Truth or Consequences: Using argumentation to reason about risk

Peter McBurney and Simon Parsons
Department of Computer Science
Chadwick Building
University of Liverpool
Liverpool L69 7ZF
United Kingdom

{P.J.McBurney,S.D.Parsons}@csc.liv.ac.uk

December 14, 1999

Abstract

Policy debates about the impact of new technologies or substances often begin with a proposal that the innovation poses potential hazards to those exposed to or using it. Such potential hazards are typically disputed by other scientists on the basis that no theoretically-sound causal mechanism exists to explain the alleged relationship. Debates often then proceed with an experiment which demonstrates a statistical correlation, but articulation of a sound causal mechanism may only follow much later. Examples of such debates include the link between smoking and lung cancer; the link between BSE and CJD; and the current debate on GM foodstuffs.

Although rarely stated as such, these debates often include implicit attacks by one side on the modes of inference used by the other. Accordingly, it would seem useful to attempt to represent the debates in formal terms, and the authors have used argumentation to do this. We present an example where at least ten distinct modes of inference are used to assert carcinogenicity, only one of which is statistical. As this example shows, an argumentation formalism enables the precise elucidation of the modes of inference deployed by various sides and their relationships to each other.

In addition, argumentation permits coherent reasoning about the consequences and likelihoods of alternative courses of action, even when expressed in qualitative terms.

1 Introduction

We seek to build computer systems which can reason autonomously about alternative actions, informed by predictions of their possible consequences. When the consequences of actions can be articulated, their relative significance expressed in quantitative terms, and the likelihoods of relevant uncertain events expressed as probabilities, classical decision theory [56, 39] provides a coherent framework for undertaking reasoning about actions; such a framework is readily automated.[1] However, estimating and agreeing quantitative probabilities and utilities (or, equivalently, losses) is not straightforward in most real-world domains.[2] We are therefore motivated to explore qualitative approaches to practical reasoning, and this paper presents an application of argumentation to this end.

The application to which we apply this is scientific reasoning about the possible carcinogenicity of some chemical substance. Typically, policy debates about the health and environmental impacts of new substances and technologies are initiated when a possible adverse health effect is observed, often on the basis of only a handful of cases.[3] This usually then leads to an experimental study which has the potential, if positive results are obtained, to demonstrate a statistical correlation between the chemical substance or the technology and these effects. However, even when such a correlation is found, there may not be a scientific explanation for the phenomenon. Indeed, a theoretically-sound causal mechanism may take a considerable time to be developed, articulated, tested and agreed by the scientific community concerned. Debates which have followed or are following this pattern include: the relationship between cigarette smoking and lung cancer; the alleged relationship between Bovine Spongiform Encephalopathy (BSE) and Creutzfeldt-Jakob Disease (CJD); and the current debate on the possibly adverse effects of Genetically-Modified foodstuffs.[4]

Although rarely stated as such, these debates often include attacks by par-

[1] As an example, Gibney *et al.* [22] use decision theory in an automated electronic auction system for allocation of telecommunications bandwidth.

[2] For instance, John Fox and his colleagues [17, 18, 36] have demonstrated the difficulty in practical risk assessment of the very first task in this formalism, that of articulating the consequences of alternative actions.

[3] For example, the toxic impacts of thalidomide were first noticed by practicing doctors. Although the drug had undergone animal and human trials before its commercial release, these trials had not included any pregnant subjects and they did not reveal any adverse effects [66]. Fatal bone marrow diseases arising from prolonged exposure to adhesives were first noticed by a Turkish physician, Muzaffer Asksoy, in treating workers in the shoe industry [26]. Recent public debate on the possible association between the use of mobile phones and brain cancer has been based on a handful of observed cases.

[4] In the case of smoking and lung cancer, there was considerable debate over the causal mechanisms involved. Indeed, a prominent psychologist even proposed mutual, genetic causes for both effects [16]. In the case of BSE-CJD, Government scientists and policy makers initially assured the public there was no link between the two diseases. When a link was acknowledged, the most commonly accepted explanation was that cows fed with infected sheep brains contracted BSE and then passed the infection onto humans, who then contracted CJD. However, it has recently been suggested that the original cause was a failed cattle fertility treatment experiment [3].

ticipants on the modes of inference used by each other, particularly before a sound causal mechanism has been agreed.[5] In many science-policy debates these inferences are not explicit. Accordingly, it would seem valuable to articulate, comprehensively and precisely, the modes of inference being used in such a debate. However, we know of no published effort to do this in a science-policy arena.[6] As we wish to automate the reasoning processes involved, we have therefore articulated, as best we are able from outside the field, the precise modes of inference used when an assertion of carcinogenicity is made.

Thus, Section 2 of this paper articulates the modes of inference used in an archetypical claim of carcinogenicity of a chemical substance, and has a brief discussion of these modes in comparison with statistical inference. Section 3 introduces the ideas of argumentation and presents the argumentation formalism which we will apply to the carcinogenicity domain. Section 4 concludes with a discussion of ongoing and further work.

What value are computer systems which can reason about risk? The Imperial Cancer Research Fund (ICRF) in London has developed systems to predict the risk of carcinogenicity of new chemicals based upon their chemical structure [67] and systems to predict the risk of individuals developing breast cancer based upon their life choices and their family disease history [14]. The ICRF were also pioneers in the application of argumentation to intelligent systems [19]. Similar systems have been deployed commercially for predicting the toxicity of pesticides [33] and the risks associated with chemicals in food [32]. However, the largest risk-prediction system known to us is Project Genoa, being developed at a cost of US$50 million by the U.S. Department of Defense to undertake automated reasoning about international geo-political risks [50, 62]. When operational, Genoa will combine evidence and reasoning across a number of relevant domains (economic, social, political, military, etc) to assist in the early identification and management of military and political crises around the globe. Genoa is also being built using argumentation.

2 Inferring Carcinogenicity

2.1 Modes of inference

On what basis do scientists claim that a chemical substance is carcinogenic? Such claims can be based upon evidence from a number of sources (adapted from [15] and [26]):

[5] For instance, the debate over the possibility that exposure to formaldehyde causes nasal cancers [25] has drawn on conflicting evidence from multiple animal and human studies, and centred on the acceptability of inferential reasoning from these various sources. Similarly, Shere [64], questions the validity of many of the inferences used in the typical environmental and health risk assessment processes of U.S. Federal Government agencies.

[6] Elucidation of inference processes and arguments used in some other domains has been undertaken. For example, McCloskey [44, 45] articulates the arguments used in economics, and demonstrates the use of traditionally rhetorical and logically-fallacious devices even in ostensibly mathematical economics research. Likewise, but for different purposes, Pera [52] explicates the modes of reasoning used in several natural scientific debates.

