Intelligent Systems to Support Deliberative Democracy in Environmental Regulation

Peter McBurney and Simon Parsons
Department of Computer Science, University of Liverpool
Chadwick Building, Liverpool L69 7ZF U.K.
{P.J.McBurney,S.D.Parsons}@csc.liv.ac.uk

Abstract

Among normative models for democracy, the Deliberative model suggests that public policy decisions should be made only following rational, public deliberation of alternative courses of action. We argue that such a model is particularly appropriate for the assessment of environmental and health risks of new substances and technologies, and for the development of appropriate regulatory responses. To give operational effect to these ideas, we propose a dialectical argumentation formalism for an intelligent system within which deliberative debates about risk and regulation can be conducted. Our formalism draws on various philosophies of argumentation, scientific and moral discourse, and communicative action, due to Toulmin, Pera, Alexy and Habermas.

1 Environmental Regulation

New technologies and substances have the potential to cause adverse and unanticipated effects, on people, on other living species, and on our environment, and these effects are increasingly global in scale. Because of such risks, most chemicals and many technological innovations now require Government regulatory approvals before they can be used or sold commercially. In developed countries, there are now typically a number of Governmental and quasi-Governmental agencies tasked with assessing the risks of new substances and technologies and of framing appropriate regulatory responses for those perceived to have high adverse risks.

Framing and implementing legal regulations for such new substances is usually very difficult. Firstly, identification of all the potential consequences of a new substance can be problematic. Although thalidomide, for example, was safely tested on both animals and humans before its commercial release, none of the experimental subjects were pregnant (Teff & Munro, 1976), presumably because the possibility of differential harm was not considered. Recently, researchers in Artificial Intelligence (AI) have proposed the use of argumentation-based procedures for this problem of possibilistic risk assessment (Krause *et al.*, 1998).

Secondly, even when the possible consequences of new substances are believed known, assessment and quantification of risks is often problematic and invariably subjective (USA EPA, 1996; Shere, 1995; Rhomberg, 1997; Toll, 1999). In many cases, the scientific evidence upon which assessment is to be based is not conclusive and the theoretical issues involved are contentious, even among scientists working in the same field. The chemical formaldehyde, for instance, was found to cause statistically-significant increases in nasal cancers in rats but not in mice, while epidemiological studies of humans whose professions exposed them to high levels of the chemical found no significant increases in such cancers (Graham, Green, & Roberts, 1988). Moreover, quantification of risk normally requires the adoption of a mathematical model linking responses to exposure levels. Different dose-response models can result in widely different assessments of risk. Two theoretically-supported models for the risks associated with aflatoxin peanuts, for example, showed human risk likelihood differing by a factor of 40,000 (Pollak, 1996).

Another major issue for environmental and health regulation of new substances is that the consequences of different regulatory options may be very different. Page (1978) noted that different groups of people may benefit or lose from regulation or non-regulation of a substance, and that their gains or losses may be very different in scope, magnitude, intensity, timing and duration. A new chemical substance, for instance, wrongly deemed by regulators to be safe and then used, may adversely impact millions of people. How does one quantify the subsequent misery or loss of life? Conversely, the same chemical wrongly deemed to be unsafe, and so never sold, may cause large financial losses to the companies which undertook the initial research. Moreover, not using the chemical, when it would be safe to do so, may adversely impact those who could benefit from its use; these people too may number in their millions and their (unrelieved) misery may also be great.

Ultimately, regulation of new substances will always involve a trade-off of alternative regulatory options, weighing the anticipated consequences of each. As mentioned,

quantification of consequences is difficult. Yet, even without quantification of consequences, different values may be assigned to different consequences: society may, for instance, prefer to forego the sunk costs of commercial development of new chemicals rather than risk the loss of life resulting from their use. It is likely that different people and groups within society will have different value-assignments in such trade-offs. As an example, Stirling & Mayer (1999) assessed the different decision-criteria and value-assignments of a sample of British experts in the current debate over Genetically-Modified Organisms (GMOs). At present, however, there is no formal mechanism for articulating and comparing these valuations in most risk regulation; instead, different interest groups make their arguments and values known to regulators through lobbying and to the public through public relations activity. Keeney (1996) has argued for values to be assigned explicitly, so that hidden agendas may be exposed, and fairer and more transparent trade-offs undertaken.

