Presenting Uncertain Information in Radiological Emergencies

Extended Executive Summary

S. French and J. Q. Smith (University of Warwick)
N. Argyris (University of Loughborough),
S. Haywood (Public Health England), and
M. Hort (The Met Office)

BACKGROUND

In 2014, the Atmospheric Dispersion Modelling Liaison Committee (ADMLC) commissioned Warwick University in collaboration with Public Health England (PHE) and the Met Office to:

- develop improved presentational techniques for representing the uncertainties and lack of knowledge in the early stages of a radiological emergency;
- build an improved, shared understanding and realistic expectations between decision-makers, scientists and communicators of what will be known in the early phase of a radiological emergency and how this knowledge, particularly relating to the areas at risk as any plume spreads, will evolve.

In Spring 2016, a report[1] was prepared and presented to ADMLC. Given its length, it was decided to produce this shorter summary document for wider circulation. All references and supporting material are provided in the longer document, hereafter PUIRE. A list of abbreviations used is provided on page 11.

[1] S. French, N. Argyris, H. Layton, J.Q. Smith, S. Haywood and M. Hort (2016) Presenting Uncertain Information in Radiological Emergencies. RISCU(14)7. Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK.

This study was funded by the Atmospheric Dispersion Modelling Liaison Committee. The views expressed in this report are those of the authors, and do not necessarily represent the views of ADMLC or any of the organisations represented on it.

1 INTRODUCTION

The project's focus was on how information would be presented to the Scientific Advisory Group for Emergencies (SAGE) and how they would prepare advice for COBR, the UK's national crisis response group within central government[2]. It did not consider the many similar issues which arise in the co-ordination of the local response. Moreover, it concentrated on the many substantial uncertainties[3] that would need to be faced in the early hours of an event.

The project involved a range of activities, including a substantial literature review; however, its key elements related to three workshops[4], all using hypothetical scenarios to focus their discussions and illustrate the many uncertainties that arise in responding to a radiation accident. During each workshop the scenario was presented, stepping through the first few hours of the accident and explaining what would be known at each time, what would not be known, what seemed most likely to happen, and what the radiological and health impacts might be. The first workshop sought to understand the current processes of information presentation and discussion within SAGE. It involved members of Government departments and agencies, who might well be called into SAGE during an actual radiological emergency. Discussion focused on how to advise COBR on the significance of the uncertainties involved in predicting the course of the plume, the impact of this on health and the likely need to prepare resources to support recovery. Building on this experience, the project developed proposals for presenting information on the potential geographical spread and impact of a radiation plume. The second workshop involved many world experts on the presentation of scientific and expert advice in high risk contexts, and aimed to challenge and criticise these proposals. The third workshop had similar attendance to the first, but this time focusing on the presentation of information using plots, graphs, and other display techniques proposed by the project to convey the uncertainty, particularly its geographical aspects, and then to reflect on how useful the different approaches were.

2 UNDERSTANDING AND COMMUNICATING RISKS

The project undertook an extensive review of current thinking on risk and decision behaviour, with attention being paid to the communication of uncertainty particularly[5]. Studies suggests that human response to uncertainty can – simplistically – be categorised in two forms: System 1 and System 2 Thinking. The former, often referred to as 'intuitive', involves superficial analysis

[2] PUIRE Sections 1.2, 3.1 and 3.2.
[3] PUIRE Section 1.3.
[4] PUIRE Chapters 4 and 5.
[5] PUIRE Chapter 2.

and interpretation of the relevant information which takes place at the fringes of or outside consciousness. System 1 Thinking essentially comprises the 'hard-wired' reactions that are generated in the face of some event, and it can lead to some very unwise judgements and behaviours. System 2 Thinking is characterised by conscious analytical thought that involves a detailed evaluation of a broad range of information, often based on a rule that is assumed to provide the 'correct' answer or solution. Formal risk analyses are examples of System 2 Thinking that have been validated against experience over many years. It is clearly this form of thinking that should be encouraged in deliberations within SAGE and between SAGE and COBR.

Amongst the many 'shortcomings' of System 1 Thinking, the following are relevant.

- The *Plausibility Effect* in which people substitute plausibility for probability. If someone imagines a scenario in some detail, then he or she will perceive it as more likely than evidence would suggest that it is. Similarly, the *Availability Effect* suggests that recalling recent events from the past will make similar events seem more likely in the future.

