymbol{Y}_T(1)$, G_2 then communicates the result, that is $\bar{u}_{T,2}(\cdot)$, to the panel overseeing the father vertex, i.e G_1. Now, G_1 performs the same steps as panel G_2 with the only difference that it communicates its result to the SB, since $\boldsymbol{Y}_T(1)$ is the root of the time slice DAG (red arrow at the bottom of Figure 2).

The expression the SB received from panel G_1 is a function of quantities with an index lower than T and consequently it can be optimized with respect to decision $\boldsymbol{D}_{T-1}$. In particular this optimized expression can be written as a linear combination of terms whose arguments can include only one of the variables associated to leaves of the DAG at time $T-1$. The SB at this point communicates each term of this linear combination to the relevant panel overseeing a leaf vertex, which adds this incoming message from the SB to its $T-1$/i-th utility function. Panels associated to leaf vertices (in the example G_3 and G_4) compute the two marginalization steps over this new quantity, similarly to the previous time slice, and communicate the result of these operations to the panel overseeing the associated father vertex in the time slice DAG. Then, as for the *T*-time slice DAG, each other panel needs to add the incoming message received from the relevant panel to its $T-1$/i-th utility function and compute the two marginalization steps. This process continues until the panel overseeing the root of the $T-1$ time slice DAG (in the example G_1) has communicated its result to the SB, who can then optimize with respect to decision $\boldsymbol{D}_{T-2}$.

The algorithm then continues to follow the same procedure described for the $T-1$ time slice, from time slice $T-2$ back to the initial one associated to time $t = 1$. Once the panel associated to the root of the first time slice DAG of the MDM has communicated its result to the SB, the expected utility is a function of the initial decision $\boldsymbol{d}_0 \in \boldsymbol{D}_0$ only, and consequently represents a numerical summary of the efficacy of a potential countermeasure. These summaries, computed through the methodology we described here, can then be used by the IDSS to rank available strategies by producing graphical outputs in order to explain to a potential DM the result of her decision analysis.

THE MDM FOR A NUCLEAR EMERGENCY

We now show through a simple example, how the algorithm developed in the previous section works within the MDM framework. Because of the length constraint we have to consider a very simple situation, since otherwise the expected utility expressions become long very quickly and so rather opaque. On the other hand, even in a very simple situation, it is possible to gain insights into the methodology. We note however that for larger systems, although the equations are much longer, they are analogous to those illustrated below and can be computed almost instantaneously in a simple algebraic form.

Thus assume a DM needs to decide whether to evacuate or not the population of a small village close to a nuclear reactor which is contaminating the surrounding area. She has only two time points in which she can decide the course of action, D_0 and D_1. Let $d_t(1)$ be the option of evacuation, while $d_t(2)$ corresponds to do nothing, $t = 0,1$. Assume also that she plans to observe the value of two continuous univariate time series $\{Y_t(1)\}_{t=1,2}$ and $\{Y_t(2)\}_{t=1,2}$ only at time $t = 1$, that is after having committed to D_0 but before having to choose D_1. Assume further that it is only considered relevant for the analysis the next value of these two series, that is $T = 2$. At time t, $Y_t(1)$ represents the amount of radioactive intake, while $Y_t(2)$ consists of the effects on human health. Note how the above decision framework corresponds to the one defined in the first part of Assumption 1, consisting of the sequence $(D_0, \{Y_1(1), Y_1(2)\}, D_1, \{Y_2(1), Y_2(2)\})$.

Assume that during decision conferences the probabilistic part of the IDSS has been agreed by the collective to be described by an MDM with a time slice DAG where $Y_t(1)$ is the parent of $Y_t(2)$, $t = 1,2$. This, together with possibly other relevant probabilistic issues (see below), corresponds to the agreement defined in Assumption 2.