By articulating and comparing the values of stakeholders, the multi-criteria scoring technique of Stirling & Mayer (1999) could be used in the formulation of regulatory policy. However, because this method is essentially quantitative and reductive, it does not incorporate the arguments for (or against) the criteria selected or the scores assigned. Our approach involves defining an intelligent system capable of qualitative representation and manipulation of arguments and values in the form of a dialogue game, and we have termed such a system a *Risk Agora* (McBurney & Parsons, 2000). Another approach, very similar in spirit to this, is the Zeno argumentation framework of Gordon & Karacapilidis (1997). This system was developed for conflict resolution and mediation in urban planning regulation, and uses an argumentation formalism based on the schema of Toulmin (1958) and the IBIS model of Rittel & Webber (1973). Our structure differs from Zeno not only in our intended application domain but also in our use of a specific philosophy of science to represent the community's understanding of scientific knowledge. This understanding may be distinct from that of any one participant, and needs to be formally represented. Another difference between our approach and that of Zeno is our deployment of dialogue moves based on speech acts specific to the domain of risk regulation, rather than more generic moves. Were environmental regulators to adopt an Agora framework for the development of risk regulatory policy, we believe this would give greater effect to a deliberative model of democracy, a concept explained in the next Section. Section 3 considers Deliberative procedures in the specific domain of risk regulation, and Section 4 outlines our formal structure for the Agora. For a recent review of the philosophical and computational issues arising in argumentation-based decision support systems see Girle *et al.* (2001).

2 Normative Models of Democracy

The term "Deliberative Democracy" was first introduced by Bessette (1980), and the subject has been the focus of much recent attention by philosophers of politics and law (Bohman & Rehg, 1997). It refers to a particular notion of democracy, one of several developed by philosophers as normative models of democratic politics.[1] To explain these, we begin with an abstract model of a democracy as consisting of just two entities: Society and the State. Society is the set of individuals, organizations and companies, together with the panoply of relationships between them, while the State is the apparatus of public-sector administration. The key normative question for democracy, then, is: *What should be the process of formation of political will?*, or *How should Society program the State?*

One model argues that a democracy is best governed when such programming is the task of a technocratic elite (who may be elected), making decisions on behalf of the general public. By contrast, rational-choice or liberal models view the process of political-will formation as akin to the workings of an economic market. In this model, political parties and interest groups act as entrepreneurs, offering alternative "products" in the form of bundles of state-instructions (or equivalently, philosophies of bundle- formation), to voters who then "purchase" their preferred bundle. That bundle with the greatest "market-share" — in the form of votes — becomes the set of instructions used to program the state.

The rational-choice model views citizens as consumers, acting in their own perceived individual self-interest and negotiating bargained compromises to political questions. Their preferences may be predetermined, uninfluenced by the process of choosing between alternative programmes. By contrast, deliberative notions of democracy see the political process as more than this, with citizens undertaking a substantial process of public deliberation to decide political questions. In this model, political-will formation is a process by which collective decisions regarding practical questions are made on the basis of rational and public reflection of the arguments for and against different courses of action. Such deliberation may well lead to the participants to change their preferences and their value assignments. As Michelman (1989) defined it:

"Deliberation . . . refers to a certain attitude toward social cooperation, namely, that of openness to persuasion by reasons referring to the claims of others as well as one's own. The deliberative medium is a good faith exchange of views — including participants' reports of their own understanding of their respective vital interests — . . . in which a vote, if any vote is taken, represents a pooling of judgments."

The benefits claimed for deliberative approaches to democracy include the legitimacy which public participation provides to political decision-making. People are

[1] Our presentation here draws primarily from Christiano (1997) and Chapter 9 of Habermas (1998).

generally more willing to accept decisions which they have had a role in forming, even when they disagree with the outcomes of those decisions. Moreover, the very act of participation may indicate, or may induce or strengthen, a concern for the welfare of the community beyond mere individual self-interest. In addition, a society with deliberative procedures may treat its citizens with more respect than it would with elitist or rational-choice procedures.