- Using chemical theoretical reasoning, on the basis of the chemical structure of the substance and the known carcinogenicity of chemicals with cognate structures.
- From mutagenic tests, applying the substance to tissue-cultures in laboratory experiments.
- From experiments involving the application of the chemical to human or animal cadavars.
- From bioassays, applying the substance to animals in a laboratory experiment.
- From epidemiological studies of humans, either case-control studies (where a case group of people exposed to the substance is matched with a control group not so exposed, and their relative incidences of cancer compared), or cohort studies (where the incidence of the cancer among people exposed to the substance is compared with that in the general population, while controlling for other potential causal and interacting factors).
- From elucidation of theoretically-sound bio-medical causal pathways.[7]

As mentioned in the Introduction, elucidation of causal pathways is generally not undertaken until evidence of an empirical nature is observed. Hence, we focus on the other categories of evidence. There are a number of comments one can make on the relative value of these different approaches. Reasoning from chemical structure is still an imprecise and immature science for most substances; indeed, automated prediction of carcinogenicity and other properties of chemicals on the basis of their structure is an active area of Artifical Intelligence research [30, 65]. Mutagenic tests may demonstrate carcinogenicity in principle, but do not reveal what will happen in a whole, living organism (with, for instance, viral defences), nor in an environment similar to that of people exposed to the substance. Experiments with cadavars have similar difficulties. Moreover, because the incidence rates of many cancers are very small, epidemiological studies may require large sample sizes, and so can be quite expensive. Also, the time-lag between exposure to typical environmental doses and the onset of a cancer can be very long (in the order of decades), so, unless undertaken retrospectively, these studies can take years to complete. For these reasons and others, the most common form of assessment of potential carcinogenicity is the bioassay.

We therefore turn our attention to animal bioassays. Because of the difficulties in inferring conclusions about humans on the basis of evidence about animal species, most cautious scientists and policy makers would not *assert* carcinogenicity to humans from a bioassay: they would, at best, only claim

[7] These are E-theories in Pera's [52, page 154] typology of scientific theories.

that there is a (perhaps high) probability of human carcinogenicity.[8] However, although it is perhaps the most contentious, the animal-to-human inference is not the only inference being deployed in concluding such a probability. It is also not the only inference deployed when quantifying the extent of risk. It therefore behooves us to examine all the modes of inference used. In doing so, we have abstracted from a number of descriptions and critiques of carcinogenic risk assessment processes [5, 7, 9, 15, 24, 25, 26, 31, 42, 46, 48, 54, 61, 64, 78], both ideal and actual. However, as mentioned earlier, we believe this is the first published attempt to develop a comprehensive list such as this.

For the purposes of exposition, we therefore suppose an archetypical animal bioassay for a chemical substance $\mathcal{X}$ is undertaken. This will involve the administration of specific doses of $\mathcal{X}$ to selected animal subjects, usually repeatedly, in a laboratory environment. Typically, two or three non-zero dose-levels are applied to the subject animals, along with a zero-dose to the control group. The rates at which cancers of a specific nature develop is then observed in each group until a pre-determined time-point (usually the natural life-span of the animal). Those animals still alive at that time are then killed, and a statistical analysis of the hypotheses that exposure to the substance $\mathcal{X}$ results in increased incidence of cancer is then undertaken. Suppose that, based on this animal bioassay, a claim is then made that $\mathcal{X}$ is carcinogenic to humans at a specified dose. For ease of expression we will notate this claim by ϕ. In asserting ϕ from the evidence of the bioassay, a number of subsidiary inferences need to be made. We have expressed these in the form of *"FROM antecedent TO consequent"*. This is short-hand for saying that an act of inference is undertaken whenever one assumes that the consequent is true (or takes a particular value) upon the antecedent being true (or, respectively, having taken a corresponding value).

The list of subsidiary inferences is as follows:

1. **FROM Administered dose TO Delivered dose.** Animal bodies defend themselves against foreign substances. Their ability to do this may be impacted by the amount of the foreign substance ingested or to which the animal is exposed. For example, chemicals applied to nasal tissues are initially repelled by defences in the tissues themselves. Larger doses may destroy this first line of defence, thereby permitting proportionately more of the chemical to enter the body's circulatory pathways than would occur for smaller doses. In other words, the dose delivered to the target tissue or organ of the body may not be proportionate to the dose administered to the animal by the experimenter.

2. **FROM A sample of animals TO A population of the same species.** Reasoning from a sample to a population from which the sample is drawn in known as statistical inference.

[8] Indeed, the USA Environmental Protection Agency guidelines [15] permit one to claim probable human carcinogenicity from (sufficiently strong) animal evidence alone. Although such a claim would be classed in the second of two categories of "probable", it is still above "possible" human carcinogenicity.

3. **FROM A genetically uniform animal population TO A genetically more diverse population.** Animal subjects used in laboratory experiments are often closely related genetically, both in order to control for the impact of genetic diversity on responses and because, for reasons of convenience, subjects are used from readily-available sources. Consequently, the animal subjects used in bioassays are often not as diverse genetically as would be a wild population of the same species.

4. **FROM An animal population TO The human population.** This is perhaps the most contentious inference-step in carcinogenicity claims from bioassays. Animals differ from humans in their physiology and in their body chemistry, so it is not surprising that they also differ from us in reactions to potential carcinogens. Indeed, they differ from each other. According to Graham *et al.* [25, page 18], writing more than a decade ago, "Several hundred chemicals are known to be carcinogenic to laboratory animals, but direct evidence of their human carcinogenicity is either insufficient or nonexistent." Formaldehyde, for instance, was found to cause significant nasal cancers in rats but not in mice [25], while epidemiological studies of humans whose professions exposed them to high levels of the chemical found no significant increases in such cancers. Conversely — and perversely — epidemiological studies did reveal significant increases in brain cancers and leukaemias, for which there was no biologically-plausible explanation [25].

5. **FROM A site specificity in bioassay animals TO A possibly different site specificity in humans.** Most chemicals are pre-carcinogens which must be altered by the body's metabolic processes into an actively carcinogenic form. This happens differently in different species, because the body-chemistries are different or because the physiology or relative sizes of organs are different. Hence, a chemical may cause liver cancer in one animal species, but not in another species, or act elsewhere in another.

6. **FROM Localised exposure TO Broader exposure.** Bioassays adminster a chemical to a specific site in a specific way to the subject animals, as for example, in bioassays of formaldehyde applied to nasal passages to test for nasal cancer. In contrast, humans exposed to it may receive the chemical in a variety of ways. Morticians exposed to formaldehyde may receive it via breathing and by direct application to their skin, for example.

