

University of Windsor

Scholarship at UWindsor

OSSA Conference Archive

OSSA 4

May 17th, 9:00 AM - May 19th, 5:00 PM

Computer Decision-Support Systems for Public Argumentation: Criteria for Assessment

Willaim Rheg
Saint Louis University

Peter Mcburney

Simon Parsons

Follow this and additional works at: <https://scholar.uwindsor.ca/ossaarchive>



Part of the [Philosophy Commons](#)

Rheg, Willaim; Mcburney, Peter; and Parsons, Simon, "Computer Decision-Support Systems for Public Argumentation: Criteria for Assessment" (2001). *OSSA Conference Archive*. 101.
<https://scholar.uwindsor.ca/ossaarchive/OSSA4/papersandcommentaries/101>

This Paper is brought to you for free and open access by the Conferences and Conference Proceedings at Scholarship at UWindsor. It has been accepted for inclusion in OSSA Conference Archive by an authorized conference organizer of Scholarship at UWindsor. For more information, please contact scholarship@uwindsor.ca.

Title: Computer Decision-Support Systems for Public Argumentation: Criteria for Assessment

Authors: William Rehg, Peter McBurney, and Simon Parsons

© 2001 William Rehg, Peter McBurney, and Simon Parsons

ABSTRACT: Computers have been used to support human decision-making for several decades, and increasingly, formal argumentation models are being used in their design. However, their use raises a number of conceptual and social-ethical questions that have yet to be fully addressed. We explore some questions raised by two current proposals for computer-mediated argumentation and decision-making, in particular the assessment of substantive quality, inclusiveness, and noncoerciveness, as these bear on legitimate policy decisions.

1. Introduction

Human reliance on computer technology is by now a well-established cultural fact. Indeed, we now rely on computers even for certain forms of decision-making: so-called “knowledge-based expert systems,” or “decision support systems,” have been developed over the last thirty or so years to support, or sometimes even replace human decision-making (Alty and Coombs 1984; Buchanan and Shortliffe 1984). As the name suggests, expert systems attempt to automate (by means of a knowledge base and inference mechanisms) the knowledge and reasoning skills of experts in a given domain, such as medical diagnosis, marketing decisions, and so on.¹ In this essay we are concerned with recent attempts to develop decision-support systems for processes of public policy argumentation. Like earlier expert systems, these argumentation-support systems incorporate both knowledge bases and inference mechanisms. Unlike the earlier systems, however, they place greater emphasis on the processes used for reasoning and inference than on the database of knowledge from which conclusions are drawn; thus we might designate them as “argumentation systems” in contrast to the earlier “knowledge systems.”² Dialectical approaches in particular have drawn attention in this developing area of Artificial Intelligence research.³ Indeed, some commentators have recommended the employment of such systems, via the Internet, to enable democratic participation in public policy decision-making processes (e.g. Ess 1996). The kind of argumentation systems we examine here are noteworthy for their attempt

¹ For an introduction, see Benfer, Brendt, and Furbee 1990.

² One can discern a natural evolution from knowledge systems to argumentation systems in the increasing concern with providing users with the underlying explanation for the system’s conclusion (or advice). Although developers of knowledge systems sought from the start to have their systems explain or justify their conclusions, early expert systems offered little in the way of explanation beyond high-level traces of the inference rules invoked in the chain of reasoning (Jackson 1986). Applications using some model of argumentation have included systems for: medical domains (Krause *et al.* 1995, Fox and Das 2000); legal domains (Gordon 1994, Verheij 1998, Bench-Capon *et al.* Forthcoming); public policy decision support (Gordon and Karacapilidis 1997, Gordon *et al.* 1997, Karacapilidis *et al.* 1997); geopolitical risk prediction (Seffers 1998); scientific discourse (McBurney and Parsons 2000,2001c); software design (Stathis 2000); and autonomous agent dialogues (Parsons *et al.* 1998, Sierra *et al.* 1998, Dignum *et al.* 2000). Carbogim *et al.* (2000) present a review of such applications

³ For example, Dung (1995) proposed dialectical argumentation within AI as a proof theory for non-monotonic reasoning; for a review of this area, see Prakken and Vreeswijk 2001. For broad overviews of developments in computational dialectics, see Walton 2000, Reed and Norman, forthcoming.

to draw explicitly upon dialectical theories of argumentation as a framework for processes of public deliberation involving multiple parties and interests.