3 Deliberation in Risk Regulation

We believe that the specific characteristics of environmental risk regulation lead to further benefits from the adoption of public deliberative approaches. Firstly, the likelihood of identifying all possible consequences of new substances and technologies is increased the greater the number of participants engaged in considering the problem. Secondly, the inconclusive nature of much of the science involved and the subjective nature of risk assessments together mean that broad debate is invaluable. In this way, assumptions can be tested, experiments replicated, and inferences subject to detailed scrutiny. Many conclusions may fall when so subjected. For instance, Wynne (1996) argues that scientific experts often possess a generalized form of knowledge, which may not always be valid in particular circumstances. An example of this is shown by the case of alleged deforestation in West Africa (Fairhead & Leach, 1998), where the traditional inhabitants have understood local environmental reality better than western scientists.[2]

Thirdly, the complexity of most important risk assessments requires contributions from a wide spectrum of expertise. For instance, a rational decision regarding the regulation of Genetically-Modified Organisms (GMOs) arguably requires expertise in: molecular biology, genetics, plant biology, entomology, ecology, medicine, agriculture, economics, statistical experiment design, statistical inference, marketing, international trade, international development and international law. In such circumstances, as Willard (1990) has argued, no one person has the breadth of expertise required, and so even experts must perforce accept arguments based on authority from outside their own domain.[3] In this context, deliberative procedures can ensure assumptions from different disciplines are all valid and consistent, and explore cross-disciplinary interactions and conclusions. Finally, deliberative procedures can ensure the articulation of the consequences associated with different regulatory alternatives, and of the different values different people and groups may place on these. For all these reasons, we believe that deliberative procedures should ensure better quality decision-outcomes in environmental regulation than can either rational-choice or elite procedures.

How can deliberative procedures be implemented in a large, modern society, with possibly millions of citizens, thousands of elected representatives and perhaps thousands of environmental decisions in train at any time? Traditional means of public consultation, such as referenda and town-hall meetings, are arguably better suited to infrequent deliberations or to small communities. Citizens' Panels, where a small group of people, acting like a jury, hear evidence from relevant experts on an issue and then determine a course of action, have been used in some recent scientific policy debates.[4] Such panels, by exposing the participants to a diversity of expert opinion and forcing a decision to be made, are undoubtedly valuable for those involved; however, not everyone can participate. The Internet has been seen by some commentators (e.g. Ess (1996)) as a means to enable greater democratic participation in public policy decision-making, both through the wider availability of information and through electronic voting systems. We believe that a Risk Agora, suitably instantiated, could represent the scientific and political uncertainty involved in an environmental risk assessment and potentially bring the benefits of participation in a citizens panel to a wider audience. The next Section outlines the formal structure of our system.

4 Agora Formalization

In order that an intelligent system is able to be used for regulation decisions regarding a new substance (or technology), it will need to be able to do the following:

- Represent possible risks of deployment or non- deployment of the substance.
- Represent scientific uncertainty over the possibility, causal mechanisms, magnitude, duration and scope of potential risks.
- Represent different regulatory options for the substance and their potential consequences of enacting or not enacting these.
- Represent values assigned to such consequences by different individuals or groups.
- Enable the questioning, contestation, defence and qualification of each of the above types of statements.

[2] In coastal ecology, the importance of such local, "indigenous", knowledge for the selection of appropriate coastal-zone management policies has been recognized by the development of the *SimCoast* expert system by the U.K. Centre for Coastal Management (`www.ccms.ac.uk/simcoast.htm`). This has been used in several countries to aid coastal policy-making, by representing and integrating both scientific and traditional knowledges of local ecology. However, despite its incorporation of indigenous knowledge, representation in *SimCoast* is undertaken within a western scientific ontology, and assumes both that integration of different knowledges is achievable and that it is possible through discussion.

[3] U.S. company Applied Biomathematics is developing software to assess the validity of arguments across different disciplines in ecological risk assessments. See: `http://www.ramas.com/interest.htm#validity`.

[4] See Bhattachary (1998) for further description of such approaches.