Further assume the collective agrees that the two series are the only attributes of the problem and that they agree on using a simple linear utility, which can be written as

$$u\big(y_1(1), y_1(2), y_2(1), y_2(2)\big) = u\big(y_1(1)\big) + u\big(y_1(2)\big) + u\big(y_2(1)\big) + u\big(y_2(2)\big)\ .$$

This utility consensus corresponds to requirement specified in Assumption 3.

Assume also that the MDM is specifically defined by the following equations:

- $Y_2(2) = \theta_2(1,2)Y_2(1) + v_2(2)$;
- $\theta_2(1,2) = \theta_1(1,2) + w_2(2)$;
- $Y_1(2) = \theta_1(1,2)Y_1(1) + v_1(2)$;
- $Y_2(1) = \theta_2(1,1) + v_2(1)$;
- $\theta_2(1,1) = \theta_1(1,1) + w_2(1)$;
- $Y_1(1) = \theta_1(1,1) + v_1(1)$.

The errors $v_t(i)$ and $w_2(i)$ are assumed by the collective during decision conferences to be independent of each other with mean zero and variance $V_t(i)$ and $W_2(i)$, respectively, for $i, t = 1,2$. Each panel individually assumed these variances to be unknown but has provided a prior mean estimate $\lambda_t(i)$ for $V_t(i)$ and $\sigma_2(i)$ for $W_2(i)$. Assume further that the panels provided prior information about the parameter vector at time $t = 1$, such that $a_1(1,j)$ is the mean estimate for $\theta_1(1,j)$, while its variance is elicited to be $\tau_1(1,j)$, for $j = 1,2$. Note that each of these can be functions of the decisions with a lower time index. For simplicity, we drop this dependence, but it is important to remember that these values might be different for each available policy

Assume further that each panel individually believes that the marginal utilities under their jurisdiction are quadratic, so that $u\big(y_t(i)\big) = -y_t^2(i), t, i = 1,2$.

Note that all the above panels' specifications are delivered separately, as required by Assumption 4. Specifically, the form of the first three equations in the list above is decided by panel G_2, while the remaining ones are specified by G_1.

An illustration of how the system updates SB beliefs

Now that the IDSS has been fully defined for this example, we can show how the algorithm works symbolically when the overarching structure is the MDM. Since the utility is polynomial and has degree two, the algorithm consists of a sequential use of the tower properties for conditional moments. Specifically, recall that the following two identities hold for any two random variables X and Y.

$$\mathbb{E}(X) = \mathbb{E}\big(\mathbb{E}(X|Y)\big),$$
$$\mathbb{V}(X) = \mathbb{E}\big(\mathbb{V}(X|Y)\big) + \mathbb{V}\big(\mathbb{E}(X|Y)\big).$$

Now, suppose the system is asked to identify the expected utility scores for an initial decision $d_0 \in D_0$. The algorithm starts from panel G_2, which simply needs to compute $\mathbb{E}(-y_2^2(2))$. Through a marginalization step over $Y_2(2)$ (which in this polynomial setting corresponds to use the above identities), they can deduce that

$$\mathbb{E}(-y_2^2(2)) = -\mathbb{E}(\theta_2(1,2)y_2(1))^2 - \mathbb{E}\big(V_2(2)\big) - \mathbb{V}(\theta_2(1,2)y_2(1)).$$

The previous expression is still a function of variables overseen by G_2. Consequently, this panel needs to compute a further marginalization step over the relevant parameter space (again this corresponds to a sequential use of the identities). After this step G_2 derives that

$$\mathbb{E}\big(-y_2^2(2)\big) = -\mathbb{E}\big(\theta_1(1,2)y_2(1)\big)^2 - \lambda_2(2) - \mathbb{V}\big(\theta_1(1,2)y_2(1)\big) - \sigma_2(2)\mathbb{E}(y_2^2(1)).$$