The use of argumentation systems to assist public deliberation raises a number of conceptual and social-ethical questions that have yet to be fully addressed.⁴ On the one hand, argumentation systems provide a deliberative forum whose results can advise and, one hopes, improve the quality of decisions. If appropriately designed, such systems should be able to assist debate by tracking the various claims and arguments, by searching databases for relevant information, and by continually updating and assessing the overall state of the debate. They thereby help participants argue in a dialectically responsible manner, while offering them ample scope for modifying aspects of the system, such as the stock of inference rules and proof standards. On the other hand, the highly qualitative and publicly controverted character of public policy argumentation poses serious questions about what would constitute an “appropriate design.” In this paper we aim to clarify the criteria for an appropriately designed argumentation system. To focus and concretize the discussion, we start by describing, in Section 2, two current proposals for computer-mediated argumentation and decision-making, the Zeno system of Gordon and Karacapilidis and their colleagues at the German National Research Centre for Information Technology and the Risk Agora of McBurney and Parsons. In Section 3 we elaborate the basic evaluative dimensions appropriate for the deliberative contexts in which such systems are to be used. In the subsequent sections we indicate how one might apply these deliberative criteria for assessing the two argumentation systems. As we shall see, the critical assessment of argumentation-support systems must draw upon an understanding of the discursive bases of responsible public legitimation and must acknowledge inherent limits on what formal design features can achieve for legitimation.

2. Two Argumentation-Based Decision Support Systems

2.1

We begin by describing the two argumentation systems. The first is the Zeno system of the German National Research Centre for Information Technology (GMD) (Gordon and Karacapilidis 1997, Gordon *et al.* 1997, Karacapilidis *et al.* 1997). Zeno was developed to support decision-making in urban planning, as part of a larger European Community-funded project to develop innovative information systems infrastructure for public collaborative environmental planning. In these domains, there are multiple interested parties, with diverse professional or private backgrounds, interests, preferences and viewpoints, and they are often geographically dispersed. Because the application domain of Zeno involves urban planning decisions, the system has to integrate information which is spatially indexed with information which is not. And because the users may be diverse and geographically distributed, the system requires intuitive and easy-to-use interfaces, provided, preferably, across an internet platform. Neither of these elements of software design was technically straightforward, but these issues will not concern us here. Our focus is on the argumentation and decision support elements of Zeno.

⁴ As Haklay notes in a forthcoming review of information technology for environmental decision-making, the philosophical and ethical aspects of information systems design are rarely made explicit or even explored (Haklay Forthcoming). Cf. also Feenberg 1999.

The developers of Zeno define their system as “a mediation system”:

“a kind of computer-based discussion forum with particular support for argumentation. In addition to the generic functions for viewing, browsing and responding to messages, a mediation system uses a formal model of argumentation to facilitate retrieval, to show and manage dependencies between arguments, to provide heuristic information focusing the discussion on solutions which appear most promising, and to assist human mediators in providing advice about the rights and obligations of the participants in formally regulated decision making procedures.” (Gordon and Karacapilidis 1997, p. 10).

The argumentation model used by Zeno is a formal version of the informal Issue-Based Information System (IBIS) model of Rittel and Webber (1973), modified for the urban-planning domain. The IBIS model identifies several atomic elements of a discourse: *Issues*, the topic about which a discussion is conducted; *Positions*, which express some statement relevant to an issue; and *Arguments*, which present statements in favour or against particular *Positions*. Thus (using an example from Karacapilidis *et al.* 1997), an Issue may be: “Which site should the airport be located?”; Positions may then be statements designating different alternative sites, or groups of possible sites; and Arguments may be positive and/or negative attributes about these site alternatives and statements, such as: “Has easy access” or “Is private land.” Each element of this model can be attached to each other element at any time, so that, for example, a new issue can be raised at any point in a discussion, thus creating a subsidiary discussion on the new issue.