- Enable the coherent comparison and manipulation of arguments for and against particular statements.
- Enable the synthesis of arguments into an overall case for a particular statement.
- Enable the taking of summary "snapshots" of a debate at any time.
- Enable the selection (and hence, legal imposition) of a particular regulatory option.

In (McBurney & Parsons, 2000), we developed a formal dialectical argumentation syntax, using a propositional language, for representing arguments over scientific claims of chemical carcinogenicity. This formalism drew upon, firstly, the philosophy of science of Pera (1994), which views scientific activity as a three-person dialogue game between a scientific investigator, Nature and a skeptical scientific community. Secondly, our structure drew upon rules for discourse in the philosophy of Discourse Ethics of Habermas (1991) and Alexy (1990), developed as a normative model for rational debate between reasonable and consenting participants. Although proposed initially for debates over moral questions, the theory has since been applied to legal and political domains, as in Habermas (1996). Thirdly, our formalism used argumentation schema of Toulmin (1958), within a dialectical framework, to enable the presentation of arguments for and against scientific claims. In other words, participants in the debate could variously posit, assert, contest, justify, rebut, undercut, qualify and retract claims, just as happens in real scientific discourse.

Moreover, by the use of dictionaries of uncertainty labels, our formalism permitted the assertion of individual degrees of belief in claims, their supporting evidences, their modes of inference, and their consequences. Participants could, for example, accept a scientific claim but label it as, say, *Plausible*, rather than as *Confirmed*. As an example of our formalism, a debate participant $\mathcal{P}_i$ could demonstrate her argument $\mathcal{A}(\rightarrow \theta)$ supporting a claim θ, an argument to which she was committed with strength D, by making the dialogue move:

show_arg$(\mathcal{P}_i : \mathcal{A}(\rightarrow \theta, D))$.

By use of such dictionaries (which could be quantitative, e.g. probability estimates), degrees of commitment and uncertainty by individual participants can be represented. Our formalization also included truth-valuation functions which assigned degrees of certainty to statements on behalf of the community as a whole, based upon the existence and strength of arguments for and against the statement and its supporting grounds. Such truth-valuation functions effectively produce an on-going representation of the dialogue community's changing views of a scientific claim, and so provide the desired "snapshot" capability.

However, as the discussion in this paper has demonstrated, in a regulatory context there are other types of statements besides scientific claims: statements of values and preferences, moral obligations and relationships, and imperatives (i.e. regulations). Therefore, in this paper, we extend our earlier formalism by incorporating expressions for these additional types of statements. We do this by drawing on other work of Habermas, his philosophy of Communicative Action (Habermas, 1984), in which he sought to understand how people collaborate rationally to achieve a common understanding of a situation or a collective action. As part of his philosophy, he proposed a typology of statements,[5] which we have adapted and relabelled for the specific context of environmental regulation. (Habermas' labels are in parantheses.)

Factual Statements (Constative Speech Acts): These are statements which seek to represent the state of the external world. In our domain, such statements include claims about scientific reality, and the scientific, economic or social consequences of particular actions. In the Agora formalism, we demarcate these different types of factual statements for clarity of exposition. Contesting such a statement means denying that it is a true description of objective, external reality.

Value Statements (Expressive Acts): These are statements which seek to represent the state of the speaker's internal world, i.e. they reveal publicly a subjective preference or value assignments. Such statements may only be contested by doubting the sincerity of the speaker.

Connection Statements (Regulative Acts): These are statements which assert some relationship between different parties, in the common world of the Agora participants. One may assert, for example, that the group of stakeholders with an interest in the regulation of a proposed new technology is broader than previously defined.

Inferential Statements (Operative Acts): These are statements which refer to the content of earlier statements made in a debate, drawing inferences from them or assessing implications. Once a scientific theory has been proposed, a specific risk assessment model and the ensuing calculations based on this theory fall into this category. Contestation of such statements can take the form of questioning the appropriateness or the validity of the inferences made.[6]

Procedural Statements (Communicative Acts): These are statements about the activity of speaking itself, such as the rules for participation and debate. In many real-life discourses, these often become the focus of debate, overtaking issues of substance.[7]

[5] He was building on the typology of Searle (1979).

[6] Our definition departs slightly from that of Habermas, in that our Inferential Statements may have "genuine communicative intent."