- The *framing* of statements can affect people's attitude to risk. If potential outcomes of events and actions are framed positively, i.e. in terms of safety, good health, profit, etc., people become more risk averse; if framed negatively, i.e. in terms of death, cancer, loss, etc., they become more risk prone.

- *Qualitative descriptions of uncertainty* – e.g. 'likely', 'improbable' – are unreliable ways of communicating. Even providing *linguistic probability lexicons* in which such words are associated with precise quantitative ranges of probability is unreliable, unless the *same* lexicon is used continually in their professional lives by all parties to the conversation.

Expertise provides no immunity to falling prey to such foibles of System 1 Thinking. Experts find it difficult to believe they have limited forecasting ability and are prone to overconfidence, hence underestimating the uncertainty in their assessments. This is particularly true when they fail to recognise the novelty in a situation; and radiation accidents are fortunately very rare and outside the experience of many experts. Moreover, they can also be prone to slips and errors in their thinking. For example in the third workshop, it was apparent some participants misinterpreted probabilities conditional on the occurrence of a second release as unconditional probabilities, at least for part of their deliberations. Continual questioning and challenge are the key to keeping experts (and others) 'within the straight and narrow of System 2 Thinking'.

The full report also contains a review of these issues in the context of geographical uncertainty[6].

3 REASONABLE WORST CASE: A CRITIQUE

There are many uncertainties inherent in the early hours of a radiation accident: lack of knowledge of the source term, its scale, energy, profile and timing; the effectiveness of engineering actions to cap the release in the immediate future; meteorological uncertainties, including the timing and location of fronts and precipitation; computational and physical approximations used in predicting the dispersion and deposition of contamination and its health effects; public response and compliance with advice; etc[7]. At present, none of these uncertainties are quantified in the information that the agencies, responders and plant operators provide to SAGE. Consequently SAGE (and COBR) currently seek to develop and work with a *reasonable worst case*. The idea of a reasonable worst case[8] is common in emergency planning, where it is defined as being designed to exclude theoretically possible scenarios which have so little probability of occurring that planning for them would lead to a disproportionate use of resources. The concept has been taken over from emergency *planning* into emergency *response* without apparent recognition that the contexts of these two activities is significantly different. The former considers the possibility, remote or otherwise, of some disaster. The latter relates to something that has most definitely happened.

While a reasonable worst case – or one might suggest, several reasonable worst cases – are essential in emergency planning to ensure sufficient, but not excessive resilience is built into a system, it is far from clear that emergency response should focus almost entirely on a single reasonable worst case. Framing issues may increase the risk proneness within SAGE and COBR, while the plausibility effect may make the effects of the reasonable worst case seem more likely. Moreover, there may be many different negative impacts (health, agricultural, economic, etc.) that could arise and some may not be visible in a single reasonable worst case. Finally the advice and assessments from SAGE are sought by COBR in relation to what should be *done*: their purpose is to support decision making. It is not clear that describing a reasonable worst case is the most helpful form of information for this. The focus of a reasonable worst case is simply on what might happen. It does not offer an analysis of what might happen were different actions taken. For these reasons, the project investigated whether SAGE would find it useful if they were presented with several scenarios spanning the scale of possible impacts from reasonable worst cases to possible,

6 PUIRE Section 2.4.
7 PUIRE Section 1.3
8 PUIRE Section 4.3

perhaps more probable outcomes with much less impact[9]. The third workshop presented four scenarios of differing scales and forms of impact and demonstrated that this widened the discussion initially, although the group did fall back to taking one of the scenarios as 'the' reasonable worst case on which to base their advice to COBR.

4 DISCUSSION AND RECOMMENDATIONS

The recommendations below, numbered as in the full report, are made from the perspective of the handling of a radiation accident. We recognise that the constitution, processes and procedures of SAGE and COBR need be designed for a much wider range of crises and that our recommendations might conflict with other needs for other events. Our recommendations fall under three headings.

4.1 Organisational and Logistical Issues

In recruiting participants to the workshops, we discovered that Civil Service procedures for career development entail staff move regularly between departments and posts. Thus many of the invited 'experts', who would be candidates for membership of SAGE during a real accident, were new to their posts and had little experience related to radiation accidents. Several of the first workshop participants had moved to unrelated posts by the time of the third workshop 14 months later. If our experience is a guide, it is possible that 25%-40% of the experts sent by agencies and ministries to form SAGE may be inexperienced and not have attended any major exercise. This potential presence of inexperience and consequent reduced expertise in the specifics of responding to a radiation accident is a significant constraint on the discussions within SAGE in relation to understanding and addressing the uncertainties, particularly since much of the uncertainty is currently unquantified.