7. **FROM Large doses TO Small doses.** At typical levels of exposure, the incidences of most individual cancers in the general population are quite small, of the orders of a few percent or much less. At equivalent dose levels, then, bioassays will require very large sample sizes to detect statistically significant increases in cancer incidence. This would be prohibitively expensive, and so most bioassays administer doses considerably greater than the equivalent doses received (allowing for the relative sizes of the animal and human species) in the environment. In order to assert

carcinogenicity, then, a conversion model — a dose-response curve — is required to extrapolate back from large to small dose levels.

While one might expect the dose-response curve to slope upwards with increasing dose levels, this is not always the case. For example, high doses of a chemical may kill cells before they can become cancerous; or a chemical may be so potent that even low doses initiate cancer in all cells able to be so initiated, and thus higher doses have no further or a lesser effect. Indeed, if the chemical is believed to be mutagenic as well as carcinogenic, then even a single single molecule of the chemical should cause an effect. The issue of whether or not a threshold level for dose exists (below which no response would be observed) is a contentious one in most cases. Fueling controversy is the fact that claims of carcinogenicity can be very sensitive to the dose-response model used. Two theoretically-supported models for the risks associated with aflatoxin peanuts, for example, show human risk likelihood differing by a factor of 40,000 [54]. Similarly, the Chief Government Medical Officer of Great Britain recently admitted that the number of people eventually contracting CJD in Britain as a result of eating contaminated beef may be anywhere between a few hundred and several million [77]. For this reason, this inference is probably the most controversial aspect of carcinogenicity claims, after that of animal-to-human inference (Inference-Mode No. 4 above).

8. **FROM An animal dose-level TO A human equivalent.** The previous paragraph used the phrase "allowing for the relative sizes of the animal and human species". But how is this to be done? Is the dose extrapolated according to relative body weights of the two species (animal and human); or skin surface area (which may be appropriate for chemicals absorbed through the skin); or relative size of the organ affected? What is appropriate if different organs are affected in different species?

9. **FROM Administered doses TO Environmental exposure.** In order to expedite response times, bioassays may adminster the chemical in a manner different to that likely to be experienced by humans exposed to it in their environment. For example, the chemical may be fed via a tube directly into the stomach of the animal subject, which is unlikely to be the case naturally.

10. **FROM A limited number of doses TO Cumulative exposure.** Some chemicals may only produce adverse health effects after a lifetime of accumulated exposure. Body chemistry can be very subtle, and a small number of large doses of a chemical may have a very different impact from a much larger number of smaller doses, even when the total dose received is the same in each case.

11. **FROM A pure chemical substance TO A chemical compound.** Most chemicals to which people are exposed are compounds of several chemicals, not pure substances. Bioassay experiments, however, need to

be undertaken with pure substances, so as to eliminate any spurious causal effects. Consequently, a bioassay will not be able to assess any effects due to interactions between substances which occur in a real environment, including any transformations which take place inside the human body.

12. **FROM The human population TO Individual humans.** Individuals vary in their reactions to chemical stimuli, due to factors such as their genetic profiles, lifestyles, and personalities. Risks of carcinogenicity may be much higher or much lower than claimed for specific groups or individuals.

To claim human carcinogenicity on the basis of evidence from a bioassay thus depends on a number of different modes of inference, each of which must be valid for the claim to stand. We could write:

"The chemical $\mathcal{X}$ is carcinogenic to humans at dose d based on a bioassay of animal species a if:

- *There is a relationship between administered dose and delivered dose in the bioassay, AND*
- *The sample of animals used for the experiment was selected in a representative manner from the population of animals, AND*
- *The animal population from which the sample was drawn is as genetically diverse as the animal population as a whole, AND*
- *The specific animal physiology and chemistry relevant to the activity of $\mathcal{X}$ is sufficiently similar to human physiology and chemistry,"*

$\vdots$

and so on, through the remaining eight inference steps.

It is important to note that even if all modes of inference were valid in a particular case, our assertion could, strictly speaking, only validly be that the chemical $\mathcal{X}$ is associated with an increase in incidence of the particular cancer. The assertion ϕ does not articulate, nor could a bioassay or epidemiological study prove, a causal pathway from one to the other. There may, for example, be other causal factors leading both to the presence of the chemical in the particular environment and to the observed carcinogenicity.[9]

For the archetypical analysis above, we began with the assumption of just one bioassay being used as evidence to assert a claim for carcinogenicity. In

[9] Statistical analysis may be used to decide between alternative causal theories, even though it could not prove any one. For example, several causal explanations have been proposed for an above-average incidence of childhood leukaemia found in a location near a British nuclear power-plant, and statistical methods have been used to decide between these [12, 13]. As economist Milton Friedman once remarked: "the role of statistics is not to discover truth. The role of statistics is to resolve disagreements between people." (quoted in [45, page 4]).

reality, however, there is often evidence from more than one experiment and, if so, statistical meta-analysis may be appropriate [53]. This may involve pooling of results across different animal species, or across both animal and human species.[10] None of these tasks are straighforward, and will generally involve further modes of inference, which we do not explore in this paper.[11]

It is possible that working biomedical scientists and scientific risk assessors would consider the list above to be an example of extreme pedantry, and that many of these modes of inferences are no more than assumptions made in order to derive usable results. We have treated them as inference-modes so as to be quite clear about the reasoning processes involved. Our purpose in doing so is to make possible the automation of these processes, for which we use an argumention formalism (presented in Section 3). However, in the case of carcinogenicity of chemicals, elucidation of the modes of inference is important for science policy reasons also. This importance is illustrated in another domain of current concern about health risks, that of Genetically-Modified (GM) foodstuffs. In this domain, there are a number of biological and agricultural experiments planned or underway to assess the environmental and health impacts of such products. A sociologist of science, Brian Wynne, has argued persuasively [79] that these experiments are subject to assumptions which may invalidate them as a basis for scientific knowledge and for science policy. For example, that the duration and geographic scope of many GM experiments is not sufficient to observe all possible impacts; that the experiments use laboratory genetic specimens untypical of industrial outputs; that the wider environmental interaction effects are ignored in experiments focused on just one GM product or effect; etc. Although not expressed in terms of modes of inference, his argument is that such assumptions mean one is not justified in making inferences from the experimental results to the world beyond. By articulating these often-unstated assumptions as formal modes of inference, we clarify exactly what inferences are being made; this also better equips us, as Wynne argues, for a public debate on their consequences.[12]

2.2 The example of statistical inference

Only one of the forms of inference listed in the previous example is Statistical Inference, that is, reasoning about a population on the basis of evidence from a sample of that population. A formal logician, arguing strictly, would say that statistical inference is unsound: true antecedents are not guaranteed to generate

[10] The U.S.A. Environmental Protection Agency Guidelines [15] deal, at a high level, with the second issue.