[7] For instance, in the scientific debate over GMOs in Britain during

Obligation Statements (Imperative Acts): These are statements which assert some obligation on the participants, for example, that they must limit the commercial sale of a new substance. Only the authorized regulator has the power to make such assertions, and once made, cannot be contested within the Agora. (In real-life, they may of course be contested in the courts.)

Given this typology of statements possible within the Agora, we can define a syntax of dialogue moves, extending the syntax for scientific reasoning of McBurney & Parsons (2000). Thus, for instance, a debate participant $\mathcal{P}_i$ could state her value assignment D_C to consequence C of action θ by means of the dialogue move:

show_value$(\mathcal{P}_i : \mathcal{V}al(\theta \rightarrow C, D_C))$.

Debate over such statements will then proceed according to the same rules for positing, proposing, contesting, qualifying, etc, statements as defined in our earlier paper.

The dialectical argumentation formalism we have presented here is related to other recent work in which we have applied argumentation in the design of intelligent systems. For example, in Fox & Parsons (1998); Parsons & Green (1999), we developed formalisms for the articulation and manipulation of statements of qualitative value, as part of calculi for qualitative decision-making. In both these papers, the argumentation formalism presented was monolectical, whereas in Amgoud, Maudet, & Parsons (2000), we presented a formalism for dialectical argumentation, involving two participants engaged in a generic debate.

5 Example

In this section, we illustrate our approach with a simplified example drawing on recent debates over Genetically Modified Organisms (Stirling & Mayer, 1999; UK ESRC Global Environmental Change Programme, 1999). We assume a debate with willing and reasonable participants denoted $\mathcal{P}_1, \ldots, \mathcal{P}_6$. For ease of understanding, we articulate the dialogue moves in plain English (not in the formal syntax), and we label each move with its type.

M1 (Factual): $\mathcal{P}_1$ asserts that foods containing GMOs may not be safe to eat.

M2 (Query): $\mathcal{P}_2$ asks $\mathcal{P}_1$ for an argument supporting Claim M1.

M3 (Factual): $\mathcal{P}_1$ presents evidence of experiments in which rodents fed GM potatoes had significantly greater tumors than a control group.

M4 (Value): $\mathcal{P}_3$ accepts the assertion M1 of $\mathcal{P}_1$, and places a large negative value on GMOs being in the human food chain.

M5 (Inferential): $\mathcal{P}_2$ asserts that assertion M1 of $\mathcal{P}_1$ only follows from argument M3 if humans and rodents are sufficiently alike biochemically.

M6 (Value): $\mathcal{P}_4$ places zero value on GMOs being in the human food chain, provided they are labeled whenever present.

M7 (Connection): $\mathcal{P}_5$ says asserts that agriculture and food distribution companies have a duty to inform consumers of the presence of GMOs in their food products.

M8 (Inference): $\mathcal{P}_2$ asserts that labeling of foods containing GMOs (as suggested in moves M6 of $\mathcal{P}_2$and M7 of $\mathcal{P}_5$) will lead to consumers rejecting such foods in favor of non-GM foods, even if they pose no dangers.

M9 (Factual): $\mathcal{P}_6$ asserts that GM foods have the potential to end hunger in the Third World.

M10 (Query): $\mathcal{P}_5$ asks $\mathcal{P}_6$ for evidence for her assertion M9.

M11 (Factual): $\mathcal{P}_6$ provides evidence of the increased yields available from GM crops.

M12 (Contestation): $\mathcal{P}_3$ contests assertion M9 of $\mathcal{P}_6$, on the grounds that the high investments required for development of GM Foods will lead to increasing concentration of corporate ownership in the agrichemicals sector, and this will increase poverty in the developing world.

M13 (Contestation): $\mathcal{P}_6$ defends assertion M9 against the attack of $\mathcal{P}_3$ by contesting the claim that increased concentration of corporate ownership will necessarily increase poverty in the developing world.

M14 (Procedural): $\mathcal{P}_5$ asserts that this debate should be limited to a discussion of the consequences of permitting GMOs only in our own country, and not elsewhere.

M15 (Value): $\mathcal{P}_4$ asserts that she would place greater value on a regulatory decision made with awareness of the global consequences than one informed only by local consequences.