Recommendation 1: Attention should be given to the effects of promotions and career development within the Civil Service and Government agencies on the expertise that may be available to SAGE during a radiation accident – and presumably other events.

In running the project, we learnt of a number of factors that limit the format and quality of the information presented to SAGE and COBR. For instance, for security reasons the software available in the meeting rooms is limited to well-tested office products. Similarly it is not easy to link to systems run remotely at the Met Office or PHE. Thus computer models such as those run within JAM[10] cannot be run nor their outputs interactively interrogated within SAGE to calculate quantities at a specific point nor to answer 'what-if' questions. Nor are there

[9] PUIRE Section 4.5

[10] The Joint Agency Model – strictly not a single model, but several agency models co-ordinated to give coherent output.

simple ways of running video sequences showing the evolution of plumes over time or 'jittering' plumes to indicate uncertainty. Materials to be presented to SAGE will come from a variety of sources: word documents, powerpoint slides, emails and, possibly, telephone messages. Drawing these together into a coherent set of materials with common use of terminology, units, formats, etc. and then summarising them for COBR is a significant challenge. The advent of guidance notes and key questions for the Chief Scientist and SAGE during a radiation accident will reduce this challenge, and the development of formal reporting templates could significantly reduce the risk of poor information capture and communication and that of failure to take into account some piece of information. Information system scientists would address these issues through an executive information system to pull together information from multiple sources, automatically using common scales, axes, etc. for graphs and plots and producing the required tables to compare different scenarios, etc. However, as noted, such systems are not currently available within the security cordon around SAGE and COBR. We recognise that such constraints may be in place for the best of reasons, but we also note that they affect the implementation of best practice in the emergency management of a civil radiation accident.

Recommendation 2: There are logistical, support and organisational issues which limit how information can be presented to SAGE and COBR. There may be benefit in reviewing whether the need to present a greater range of information e.g. as in the case of JAM, requires some modification of the structure and organisation of the communication and information presentation within SAGE and COBR.

No accident ever goes as 'planned'. The academic members of our team were quite surprised to discover that most UK national exercises only rehearse response to design-basis accidents. Moreover, those that do consider more substantial events may not do so in full detail. Nor is it clear that there are enough exercises to provide sufficient experience to all potential members of SAGE. Our exercises went substantially beyond design-basis and thus stretched the participants' thinking, gaining new insights. For instance, in both workshops there was a realisation that the need for stable iodine might exceed its local availability and hence, as a contingency, it would be worth moving supplies to the area from the national stockpile.

Recommendation 3: There may be benefit in exercising SAGE (and other bodies) with more significant accident scenarios than are conventionally used.

As noted above, experts are not immune from making slips and errors in interpreting or analysing information or making judgements. Facilitators of problem solving workshops use gentle but insistent challenging questions and interventions to counter such slips into System 1 Thinking. However, although emergency management is undoubtedly a problem solving context, carefully-paced reflective and challenging processes are too slow to be incorporated into the workings of SAGE. Nonetheless, deliberations within SAGE should incorporate as much challenge as possible; and they do incorporate challenge. Scientists invariably question the evidence and reasoning behind their colleagues' statements. However, it is important that guidance documents for chief scientists recognise this and encourage challenge whenever possible.

Recommendation 4: Process briefing documents for chief scientists and participants in SAGE should recognise the importance of bringing 'challenge into the room' to reduce the risk of errors, slips and misinterpretation.

4.2 Presenting Uncertainty and other Information to SAGE and thence to COBR

A common observation made in all three workshops and indeed in our preparations for them was that departments, agencies and, particularly, software output used different conventions, units, map scales, colours and symbols, etc. for reporting consequences. It can be difficult to align maps to show different aspects of the release and its management. The urgency of any emergency means that there is no time to explain notation, the choice of units, integration periods, re-plot maps to common scales, develop legends and explanations of graphs, etc. It should be possible to reduce the risk of confusion within SAGE and subsequently COBR by developing templates for capturing the information, along with standardised explanation of the terms for both bodies. Input templates would have the additional advantage of acting as an aide memoire of the minimum information that each body would be expecting to receive.

Recommendation 5: Standard templates, legends and explanations relating to all maps, plots, tables for both SAGE and COBR should be developed in advance.