[11] In addition, most chemical substances which adversely impact the body cause a number of effects, e.g.: cell mutation; malignant tumours; benign tumours; toxicity to cells; cell death; cell replication; suppression of the immune system; endocrine disturbances, etc. Some of these clearly interact — dead cells cannot then become cancerous, for instance — and the extent of interaction may be a non-linear function of the dose levels delivered. Simple claims about carcinogenicity often ignore these other effects and their interactions with the growth of malignant tumours ("carcinogenicity"). We do not deal with this issue here.

[12] Wynne also argues that public scepticism of scientific claims is justified in the light of science's failure to recognize the limitations of inference from unrealistic scientific experiments.

true consequents. However, the key achievement of mathematical statistics this century has been to place a bound on the extent of unsoundness: if we know the probability distribution of the variable of interest in the population, and that the mechanism which generated the sample was random (or, if not, the extent to which it is not), then we can estimate the probability that the inference from sample to population is incorrect. For example, we may conclude from particular functions of the sample values that there is a 95% chance that a certain interval contains the mean of the population.[13] This form of inference is still unsound (i.e. we still cannot guarantee the truth of a claim about a population parameter, given the truth of a claim about a sample parameter), but we now have an estimate of the upper bound on the extent of unsoundness. If we (as a society) make decisions based on the sample data, we still do not know which decisions are correct and which wrong, but we can estimate how many of the latter there will be at most. We are better off as a result.

The same would be true for the other modes of inference listed above. None of the modes listed is sound, in the sense of being able to be proven to guarantee the preservation of truth (from antecedent to consequent) in all circumstances. But, just as with statistical inference, if we were to have a quantitative bound on the extent of possible error in inference, then we would be better off than we are presently without it. Moreover, if such bounds existed for all the inferential modes listed, it may be possible to combine these bounds in an appropriate way, thereby generating a bound for the overall assertion of carcinogenicity from a bioassay. Estimating the soundness of each type of inference could be a matter of detailed examination of all the experimental and theoretical evidence (which may be a considerable undertaking) and then using this to develop a framework for theory development relevant to the mode of inference. Such theories would be analogous to the theories (e.g. Neyman-Pearson, Bayesian, Fiducial Probability) supporting the use of statistical inference. We do not pursue this idea further in this paper.

3 Argumentation Frameworks

3.1 Monodic and dialectical argumentation

An argument for a claim may be considered as a tentative proof for the claim. The philosopher Stephen Toulmin [68] proposed a generic framework for the structure of arguments which has been influential in the design of intelligent systems which use argumentation [19, 35, 75]. Our analysis, informed by Toulmin's structure, considers an argument to have the form of a proof, without necessarily its force.

Suppose ϕ is a statement that a certain chemical is carcinogenic at a specified level of exposure. Then an argument for ϕ is a finite, ordered sequence of

[13] There are different and contending views within statistical theory about the meaning of a statement such as this, a debate we do not enter. The interested reader is referred, for example, to [10, 63].

inferences $G_\phi = (\phi_0, \phi_1, \phi_2, \ldots, \phi_{n-1})$. Each sub-claim ϕ_i is related to the preceding sub-claim ϕ_{i-1} in the sequence as result of the application of an inference rule, R_i. These rules correspond to warrants in Toulmin's schema.[14] Note that R_i and R_j may be the same rule for i and j different. The modes of inference listed in Section 2.1 are examples of such rules. We may present this sequence graphically as follows:

$$\phi_0 \xrightarrow{R_1} \phi_1 \xrightarrow{R_2} \phi_2 \longrightarrow \ldots \longrightarrow \phi_{n-1} \xrightarrow{R_n} \phi.$$

If any of these rules were rules of inference generally considered valid in deductive logic (Modus Ponens, say), then we would be confident that truth would be preserved by use of the rule. In other words, using a valid rule of inference at step i means that whenever ϕ_{i-1} is true, so too is ϕ_i. If all the rules of inference were valid in this sense, then the argument G_ϕ would constitute a deductive proof of ϕ. The situations of interest to us, however, are when some or all of the inference rules are not valid in this sense, such as those of Section 2.[15]

In pure mathematics in general, once a theorem has been proven true, further proofs do not add to its truth, nor to the extent to which we are willing to believe the theorem to be true.[16] With arguments, however, alternative arguments may be of great interest. The greater the number of independent arguments that exist for a claim, the stronger is the case for it, and the stronger may be our belief in its truth. However, in arriving at a considered view as to our belief in the truth of a claim ϕ, we also need to consider the arguments against it, the arguments in favour of its negation $\neg\phi$ (which may not be the same thing), and any arguments which attack its supporting sub-claims, ϕ_i.[17]

[14] They are called step-warrants in Verheij's legal argumentation system [75].

[15] Note that our arguments in this Section could be formalised with use of valuation functions and models, as these terms are used within mathematical logic [34, 55], and we are attempting this in [43].

[16] However, even pure mathematicians may have variable belief in an assertion depending upon the means used to prove it. For example, constructivist mathematicians (e.g. [6, 70]) do not accept inference based on proof techniques which purport to demonstrate the existence of a mathematical object without also constructing it. Typically, such proofs use a *reductio ad absurdum* argument, showing that an assumption of non-existence of the object leads to a contradiction. Thus, constructivist mathematicians will seek an alternative proof for an assertion which a non-constructivist mathematician would accept as already proven. Although originally a contentious notion within pure mathematics, constructivist mathematics has obvious applications to computing, and has recently been proposed as a medium for the foundations of quantum physics [8]. Likewise, in another example, not all mathematicians accept the use of computers in proofs, or may do so only for some proofs. Computers have been used, for instance, to prove the Four Color Map Theorem [1] and to demonstrate the non-existence of projective planes of order 10 [38]. Verheij [73, note 9, pages 22-23] also remarks on these aspects of mathematical proof. For an interesting deconstruction of mathematical proofs as "objectively existing real things" see Appendix D of [23].

[17] Lakatos [37] proposed a model of successful mathematical discovery in which an attempted, but failing, proof of an assertion leads to successive re-formulations of the assertion and the proof, in response to undercutting counter-examples. This process continues until an assertion-proof pair is found in which the proof succeeds. Lakatos' model is similar to Naess' concept of "precizating" assertions [47], in which participants in a debate progressively elim-

Given these different arguments and counter-arguments, it is possible to define a symbolic calculus, called a Logic of Argumentation, which enables the combination ("flattening") of arguments for and against a proposition [20]. Since an argument is a tentative proof of a claim, our degree of belief in the claim will likely depend upon the argument advanced for it. Thus, for each pair (ϕ, G_ϕ) consisting of a claim and an argument for it, we can associate a measure α_ϕ of our strength of belief in ϕ given G_ϕ. We represent this as a triple $(\phi, G_\phi, \alpha_\phi)$, which we call an assessed argument.[18] The belief-indicator may be a quantitative measure, such as a probability, or an element from a qualitative dictionary, such as $\{$*Likely, Unlikely*$\}$. In either case, we can define algebraic operations on the set of belief-indicators (the "denotation dictionary for belief") enabling us to generate the degree of belief in a combined argument, when we know the degrees of belief of the subsidiary arguments. In addition to belief-indicators, one can also define other labels for claim-argument pairs, such as the values of world-states and the consequences of actions arising from the claim [20]. With such formal calculi, argumentation can be used in intelligent computer systems, and has been so used, particularly in the medical and legal domains (e.g. see [20, 36, 75]).