. . . and so on.

As this very simplified example shows, the scientific, economic and social issues involved may be quite complex, and the values placed on outcomes by different participants possibly very discordant. There is no guarantee of resolution of such differences, as Stirling & Mayer

1999, an argument between the medical journal *The Lancet* and The Royal Society ensued over whether the latter was entitled to comment on a paper submitted to the journal while it was still under consideration for publication (Lancet Editorial, 29 May 1999).

(1999) found in their application of multi-criteria scoring to the same issue of GM foods. However, even without a guarantee of resolution, representation of a debate within such a formalism, forces greater clarity in the statements articulated, and this will surely facilitate any attempt at trade-offs between different alternatives.

6 Discussion

In this paper, we have proposed a novel dialectical argumentation formalism for an intelligent system within which deliberative debates about possible environmental risks and regulatory alternatives can be conducted. As we have explained, our formalism draws on Toulmin's theory of argumentation, Pera's philosophy of science, the Discourse Ethics of Habermas and Alexy, and Habermas' theory of Communicative Action. We see this approach as potentially giving practical effect to notions of Deliberative Democracy, enabling rational, public and transparent consideration of decision alternatives prior to deciding on a course of action. As well as effecting deliberative notions of democracy, this approach, we have argued, provides particular benefit in the domain of regulation of environmental and health risks. This is due, firstly, to the typically difficult, subjective and contested nature of risk assessment and the science on which it is based. Secondly, it is because regulatory alternatives may impact different groups with markedly different consequential outcomes, outcomes to which individuals and groups may assign very different values. Once instantiated with the details of a specific risk debate, the Risk Agora could be used in a number of ways:

1. To understand the logical implications of the scientific knowledge relating to the particular issue, and the arguments concerning the consequences and value-assignments of alternative regulatory options.

2. To consider the various arguments for and against a particular claim (including regulatory options), how these arguments relate to each other, their respective degrees of certainty, and their relative strengths and weaknesses.

3. To develop an overall case for a claim, combining all the arguments for it and against it.

4. To enable interested members of the public to gain an overview of the debate on an issue.

5. To support group deliberation on the issue, for example in Citizens Panels.

6. To support risk assessment and regulatory determination by government regulatory agencies.

As the last three in this list demonstrate, the Risk Agora potentially gives effect to the ideals of reasoned, public decision-making, and thus supports notions of deliberative democracy. We believe the nature of decisions involved in the assessment and regulation of risk mean that the adoption of such processes may improve the quality and fairness of decisions made in this domain.[8]

References

Alexy, R. 1990. A theory of practical discourse. In Benhabib, S., and Dallmayr, F., eds., *The Communicative Ethics Controversy*. Cambridge, MA, USA: MIT Press. 151–190. (Published in German 1978).

Amgoud, L.; Maudet, N.; and Parsons, S. 2000. Modelling dialogues using argumentation. In Durfee, E., ed., *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, 31–38. Boston, MA, USA: IEEE Press.

Bessette, J. 1980. Deliberative Democracy: The majority principle in republican government. In Goldwin, R. A., and Schambra, W. A., eds., *How Democratic is the Constitution*. Washington, DC, USA: AEI. 102–116.

Bhattachary, D. 1998. *Enviromental values: a critical appraisal of cost-benefit analysis and discursive forms of decision-making*. Report, Oxford Centre for Environment, Ethics and Society, Mansfield College, Oxford, UK.

Bohman, J., and Rehg, W., eds. 1997. *Deliberative Democracy: Essays on Reason and Politics*. Cambridge, MA, USA: MIT Press.

Christiano, T. 1997. The significance of public deliberation. In Bohman, J., and Rehg, W., eds., *Deliberative Democracy: Essays on Reason and Politics*. Cambridge, MA, USA: MIT Press. 243–277.

Ess, C. 1996. The political computer: Democracy, CMC, and Habermas. In Ess, C., ed., *Philosophical Perspectives on Computer-Mediated Communication*. Albany, NY, USA: State University of New York Press. 197–230.

Fairhead, J., and Leach, M. 1998. *Reframing Deforestation: Global Analysis and Local Realities: Studies in West Africa*. London, UK: Routledge.