The advent of JAM provides an opportunity for all geographical plots to be provided for SAGE and COBR from a common source, allowing consistent use of scale, colour, etc. JAM could also produce automatically all the maps, plots and tables that would be needed by SAGE. Plots that contour emergency reference levels and thus are focused on potential actions and countermeasures may be more helpful than simple maps of dose or deposition. Several authors, identified in our literature review, also emphasised the need for action-oriented plots.

Recommendation 6: The presentation of observational and modelling data should be implemented with consistency in the use of scales, units, colour, etc. This is particularly true of geographical information, which should be presented using maps that can be easily aligned. Ideally once the source term and meteorology have been set for a scenario the output should be developed and produced automatically by the system providing an agreed set of maps, tables and plots for SAGE without further intervention or collation. Where possible, these should be designed to support discussion of potential countermeasures, rather than simply show contours of dose or deposition.

Initially we believed the project's focus would be on communicating geographical uncertainty about the spread of an atmospheric plume of radioactive contamination and the consequent uncertain predictions of dose maps; and indeed much of our focus has been on understanding the likely spread. However, we rapidly understood that many of the key uncertainties that are discussed within SAGE relate to non-geographical events and parameters, in particular the source term. Moreover, current practice does not use a formal way for addressing these uncertainties. Probabilities are not offered to SAGE. The two sets of key uncertainties relate to the weather and the source term. Although many atmospheric dispersion models are essentially stochastic, the output currently prepared for SAGE and COBR is deterministic, depicting a single plume. In the case of the source term, the emphasis seems to be on describing the physical situation, the fault(s) and the engineering actions being taken, along with best guesses of when the release might be capped, and its scale and evolution until then. Members of SAGE have to internalise this information to form their own assessment of the uncertainty.

It was suggested to us several times that a probability lexicon might provide a way of communicating and discussing uncertainties, particularly in the early phase when uncertainties abound and information is sparse. However, unless members of SAGE are fully familiar with a standardised use of the same lexicon and use it regularly in their daily lives, this would be very likely to cause more confusion and miscommunication than it resolves. Any steps in this direction should be taken only after such an approach has been fully adopted across government.

Recommendation 7: SAGE should not adopt a probability lexicon to give quantitative meaning to everyday expressions of uncertainty unless and until a common lexicon is adopted and used consistently across all government departments and agencies in their day-to-day activities.

It is possible that we are being too defeatist in thinking that it is impossible to get some quantitative probabilities for some of the key early uncertainties. Might the operators, ONR or some others be prepared to give some very rough probabilities? The operators were not at all involved in this project and ONR was only peripherally involved as workshop attendees. Of course, anyone who argues that a probability lexicon could be implemented is also arguing implicitly that it is possible to give rough probabilities, because that is what a lexicon is based on. To develop a range of 3-5 scenarios, as we recommend below, it would be sufficient to have some rough indication on the balance of probabilities between small and large releases, its duration and roughly what radionuclides might be present. So this would be worth exploring. It might also be the case that SAGE is able to work with quantitative probabilities, recognising the limits of their accuracy at this stage; but that it would be unwise to pass such rough numbers onto COBR.

Recommendation 8: Discussions with the operators, ONR and other relevant parties should take place to see if it were possible to get some very rough quantitative probabilities relating to the source term in the early stages of the event.

A key point made in both the first and third workshop relates to the importance of setting clear expectations of when further information will be available and uncertainties be resolved or reduced. Members of COBR and to a lesser extent some members of SAGE may have inaccurate perceptions of what is realistically achievable in the early hours, leading to a loss of confidence in scientific advice as the information picture evolves. Thus the initial report to SAGE and that subsequently provided to COBR should provide a timeline for information flows.

Recommendation 9: Timelines relating to the availability of further information in respect of each key uncertainty should be provided to SAGE and COBR.

4.3 Discussion of Uncertainty within SAGE and Reasonable Worst Case

We are concerned that focusing on a single reasonable worst case might lead to flawed thinking and an overemphasis of the risks of significant escalation. Moreover, we have noted that a reasonable worst case describes what might happen if nothing is done, but it does not offer the basis to decide between different possible strategies. However, it seems clear that the use of a reasonable worst case is embedded in the processes used by COBR. That does not mean that it need be embedded to the exclusion of all other possibilities from discussions within SAGE. It is important that SAGE prepare balanced advice for COBR which reflects reasonable expectations of the evolution of the accident, as well as giving guidance on what resources might need to be put in place if that evolution is at the worse end of the spectrum. If public confidence is to be maintained, it is important that the authorities are seen to be anticipating and mitigating the possible course of the accident, including that which actually occurs. Moreover, even if COBR's attention in the very early phase is focused on short term risks to human health from direct exposures, some indication of potential longer term health risks, e.g. from food and water, and the scale of the countermeasures needed should also be given. Given that politicians need to be seen to have a comprehensive view of the potential consequences of the accident if public trust is to be maintained, SAGE should briefly consider whether any specific significant long term economic or environmental impacts might occur and include brief mention in their report to COBR.