This view of argumentation presents the arguments as disembodied cases for and against a proposition. It is as if there were just one person in the debate, weighing the pros and cons with him or herself to arrive at a considered conclusion. Indeed, the carcinogenicity risk assessment guidelines of the U.S. Government Environmental Protection Agency [15], which are rules for the combination of evidence from disparate sources, have the appearance of an algorithm for the dispassionate weighing of evidence. We term this monodic (single-voice) argumentation.[19] However, in real life, there are usually many voices, each arguing for and against a proposition from differing perspectives, and sometimes with different views as to what constitutes acceptable rules of inference.

In seeking to model this rational cacophony, we have therefore turned to dialectical argumentation, a branch of philosophy dealing with the conduct of debate and discourse [60, 72]. One influential framework for dialectical argumentation has been that proposed by the German philosopher Jürgen Habermas [27]. Originally seeking to understand how ethical norms could be agreed between different people, and building on Toulmin's work [68], Habermas proposed a framework in which consenting members of a community can engage in a civil discourse. The key difference between monodic and dialectical argumentation is the presence in the latter of an audience. An audience needs to be persuaded, and may withhold her (or his or their) agreement to the claims being advanced

inate disagreement by qualifying their statements so as to make them more precise. It would be interesting to explore the relationship of these models to the argumentation approaches adopted in Artificial Intelligence, but we have not done this here.

[18] The use of "assessment" here is analogous to the concept of valuation in mathematical logic [55].

[19] Rehg [58], following Habermas [27], has referred to this as "monological" argumentation, a term we have not used because it may give the misleading impression of sound inference.

by a proponent. Indeed, members of an audience may advance counter-claims of their own, or rebuttals and undercutting arguments, or may question the premises or modes of inference used by a proponent. Habermas sought to identify rules under which such discourse could occur in a civil manner and so that all reasonable participants would feel satisfied with the process of discourse.[20] In [27], Habermas gave examples of the sort of rules his framework would include, for instance: "Different speakers may not use the same expression with different meanings" and "Everyone is allowed to question any assertion whatever".

Habermas has applied his framework to discourse in political, legal and social science arenas [28, 29]. Dialectical argumentation has also been applied by philosophers of science to the natural sciences. Marcello Pera [52], for example, models scientific discourse as a three-person dialogue, involving the scientific investigator, Nature and a sceptical scientific community. In Pera's model, the investigator proposes theoretical explanations of scientific phenomena and undertakes scientific experiments to test them. These experiments lead to "replies" from Nature in the form of experimental evidence. However, Nature's responses are not given in a direct or pure form, but are mediated through the third participant, the scientific community, which interprets the evidence, undertakes a debate as to its meaning and implications, and eventually decides in favour or against proposed theoretical explanations.[21] Drawing on the work of Pera and Habermas, William Rehg [59] proposes a form of dialectical argumentation for the debate which occurs within the scientific community (and between it and the experimenter), arising from Nature's responses to experiments. One of Rehg's aims is to capture the fact that even though the resolution of scientific questions may be objective and rational, as these terms are commonly understood, such resolution, by its nature as a human activity, takes place within a particular social, cultural and institutional context which invariably influences the course of resolution.[22]

3.2 A framework for dialogue

Motivated by these approaches, we seek to develop a dialectical argumentation framework for claims of carcinogenicity.[23] As does Habermas, we seek to codify rules of engagement. As do Pera and Rehg, we desire this to be a realistic model of scientific debate in the natural sciences, at least in our specific domain of chemical carcinogenicity. Our further aim, as was mentioned, is to encode

[20] We are here using "process" in its usual sense, not in Habermas' [27] specialized sense.

[21] An elegant over-simplification would be to think of the investigator's acts as locutions, Nature's responses as illocutions, and the scientific community's judgments as perlocutions, following Austin [2].

[22] As just one example, Jamieson [31] has argued that, in science policy debates over environmental and health risks, even uncertainty itself is, at least partly, socially constructed, with debate participants establishing, maintaining and using it to further particular agendas.

[23] Note that Verheij [75], building systems for legal appplications, uses the term "dialectical argument" to refer to a monodic argument which incorporates undercutting exceptions. We are using dialectical argumentation not in this sense, but to refer to a debate involving different views.

this framework in an intelligent computer system. This requires us to be explicit and comprehensive in the framework structure and rules we propose. Our work is still in progress, so we only suggest here the direction in which we are headed. We make no claims yet to comprehensiveness, and we have not studied the formal properties of the system we propose.

We define an "agora" (from the Greek for "meeting place") as a space in which the dialogue will occur, and we use this term also for the dialogue itself. Thus a "ϕ-agora" is a debate about the claim ϕ. A ϕ-agora consists of the following elements:[24]

- A database ∇ of well-formed formulae of a symbolic propositional language, with the usual connectives, in which atomic propositions are denoted ϕ_i.
- A set of different modes of inference, each denoted R_j.
- A set of debate participants, each denoted $\mathcal{P}_k$.
- A set of rules for asserting, supporting, questioning, denying, rebutting, undercutting, assumption-denying, mode-of-inference-denying of a claim. (i.e. which argument-moves are valid, when; which responses are valid, when.)
- A set of rules for summarising, combining and manipulating arguments.
- A presentation and advice module (so this can be presented to the user).

We see the agora framework presented here being embodied in an intelligent computer system which advises a human user or users. In our vision, the user will interact with the system in a number of ways, depending on his or her current purposes. For instance, the user may wish to explore both the arguments for and the against a particular claim, as when exploring the consequences of a particular action. Or, the user may wish to marshall together all the arguments for the claim, ignoring or rebutting any arguments against it.[25] The structure we propose is intended to allow for such variety of user purposes, with the system undertaking both automatic reasoning and argument mediation, in the terminology of Bart Verheij [74].[26]

Following Aristotle, Habermas [27] proposed a three-level structure for his dialectical argumentation framework. This structure ties in well with our previous work on argumentation.

[24] From the perspective of software functionality specification, our dialogue space is similar to the negotiation spaces of electronic commerce auction systems [40].