Now this expression consists only of constant quantities, terms that are overseen by G_1, or elements associated to the previous time slice. Therefore G_2 communicates the above expression to G_1, which needs to add this incoming message to $u\big(y_2(1)\big)$. By performing the same steps as G_2, G_1 can derive that

$$\mathbb{E}\big(-y_2^2(2)\big) + \mathbb{E}\big(-y_2^2(1)\big) = -\lambda_2(1) - \mathbb{E}\big(\theta_1(1,2)\theta_1(1,1)\big)^2 - (\sigma_2(1) + \lambda_2(1))\mathbb{E}(\theta_1(1,2)^2) - \mathbb{V}\big(\theta_1(1,2)\theta_1(1,1)\big) - (1 + \sigma_2(2))(\mathbb{V}\big(\theta_1(1,1)\big) + \sigma_2(2) + \lambda_2(1) + \mathbb{E}(\theta_1(1,1))^2). \quad (6)$$

Note that these first steps simply consisted of a symbolic substitution of some of the unknown quantities included in the incoming message from the father vertex.

Although it is not explicit in the previous expression, since we dropped the dependence on the decisions, the previous expression does not include any random element which is a function of D_1. Thus, following the algorithm described above, these steps should be then performed:

1. G_1 communicates the above expression to the SB;
2. The SB identifies an optimal decision for D_1, as a function of the previous decision and of the values of the two series at the first time point;
3. The SB communicates to G_2 the optimized version of the rhs of Equation (6);

By recursively applying the identities for the first time slice DAG, the collective can then deduce that the overall expected utility factorization can be written as

$$\mathbb{E}\left(u\big(y_1(1), y_1(2), y_2(1), y_2(2)\big)\right) = -2a_1^2(1,2)\tau_1(1,1) - \tau_1(1,2)\big(\,a_1(1,1)^2 + \tau_1(1,1)\big) - \lambda_1(2) - 2a_1^2(1,2)a_1^2(1,1) - \lambda_1(1)(\,a_1(1,2)^2 +$$

$$\tau_1(1{,}2)) - \lambda_2(2) - \big(1 + \sigma_2(2)\big)\big(\sigma_2(1) + \lambda_2(1) + \tau_1(1{,}1) + a_1(1{,}1)\big) - \big(\sigma_2(1) + \lambda_2(1)\big)\big(\tau_1(1{,}2) + a_1^2(1{,}2)\big) - \tau_1(1{,}1) - \lambda_1(1) - a_1^2(1{,}1).$$

From this expression the SB can then deduce the expected utility scores of the initial decisions, assuming a subsequent optimal behavior. From distributed panel opinions the collective therefore now has a quick and defensible way of combining the different probabilistic beliefs of informed experts and agreed utilities into simple expected utility scores of options to present for discussion and examination within an incident scenario. But the distributed nature of the system has more than computational advantages. It gives a route along which the relevant panels can be queried as to the reasons for their contribution to the score. Such route is described by the diagram in Figure 3. Assume the IDSS is able to compute the expected utility scores of the available countermeasures, as stated in the top central box of the diagram. A DM might either be happy with the result produced by the IDSS and proceed with the suggested countermeasure or proceed with a further inspection of the factors that are driving the decision making. In the latter case a user might want to query the IDSS to provide the scores associated to the different attributes for a subset of the policies (second box in the central column). The IDSS can easily provide such scores (through graphical representations for example) which are then observed and explored by the user, who can then either adopt the suggested action, reappraise the model, or proceed with further inspections of the current model. In the latter case, within the IDSS framework, panels can be directly and separately addressed to provide the reasoning behind their delivered judgments. Thus, for example, assume that a user of the IDSS believes that the expected utility score associated with the amount of radioactive intake at the first time point is unreasonable. The appropriate panel (in this case panel G_1) will then be contacted (possibly on-line). Of course the explanation the panel provides might come in several different ways (i.e. graphs, natural language etc.), but we note here that symbolically the expected utility associated to that attribute is simply $-(\lambda_1(1) + a_1^2(1{,}1) + \tau_1(1{,}1))$. Since these are prior values elicited by the panel, they will need to justify their values' choices based on common and scientific knowledge. Such additional information might convince the user about the adequacy of the model used, in which case she then might adopt the suggested policy. Alternatively, she might believe that the model used is inadequate and modify the analysis on the basis of the further insights she got while exploring the system.