We believe that this can best be achieved through the use of multiple scenarios, perhaps 3-5. These should include a likely case to set reasonable expectations, 1 to 3 reasonable worst cases, and a best case to provide a counter to the pessimism of the latter. The need to ensure that national resources are prepared does mean that there has to be a bias towards reasonable worst cases. The second and third workshops confirmed us in the belief that this approach would support a more balanced deliberation within SAGE, while still enabling SAGE to provide the form of advice that COBR require. Current processes and timescales would probably limit the number to 3-5 scenarios, although we can imagine circumstances in which even 5 would not fully scope the possibilities.

Recommendation 10: SAGE should be provided with 3-5 scenarios which together provide an overview of the range of possible impacts that might result from the accidental release.

Our presentation of the scenarios and the consequences of the accident were essentially based on maps in the third workshop. Several participants made the very sensible suggestion that it would have been helpful to provide tabular summary information to compare the potential impacts of the different scenarios.

Recommendation 11: SAGE should look at all scenarios prepared to explain the range of possible impacts. To aid in this, the geographical plots prepared for each scenario should be supplemented by a brief list of the key impacts in tabular or bulleted form. A template for doing this should be prepared. Moreover, the design of any supporting IT systems such as JAM should provide the key tables, though some of the more qualitative comparisons will need to be summarised by hand.

The Report[11] suggests a constructive way of developing scenarios, though it is not an easy task. Local risk registers should contain pointers to features which might have implications for the seriousness of different plume paths. Even though there will be huge pressures on time, it is important to recognise that the procedure for developing scenarios is likely to be iterative rather than linear and to involve selection from rather more scenarios than the 3-5 to be presented to SAGE. Once JAM is developed and commissioned, it should be possible to generate scenarios relatively quickly; and the production of comparative tables should help in the selection of 3-5. Nonetheless, there is a need to develop fuller procedures and guidance to produce the scenarios for SAGE; and this should be undertaken in collaboration with the development of JAM.

Recommendation 12: Procedures and guidance for constructing the 3-5 scenarios to present to SAGE should be developed. These procedures should be developed and exercised in collaboration with the designers and developers of supporting IT systems such as JAM.

When JAM is implemented, the teams at the Met Office, PHE and other agencies running JAM should have the responsibility for developing and selecting scenarios for SAGE. Many approaches to problem solving would suggest that the problem owners, in this case SAGE, should also be involved in the development of the scenarios in order to explore their concerns. Given the urgency of this context, that is clearly not practical. Nonetheless, if it becomes clear during the SAGE meeting that a further possible scenario should be examined, it should be possible for that scenario to be generated.

Recommendation 13: It should be the responsibility of the teams using supporting IT systems, e.g. JAM, to identify and develop the scenarios to present to SAGE. Ideally, if SAGE wish to see a further scenario, it should be possible for a request to be made from within SAGE, the necessary runs made and the results sent back into SAGE.

11 PUIRE Section 4.5

Several participants at the workshop expressed a wish for some rough probabilistic assessment of the risks. There was not capacity in this project to explore that idea to any depth, although some structuring and discussion was given to the idea[12]. There is merit in exploring this suggestion further, recognising that such a development would require the use of structured expert judgement.

Recommendation 14: Consider an exploration in the longer term of the potential for providing SAGE with probabilities.

ABBREVIATIONS AND ACRONYMS

ADMLC	-	Atmospheric Dispersion Modelling Liaison Committee
COBR	-	UK's national crisis response group (known from its location: Cabinet Office Briefing Room)
JAM	-	A shorthand for the current development of the Joint Agency Modelling procedures and processes to provide timely plots and predictions to SAGE, drawing together the output of several agencies.
ONR	-	Office for Nuclear Regulation
PHE	-	Public Health England
SAGE	-	Scientific Advisory Group for Emergencies, the group of experts who directly advise COBR.

[12] PUIRE Section 4.5