[25] This particular contrast in aims is similar to a distinction drawn by Naess [47]. He distinguished between arguments which were *pro-et-contra*, which considered the cases for and against a claim without reaching a conclusion, and arguments which were *pro-aut-contra*, which gave a conclusion and an overall, consistent, case for it.

[26] Subsequently, Verheij [75] re-defined argument-mediation systems as only those which assist more than one user. Our usage would include single-user applications, with the user constructing and manipulating multiple arguments for and against a claim.

Logical Level: At this level we seek to understand what are the logical implications of the knowledge base, to understand the "inner links" of an argument in Toulmin's [69] phrase. We need to do this because, in general, we do not know the consequences of our own knowledge, that is we do not know what we know. At this level, the system would be undertaking automated reasoning generating all the possible arguments that may be constructed from its knowledge in exactly the same way that the argumentation systems described in [20, 49] construct arguments.

Dialectical level: At the dialectical level, we are considering the cases for and against some proposition, the pros and the cons, and we need to be able to combine and flatten arguments generated by the first level in a way that is similar to that described earlier. This may be by simply looking at the strengths of the arguments [20, 49], or by looking at the relationships between them [51]. As indicated earlier, it is important for the carcinogenicity domain that the modes of inference themselves are able to be the subject of argument (for instance, being attacked or denied). Recent argumentation systems, such as those developed for legal applications by Verheij [73, 75], permit argument about inference rules.

Rhetorical level: At the rhetorical level, we are concerned with convincing an audience of a particular case. In terms of our system, we see this as a presentation layer, where the user can manipulate the activities of the other two layers for self-elucidation or for presentation to others. For example, this layer permits the user, within the permitted rules of the dialectical framework, to interrogate the arguments articulated by the system, to proposed rebuttals and undercutting arguments, etc. In other words, the user is permitted to act as a participant in the debate. Reed [57], for example, has explored the use of rhetoric for persuasion purposes in argumentation systems.[27]

Several argumentation theorists have proposed structures for dialectical argumentation which permit combination and resolutions of different arguments. These include the conflict resolution procedures of Rupert Crawshay-Williams [11], Arne Naess' rules of 'precizating' statements [47], the Dialectical Argumentation of Nicholas Rescher [60], the Dialogue Logic of Paul Lorenzen and his colleagues [41], Else Barth and Erik Krabbe's Formal Dialectics [4], and the Pragma-Dialectics of Frans van Eemeren and Rob Grootendorst [71]. We are currently examining these various approaches to see which, if any, is appropriate for application to design of intelligent risk systems. In addition, Verheij [75] argues that new kinds of user interfaces are required for argumentation systems, and he reports on several approaches in this area. This too will be an area requiring further investigation.

[27]Drawing on theories of legal mediation, Fuller [21, page 312] proposes three roles for someone seeking to assist conflicting parties to resolve a dispute in a knowledge domain: Facilitator; Negotiator; or Arbitrator. Our proposed system could certainly perform the first two of these roles, and may also support the third by, for example, identifying areas where scientific experiments may resolve disagreement.

4 Conclusions and Further Work

The Asian scholar Arthur Waley once wrote [76, Introduction, pages 96-7]:

"All argument consists in proceeding from the known to the unknown, in persuading people that the new thing you want them to think is not essentially different from or at any rate is not inconsistent with the old things they think already. This is the method of science, just as much as it is the method of rhetoric and poetry. But, as between science and forms of appeal such as poetry, there is a great difference in the nature of the link that joins the new to the old. Science shows that the new follows from the old according to the same principles that built up the old. 'If you don't accept what I now ask you to believe,' the scientist says, 'you have no right to go on believing what you believe already.' The link used by science is a logical one. Poetry and rhetoric are also concerned with bridging the gap between the new and the old; but they do not need to build a formal bridge. What they fling across the intervening space is a mere filament such as no sober foot would dare to tread. But it is not with the sober that poetry and eloquence have to deal. Their tê, *their essential power, consists in so intoxicating us that, endowed with the recklessness of drunken men, we dance across the chasm, hardly aware how we reached the other side."*

As engineers, our purpose in the work reported here is to erect the bridge enabling a computer system to reason its way across the chasm of carcinogenicity claims. To do this, we need to understand the way in which science itself crosses this chasm, and we have attempted this by explicating all the modes of inference deployed when bioassay evidence is used to support claims of carcinogenicity.

Claims of carcinogenicity are often contentious. Dialectical argumentation, as in the form proposed by Habermas to enable civil discourse on ethics, can be used as a framework for an automated debate on carcinogenicity. We have sketched, at a high-level, the outline of such a framework. Development of the detail needed to fully specify this framework is on-going work. In addition, we are exploring the links with modal logic [43] in order to formalize these ideas, and so provide means to apprehend the properties and behaviour of the resulting computer systems. This approach will, we believe, then enable us to incorporate and manipulate domain-specific qualitative information about the values of actions and their consequences, building on recent joint work done by one of us in extending logics of argumentation [20, 49].

In dealing with the truth of claims and the consequences of actions in the domain of chemical carcinogenicity, we are making a small step towards the goal, first proposed by Toulmin in 1958 [68], of explicating the actual forms and styles of acceptable arguments used in specific knowledge domains. We have considered just one domain, but our approach is clearly applicable more widely. Although our specific agenda here is the development of intelligent risk systems, the delineation of modes of inferences and the modeling of arguments used should benefit whichever is the community concerned with, and debating, the claim at issue. In the case of carcinogenicity of chemicals that community is all of us.

Acknowledgments

The work reported here was undertaken while the authors were with the Department of Electronic Engineering at Queen Mary and Westfield College, University of London. It was partially supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant GR/L84117 and a studentship, and this support is gratefully acknowledged. We also thank Heather Dickinson, Christoph Helma, Alessio Lomuscio, William Rehg and Bart Verheij for providing us with copies of their papers, and we are grateful for comments on an earlier draft received from Heather Dickinson, Susan Haack, William Rehg, Martijn Schut and Bart Verheij.

References

[1] K. Appel and W. Haken. Every planar map is four colorable. *Bulletin of the American Mathematics Society*, 82:711–712, 1976.

[2] J. L. Austin. *How To Do Things With Words*. Clarendon Press, Oxford, UK, 1962.

[3] A. Barnett and P. Wintour. Revealed: Science blunder that gave us BSE. *The Observer (London, UK)*, page 1, 8 August 1999.

[4] E. M. Barth and E. C. W. Krabbe. *From Axiom to Dialogue: A Philosophical Study of Logics and Argumentation*. Walter de Gruyter, Berlin, Germany, 1982.

[5] D. J. Benford and D. R. Tennant. Food chemical risk assessment. In D. R. Tennant, editor, *Food Chemical Risk Analysis*, pages 21–56. Blackie Academic and Professional, London, UK, 1997.