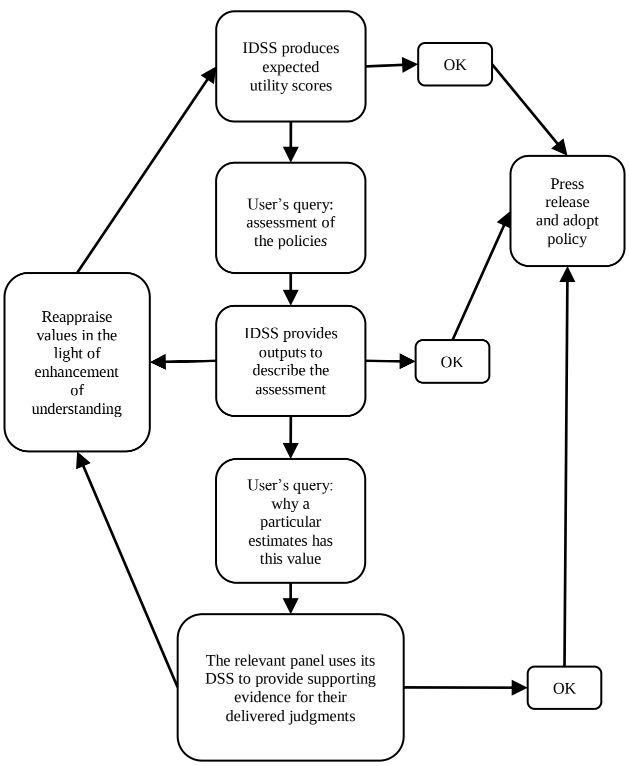

Fig. 3. Diagram describing the use of an IDSS.

DISCUSSION

The implementation of Bayesian methods for group decision analysis has often been considered too difficult to be developed. In this paper we have considered a framework in which, from a theoretical point of view, it is possible to deal with such class of problems. Of course, as shown by the above example, it is a daunting task to develop a system which can actually provide operational guidelines in nuclear emergency management. Note that even in this simple example, the panels need to provide a rather large number of values of hyper-parameters (recall that these are functions of the previous decisions and data). With only two decisions, two time series at two time points, one observational variance and only one regression parameter for each regression, there are 30 values the panels need to elicit (15 each). For the development of real systems it would be necessary the elicitation of a way larger number of values and processors that can deal with complex computations and huge amount of data. However, these quantities are provided by the complex IDSS. So although there are many of these, the actual computations are trivial ones, calculated on known outputs.

The authors are also currently investigating the use of computer algebra to study the polynomials representing the expected utility function. Such techniques might be able to identify the key variables of the polynomial, which are driving the decision making process. Once such elements are recognized, then the relevant panels can focus their attention on the elicitation of such quantities, enabling them to provide optimally useful outputs for the end user.

In this paper we have not dealt with the introduction of experimental and observational (from similar accidents) data into the IDSS before the development of the actual crisis. Of course this is a fundamental aspect an IDSS also needs to address, since data can reduce the involved uncertainties and consequently provide a better support. With sufficient care we are able to show that methods can be developed to deal with this. This will be the subject of a forthcoming paper.

ACKNOWLEDGMENTS

One of the authors is funded by the EPSRC grant EP/K007580/1 on “Management of Nuclear Risk Issues: Environmental, Financial and Safety”, within the NRFES project, and by the EPSRC grant EP/K039628/1 on “Coherent inference over a network of probabilistic systems for decision support with applications to food security”.