The IBIS model is well-suited for the display of a discussion as a hierarchical graph; the implementation of hypertext links enables a user to move easily around this graph from one thread of a discussion to another, or to access background data, supporting documents or contextual information, etc, associated with any element of the discussion. Among the objectives of the Zeno project is to enable snapshots of a debate: “*One important goal is to provide easy access to the current state of the planning process, at any time.*” (Gordon *et al.* 1997, Section 2). In order to provide for these, the designers of Zeno modified the IBIS model to permit the expression of preferences. Positions have the form of logical propositions but they do not have a context-independent truth status. Their meaning is defined by their role in a particular thread of discussion. Preferences are defined as particular types of positions with an internal structure of comparison of two (non-preference) positions. For example, two planning options for siting an airport, such as “public site” and “easy access” might enter into a preference that considers easy access as “more important than” a public site. This preference then constitutes a qualitative constraint, which may or may not be supported by further arguments (i.e., positions), and which may or may not be consistent with other constraints. These preferences and constraints, being positions, may themselves be the subject of discussion, via the articulation of arguments and the raising of further issues.

Zeno provides users with an overview of the argumentative status of positions, preferences, and constraints. By considering to extent to which each position satisfies the articulated constraints, Zeno permits positions to be labeled as acceptable or not at any time in a discussion. These position-labels can then be aggregated in various ways to assign labels to Issues, and Zeno does this to indicate the extent to which their current argument support meets defined standards of proof. The Zeno developers argue that no set of proof-standards is applicable across all application domains, and so they adopted a set of five labels from the field of jurisprudence, namely: (1) *Scintilla of Evidence*; (2) *Preponderance of Evidence*; (3) *No Better Alternative*; (4) *Best Choice*; and (5) *Beyond a Reasonable Doubt* (Gordon and Karacapilidis 1997). Each label is provided with a formal definition in terms of the presence or absence of positions and

arguments (although these definitions are not claimed to be instantiations of the legal definitions of these terms). With this argumentation structure, it is then straightforward for the system to present to a user the current status of a discussion and to show this as it changes.

The designers of Zeno identify the generic speech acts involved in contributing to a discussion, for example, “Raise an issue”, “Assert a position”, “State a preference”, etc (Gordon and Karacapilidis 1997). However, they do not, in the work published to date, articulate a definitive list of such speech acts or the rules that govern their use. By contrast, our second example of an argumentation system, the Risk Agora of McBurney and Parsons (2000,2001b,2001c) is fully specified in this manner, and we now discuss this system.

2.2

The Risk Agora has been proposed as a system to support deliberations over the potential health and environmental risks of new chemicals and substances, and the appropriate regulation of these substances. Determination of these issues typically first involves debate within the relevant scientific community over whether or not a significant correlation exists between the putative chemical cause and any observed health effects, and then, if a significant relationship is identified, a subsequent debate in the wider community over the consequences of alternative regulatory options then occurs. Initial development of the system (McBurney and Parsons 2000, 2001c) has focused on the first of these debates, the scientific debate, where the designers adopted explicit philosophies of science and of rational discourse. The philosophy of science they draw on is Pera’s three-person game model of science, where progress is made through the work of, firstly, scientists undertaking experiments, whose experiments provoke reactions from, secondly, Nature, whose responses are in turn mediated through, thirdly, the scientific community (Pera 1994). Because the members of the scientific community are (presumed to be) rational and willing participants in the process, and because the assertions in these scientific discourses are all subject to contestation and defence, the authors adopt a philosophy of rational discourse for such debates. For this they draw on Habermas’s theory of discourse ethics (Habermas 1983/1991), whose rules were first fully articulated by Alexy (1978/1990), and the principles of rational mutual inquiry of Hitchcock (1991).

Within this framework, the authors then articulate the locutions and rules of a dialogue-game, in the style of Hamblin (1971) and MacKenzie (1979), specifying the pre-conditions necessary for the execution of each locution and the changes each locution effects. The locutions of this game permit assertion, contestation and defence of propositions, modes of inference, assumptions and consequences, in what is essentially a persuasion dialogue (Walton and Krabbe 1995). Certain locutions incur requirements on the speaker to defend a statement if subsequently contested, and the formalism permits this to be in the form of an argument for the statement. Statements, assumptions, consequences and modes of inference may all be asserted with an attached modality, expressing the speaker’s degree of confidence in the assertion. As in the Hamblin games, Commitment stores track assertions made in the course of the debate. McBurney and Parsons have also begun the task of the detailed specification of a system to support discussion over regulatory options (McBurney and Parsons 2001b). In this work, they have drawn on Habermas’s theory of communicative action (Habermas 1981/1984) to define types of locutions (i.e., speech acts) appropriate for such discussions. Because these discussions are about actions, their formalization as dialogue games requires models of deliberation dialogues (Walton and Krabbe 1995), which is the subject of another paper at this OSSA meeting

(Hitchcock, McBurney and Parsons 2001). Thus, the work of specifying the Agora to support regulatory debates is on-going.