[6] E. Bishop. *Foundations of Constructive Analysis*. McGraw-Hill, New York City, NY, USA, 1967.

[7] C. P. Boyce. Comparison of approaches for developing distributions for carcinogenic slope factors. *Human and Ecological Risk Assessment*, 4 (2):527–577, 1998.

[8] D. S. Bridges. Can Constructive Mathematics be applied in physics? *Journal of Philosophical Logic*, 28:439–453, 1999.

[9] C. L. Broadhead, R. D. Combes, and M. Balls. Risk assessment: alternatives to animal testing. In D. R. Tennant, editor, *Food Chemical Risk Analysis*, pages 133–162. Blackie Academic and Professional, London, UK, 1997.

[10] D. R. Cox and D. V. Hinkley. *Theoretical Statistics.* Chapman and Hall, London, UK, 1974.

[11] R. Crawshay-Williams. *Methods and Criteria of Reasoning: An Inquiry into the Structure of Controversy.* Routledge and Kegan Paul, London, UK, 1957.

[12] H. O. Dickinson and L. Parker. Quantifying the effect of population mixing on childhood leukaemia risk: the Seascale cluster. *British Journal of Cancer*, 81:141–151, 1999.

[13] R. Doll. The Seascale cluster: a probable explanation. *British Journal of Cancer*, 81 (1):3–5, 1999.

[14] J. Emery, R. Walton, A. Coulson, D. Glasspool, S. Ziebland, and J. Fox. Computer support for recording and interpreting family histories of breast and ovarian cancer in primary care (RAGs): qualitative evaluation with simulated patients. *British Medical Journal*, 319:32–36, 1999.

[15] U.S.A. Environmental Protection Agency. Guidelines for carcinogen risk assessment. *U.S. Federal Register*, 51:33991–34003, 24 September 1986.

[16] H. J. Eysenck. *Smoking, Personality and Stress.* Springer Verlag, Berlin, Germany, 1991.

[17] J. Fox. Will it happen? Can it happen? A new approach to formal risk analysis. *Science and Public Affairs*, pages 45–49, Winter 1997.

[18] J. Fox. Will it happen? Can it happen? A new approach to formal risk analysis. *Risk, Decision and Policy*, 4, August 1999.

[19] J. Fox, P. Krause, and S. Ambler. Arguments, contradictions and practical reasoning. *Proceedings of the European Conference on Artificial Intelligence 1992, Vienna, Austria*, 1992.

[20] J. Fox and S. Parsons. Arguing about beliefs and actions. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, pages 266–302. Springer Verlag (Lecture Notes in Artificial Intelligence 1455), Berlin, Germany, 1998.

[21] S. Fuller. *Philosophy, Rhetoric, and the End of Knowledge: The Coming of Science and Technology Studies.* University of Wisconsin Press, Madison, WI, USA, 1993.

[22] M. A. Gibney, N. R. Jennings, N. J. Vriend, and J. M. Griffiths. Market-based call routing in telecommunications networks using adaptive pricing and real bidding. In *Proceedings of the 1999 Conference on Intelligent Agents for Telecommunications Applications*, Stockholm, Sweden, August 1999.

[23] J. Goguen. An introduction to algebraic semiotics, with application to user interface design. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy, and Agents*, pages 242–291. Springer Verlag (Lecture Notes in Artificial Intelligence 1562), Berlin, Germany, 1999.

[24] J. D. Graham. Historical perspective on risk assessment in the Federal government. *Toxicology*, 102:29–52, 1995.

[25] J. D. Graham, L. C. Green, and M. J. Roberts. *In Search of Safety: Chemicals and Cancer Risk.* Harvard University Press, Cambridge, MA, USA, 1988.

[26] J. D. Graham and L. Rhomberg. How risks are identified and assessed. *Annals of the American Academy of Political and Social Science*, 545:15–24, 1996.

[27] J. Habermas. *Moral Consciousness and Communicative Action.* MIT Press, Cambridge, MA, USA, 1991. *(Translation by C. Lenhardt and S. W. Nicholsen of: Moralbewusstsein un kommunikatives Handeln. Suhrkamp Verlag, Frankfurt-am-Main, Germany, 1983.).*

[28] J. Habermas. *Justification and Application.* MIT Press, Cambridge, MA, USA, 1993.

[29] J. Habermas. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy.* MIT Press, Cambridge, MA, USA, 1996 *(Translation by W. Rehg of: Faktizität und Geltung. Beiträge zur Diskurstheorie des Rechts und des demokratischen Rechtsstaats. Suhrkamp Verlag, Frankfurt-am-Main, Germany, 1992.).*

[30] C. Helma, S. Kramer, and B. Pfahringer. Carcinogenicity prediction for noncongeneric compounds: experiments with the machine learning program SRT and various sets of chemical descriptors. In *Proceedings of the Twelfth European Symposium on Quantitative Structure-Activity Relationships: Molecular modelling and prediction of bioactivity*, Copenhagen, Denmark, 1998.

[31] D. Jamieson. Scientific uncertainty and the political process. *Annals of the American Academy of Political and Social Science*, 545:35–43, 1996.

[32] P. N. Judson. Expert systems for hazard evaluation. In D. Tennant, editor, *Food Chemical Risk Analysis.* Chapman and Hall, London, UK, 1997.

[33] P. N. Judson and R. D. Coombes. Artificial intelligence systems for predicting toxicity. *Pesticide Outlook*, 7 (4):11–15, 1996.

[34] M. Kracht. *Tools and Techniques in Modal Logic.* North Holland, Amsterdam, The Netherlands, 1999.

[35] P. Krause, J. Fox, and P. Judson. An argumentation based approach to risk assessment. *IMA Journal of Mathematics Applied in Business and Industry*, 5:249–263, 1994.

[36] P. Krause, J. Fox, P. Judson, and M. Patel. Qualitative risk assessment fulfils a need. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, pages 138–156. Springer Verlag (Lecture Notes in Artificial Intelligence 1455), Berlin, Germany, 1998.

[37] I. Lakatos. *Proofs and Refutations: The Logic of Mathematical Discovery.* Cambridge University Press, Cambridge, UK, 1976.

[38] C. W. H. Lam, S. Swiercz, and L. Thiel. The non-existence of finite projective planes of order 10. *Canadian Journal of Mathematics*, 41:1117–1123, 1989.

[39] D. V. Lindley. *Making Decisions.* John Wiley and Sons, London, UK, second edition, 1985.

[40] A. Lomuscio, M. Wooldridge, and N. Jennings. Automated negotiation in agent-mediated electronic commerce. Unpublished manuscript, Department of Electronic Engineering, Queen Mary and Westfield College, University of London, London, UK, 1999.

[41] P. Lorenzen and K. Lorenz. *Dialogische Logik.* Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1978.