Fox, J., and Parsons, S. 1998. Arguing about beliefs and actions. In Hunter, A., and Parsons, S., eds., *Applications of Uncertainty Formalisms*. Berlin, Germany: Springer Verlag (LNAI 1455). 266–302.

[8]We are grateful for partial funding from the British EPSRC, under grant GR/L84117 and a PhD studentship. We also thank William Rehg and Trevor Bench-Capon for discussions on these topics.

Girle, R.; Hitchcock, D.; McBurney, P.; and Verheij, B. 2001. Decision support for Practical Reasoning: a theoretical and computational perspective. In Reed, C., and Norman, T., eds., *Bonskeid Symposium on Argument and Computation*. (In press).

Gordon, T. F., and Karacapilidis, N. 1997. The Zeno argumentation framework. In *Proceedings of the Sixth International Conference on AI and Law*, 10–18. ACM Press.

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

Habermas, J. 1984. *The Theory of Communicative Action: Volume 1: Reason and the Rationalization of Society*. London, UK: Heinemann. (Published in German 1981).

Habermas, J. 1991. *Moral Consciousness and Communicative Action*. Cambridge, MA, USA: MIT Press. (Published in German 1983).

Habermas, J. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Cambridge, MA, USA: MIT Press. (Published in German 1992).

Habermas, J. 1998. *The Inclusion of the Other: Studies in Political Theory*. Cambridge, MA, USA: MIT Press. Edited by C. Cronin and P. De Greiff.

Keeney, R. L. 1996. The role of values in risk management. *Annals of the American Academy of Political and Social Science* 545.

Krause, P.; Fox, J.; Judson, P.; and Patel, M. 1998. Qualitative risk assessment fulfils a need. In Hunter, A., and Parsons, S., eds., *Applications of Uncertainty Formalisms*. Berlin, Germany: Springer Verlag (LNAI 1455). 138–156.

Lancet Editorial. 29 May 1999. Health risks of genetically modified foods. *The Lancet* 353(9167).

McBurney, P., and Parsons, S. 2000. Risk agoras: Dialectical argumentation for scientific reasoning. In Boutilier, C., and Goldszmidt, M., eds., *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, 371–379. Stanford, CA, USA: UAI.

Michelman, F. I. 1989. Conceptions of democracy in American Constitutional argument: the case of pornography regulation. *Tennessee Law Review* 291:293.

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

Parsons, S., and Green, S. 1999. Argumentation and qualitative decision making. In Hunter, A., and Parsons, S., eds., *The Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU99)*. Berlin, Germany: Springer Verlag (LNAI 1638). 328–339.

Pera, M. 1994. *The Discourses of Science*. Chicago, IL, USA: University of Chicago Press.

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

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

Rittel, H. W. J., and Webber, M. M. 1973. Dilemmas in a general theory of planning. *Policy Sciences* 155–169.

Searle, J. 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge UK: Cambridge University Press.

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

Stirling, A., and Mayer, S. 1999. *Rethinking Risk: A Pilot Multi-Criteria Mapping of a Genetically Modified Crop in Agricultural Systems in the UK*. Report, SPRU, University of Sussex, Brighton, UK.

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

Toll, J. E. 1999. Elements of environmental problem-solving. *Human and Ecological Risk Assessment* 5(2):275–280.

Toulmin, S. E. 1958. *The Uses of Argument*. Cambridge, UK: Cambridge University Press.

UK ESRC Global Environmental Change Programme. 1999. *The Politics of GM Food: Risk, Science and Public Trust*. Special Briefing No. 5. Brighton, UK: University of Sussex.

USA EPA. 1996. Environmental Protection Agency: Proposed guidelines for carcinogen risk assessment. *U.S. Federal Register* 61:17960–18011 (23 April 1996).

Willard, C. A. 1990. Authority. *Informal Logic* 12:11–22.

Wynne, B. 1996. May the sheep safely graze? A reflexive view of the expert-lay knowledge divide. In Lash, S.; Szerszynski, B.; and Wynne, B., eds., *Risk, Environment and Modernity: Towards a New Ecology*. London, UK: Sage. 44–83.