REFERENCES

1. A.E. ABBAS, "General decomposition of multiattributeutility functions with partial utility independence," *Journal of Multi-Criteria Decision Analysis*, **17**.1-2 (2010).
2. O. ANACLETO, C.M. QUEEN, C.J. ALBERS, "Multivariate forecasting of road traffic flows in the presence of heteroscedasticity and measurement errors," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, (2012).
3. J. BARTZIS ET AL. "RODOS: Decision support for nuclear emergencies," *In Decision Making: Recent Developments and Worlwide Applications*, 381-395. Springer US, (2000).
4. R.M. COOKE, *Experts in uncertainty: opinion and subjective probability in science*, (1991).
5. L. COSTA, J.Q. SMITH, T. NICHOLS, "On the selection of Multiregression Dynamic Models of fMRI networked time series," Crism Report, (2013).
6. R.G. COWELL ET AL. *Probabilistic networks and expert systems*, Springer, (1999).
7. P. DAWID, V. DIDELEZ, "Identifying the consequences of dynamic treatment strategies: a decision-theoretic overview," Statistics Surveys, **4** (2010).
8. S. DE, A.E. FARIA, "Dynamic spatial Bayesian models for radioactive deposition," Journal of Time Series Analysis, **32**.6 (2011).
9. A.E. FARIA, J.Q. SMITH, "Conditionally externally Bayesian pooling operators in chain graphs," The Annals of Statistics, **25**.4 (1997).
10. S. FRENCH, "Aggregating expert judgement," Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales. Serie A. Matematicas, **105**.1 (2011).
11. S. FRENCH, "Uncertainty modelling, data assimilation and decision support for management of off-site nuclear emergencies," Radiation Protection Dosimetry, **73**.1-4 (1997).
12. S. FRENCH, D RIOS INSUA, *Statistical decision theory*, Kendall's Library of Statistics 9. Arnold, (2000).
French, S, D C Ranyard, and J Q Smith. Uncertainty in RODOS. University of Leeds, School of Computer Studies, 1995.
13. S. FRENCH ET AL. *Decision behaviour, analysis and support*, Cambridge Press, (2009).
14. C. GONZALES, P. PERNY, "GAI networks for decision making under uncertainty," *IJCAI-05* (2005).
15. F.V. JENSEN, T.D. NIELSEN, *Bayesian networks and decision graph*, Springer (2007).
16. R.L. KEENEY, H. RAIFFA, *Decision with Multiple Objectives: Preferences and Value Tradeoff,*. Cambridge Press (1993).
17. M. LEONELLI, J.Q. SMITH, "A Graphical and Probabilistic Integrating Decision Support System for Optimal Uncertainty Handling," Crism Report, (2013).
18. M. LEONELLI, J.Q. SMITH, "Using graphical models and multi-attribute utility theory for probabilistic uncertainty handling in large systems, with application to the nuclear emergency management," *ICDEW* (2013a).
19. S.A. MURPHY, "Optimal dynamic treatment strategies," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, (2003).
20. A. O'HAGAN, "Bayesian analysis of computer code outputs: a tutorial," *Reliability Engineering & System Safety*, **91**.10 (2006).
21. P.K. MURPHY, "Dynamic Bayesian networks: representation, inference and learning," PhD Thesis, (2002).
22. C.M. QUEEN, "Using the multiregression dynamic model to forecast brand sales in a competitive product market." *The Statistician*, (1994).
23. C.M. QUEEN, J.Q. SMITH, "Multiregression dynamic models." *Journal of the Royal Statistical Society. Series B (Methodological)*, (1993).
24. K.F. RICHTER ET AL. "Data assimilation for assessing radioactive contamination of feed and foodstuff." *In Devolpment and application of computer techniques to environmental studies*. WIT Press (2002).
25. S. THYKIER-NIELSEN ET AL, “Description of the atmospheric dispersion module RIMPUFF”. Riso National Laboratory, (1999).
26. M. WEST, J. HARRISON, *Bayesian forecasting and dynamic models*, Springer, (1997).
27. B. WISSE ET AL. "Expert judgment combination using moment methods." *Reliability Engineering & System Safety*, **93**.5 (2008).