[42] D. P. Lovell and G. Thomas. Quantitative risk assessment. In D. R. Tennant, editor, *Food Chemical Risk Analysis*, pages 57–86. Blackie Academic and Professional, London, UK, 1997.

[43] P. McBurney and S. Parsons. Modal argumentation. Technical report, Department of Electronic Engineering, Queen Mary and Westfield College, University of London, London, UK, 1999 *(In preparation).*

[44] D. N. McCloskey. *The Rhetoric of Economics.* Wheatsheaf Books, Brighton, UK, 1985.

[45] D. N. McCloskey. *Knowledge and Persuasion in Economics.* Cambridge University Press, Cambridge, UK, 1994.

[46] S. H. Moolgavkar. Stochastic models of carcinogenesis. In C. R. Rao and R. Chakraborty, editors, *Handbook of Statistics, Volume 8: Statistical Methods in Biological and Medical Sciences*, pages 373–393. North-Holland, Amsterdam, The Netherlands, 1991.

[47] A. Naess. *Communication and Argument: Elements of Applied Semantics.* Allen and Unwin, London, UK, 1966. *(Translation of: En del Elementaere Logiske Emner. Universitetsforlaget, Oslo, Norway, 1947.).*

[48] T. Page. A generic view of toxic chemicals and similar risks. *Ecology Law Quarterly*, 7 (2):207–244, 1978.

[49] S. Parsons and S. Green. Argumentation and qualitative decision making. In A. Hunter and S. Parsons, editors, *Proceedings of the 5th European Conference on Quantitative and Symbolic Approaches to Reasoning and Uncertainty (ECSQARU99)*, pages 328–339. Springer Verlag (Lecture Notes in Artificial Intelligence 1638), Berlin, Germany, 1999.

[50] S. Parsons and P. McBurney. Intelligent systems which reason about risk. *Journal of Long Range Planning*, 1999 *(Submitted for publication)*.

[51] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261—292, 1998.

[52] M. Pera. *The Discourses of Science.* University of Chicago Press, Chicago, IL, USA, 1994.

[53] D. B. Petitti. *Meta-Analysis, Decision Analysis and Cost-Effectiveness Analysis: Methods for Quantitative Synthesis in Medicine.* Oxford University Press, Oxford, UK, 1994.

[54] R. A. Pollak. Government risk regulation. *Annals of the American Academy of Political and Social Science*, 545:25–34, 1996.

[55] S. Popkorn. *First Steps in Modal Logic.* Cambridge University Press, Cambridge, UK, 1994.

[56] H. Raiffa. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty.* McGraw-Hill, New York, 1968.

[57] C. Reed. *Generating Arguments in Natural Language.* PhD thesis, University College, University of London, London, UK, 1998.

[58] W. Rehg. Reason and rhetoric in Habermas's Theory of Argumentation. In W. Jost and M. J. Hyde, editors, *Rhetoric and Hermeneutics in Our Time: A Reader*, pages 358–377. Yale University Press, New Haven, CN, USA, 1997.

[59] W. Rehg. Argumentation theory and the philosophy of science since Kuhn: the rationality of theoretical advance. Paper presented to Seminaret i Vitenskapsteori, Universitetet i Oslo, Oslo, Norway, August 1998.

[60] N. Rescher. *Dialectics: A Controversy-oriented approach to the Theory of Knowledge.* State University of New York Press, Albany, NY, USA, 1977.

[61] L. R. Rhomberg. A survey of methods for chemical health risk assessment among Federal regulatory agencies. *Human and Ecological Risk Assessment*, 3 (6):1029–1196, 1997.

[62] G. I. Seffers. Crisis system for President: DARPA's Project Genoa would speed U.S. decision-making. *Defense News*, 13 (13):1, 30 March, 1998.

[63] T. Seidenfeld. *Philosophical Problems of Statistical Inference.* Reidel, Boston, MA, USA, 1979.

[64] M. E. Shere. The myth of meaningful environmental risk assessment. *Harvard Environmental Law Review*, 19 (2):409–492, 1995.

[65] A. Srinivasan, S. H. Muggleton, M. J. E. Sternberg, and R. D. King. Theories of mutagenicity: a study in first-order and feature-based induction. *Artificial Intelligence*, 85:277–299, 1995.

[66] H. Teff and C. R. Munro. *Thalidomide: The Legal Aftermath.* Saxon House, Westmead, Farnborough, Hampshire, UK, 1976.

[67] C. A. G. Tonnelier, J. Fox, P. N. Judson, P. J. Krause, N. Pappas, and M. Patel. Representation of chemical structures in knowledge-based systems: The StAR system. *Journal of Chemical Information and Computer Sciences*, 37:117–123, 1997.

[68] S. E. Toulmin. *The Uses of Argument.* Cambridge University Press, Cambridge, UK, 1958.

[69] S. E. Toulmin. Logic, rhetoric and reason: redressing the balance. In F. H. van Eeemeren, R. Grootendorst, J. A. Blair, and C. A. Willard, editors, *Argumentation Illuminated*, pages 3–11. SICSAT, Amsterdam, The Netherlands, 1992.

[70] A. S. Troelstra and D. van Dalen. *Constructivism in Mathematics: An Introduction (Two Volumes).* North-Holland, Amsterdam, The Netherlands, 1988.

[71] F. H. van Eemeren and R. Grootendorst. *Argumentation, Communication and Fallacies: A Pragma-Dialectical Perspective.* Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1992.

[72] F. H. van Eemeren, R. Grootendorst, and F. S. Henkemans. *Fundamentals of Argumentation Theory.* Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.

[73] B. Verheij. *Rules, reasons, arguments: Formal studies of argumentation and defeat.* PhD thesis, Universiteit Maastricht, Maastricht, The Netherlands, 1996.

[74] B. Verheij. Argue! - an implemented system for computer-mediated defeasible argumentation. In H. La Porte and J. van den Herik, editors, *Proceedings of the Tenth Netherlands-Belgium Conference on Artificial Intelligence*, pages 57–66, Amsterdam, The Netherlands, 1998. CWI.

[75] B. Verheij. Automated argument assistance for lawyers. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law, Oslo, Norway*, pages 43–52, New York City, NY, USA, 1999. ACM.

[76] A. Waley. *The Way and Its Power: A Study of the Tao Tê Ching and Its Place in Chinese Thought.* George Allen and Unwin, London, UK, 1934.

[77] N. Watt. Millions still at risk from CJD. *The Guardian (London, UK)*, page 1, 22 September 1999.

[78] B. Wynne, editor. *Risk Management and Hazardous Waste: Implementation and the Dialectics of Credibility.* Springer Verlag, Berlin, Germany, 1987.

[79] B. Wynne. Bitter fruits. *The Guardian (London, UK)*, page S2, 16 September 1999.