Unlike the Zeno system, the Risk Agora is not intended to support debates in real-time, and the designers have also not yet developed the intuitive, graphical interfaces present in Zeno. Rather, the Agora is intended to formally model and represent debates in the risk domain, so as:

- “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.”
- (McBurney and Parsons 2001b).

Although the Risk Agora has different aims from Zeno, its designers also desire to enable snap-shots of a debate to be taken at any time. This requirement has added importance in the risk domain, where regulatory decisions must be made even though a final determination of the state of scientific knowledge on a particular issue is not possible. For this reason, and like Zeno, the Agora defines a set of labels which are attached to each claim on the basis of the arguments presented for and against the claim up to that time. *Probable* claims, for example, are those for which an argument has been presented in the Agora, but for which no rebutting arguments (arguments for the negation of the claim) or undercutting arguments (arguments for negations of an assumption or an intermediate premise) have been presented. In this way, the dialectical status of a claim can be assessed at any time, thus providing a snapshot of the debate to that time. The designers of the Agora then examine the likelihood that a snapshot, taken at some finite time after commencement of a debate, is indicative of the longer-term state of the debate, assuming such a stable state is achieved. They show that the Agora has desirable properties when used for inference from finite snapshots to longer-term states in this way (McBurney and Parsons 2001c).

3. Dimensions for Evaluating Argumentation Systems

Along what dimensions might one analyze and evaluate such systems? For software systems designed for well-defined, decomposable and measurable tasks, such as those for the production of bills for use of an electricity network, for example, assessment of system competence and quality is straightforward; standardized methods have been developed, are in widespread use, and influence good software design (e.g. Kirwan and Ainsworth 1993). However, as Parker (2000) observes, these methods are inapplicable for most decision support systems, since most decisions and decision-processes are not amenable to such reductionist task analysis. She further notes that no methods have been developed for decision support systems.⁵

⁵ However, see the references in note 6 below. Parker (2000) proposes a design methodology based on the types of questions likely to be asked by a user of the system. While promising, this approach seems more

Perhaps it is because decisions are not amenable to reductionist analysis that relatively little attention has been paid in the Artificial Intelligence community to the question of the quality of decision support systems. Groothuis and Svensson, in research presented as recently as December 2000 and investigating the quality of computer-supported welfare assistance decisions in the Netherlands, say: “*To our knowledge this is the first investigation into the extensive use of expert systems in the daily practice of handling a very complex administrative task*” (Groothuis and Svensson 2000, p. 9).

It is true that designers of what we have termed knowledge systems — expert systems that encode some body of expertise — will typically compare performance of the system against a group of human experts through a number of test cases.⁶ However, such comparisons are fraught with difficulties. In some domains, what counts as “expertise” may be open to dispute, or subject to cultural and contextual variation.⁷ In other domains, normative decision theory (e.g., Lindley 1985) does not reflect the decision methods people actually use. Although one may simply take this as an illustration of deficiency of human decision-makers, in some cases one may not have any means of assessing the normative methods as superior.⁸ Moreover, normative decision theories have tended to focus only on those elements that are quantifiable, and so may ignore much that is salient to good decision-making (cf. Rehg 1997a). Assessing the quality of advice-giving systems poses further specific difficulties. If advice is not taken, is it necessarily of low quality or unhelpful to the decision-maker? And how does one assess the quality of advice if the world changes in a salient way between the giving of the advice and its execution?⁹

appropriate for what we have termed knowledge systems (those encoding expert knowledge) than for argumentation systems.

⁶ E.g., Buchanan and Shortliffe 1984; in chap. 30 of this work, the authors consider quality evaluation methods for decision support systems. Greenberg (1987) also considers evaluation issues for such systems, focusing his discussion on the validation of the inference rules used by the system. Taylor (1991) briefly discusses evaluation questions as part of an exploration of the wider organizational issues associated with deployment of these systems.

⁷ For example, the SimCoast expert system of the UK’s Centre for Coastal and Marine Sciences, for example, was developed to provide assistance to marine coastal environment decision-makers in developing nations (McGlade 1999). However, the expertise encoded in the system embodies a specific *Weltanschauung*, that of a standardized western scientific ontology, which may not accord with the worldviews or ontologies of non-scientist users in the developing world. Moreover, Western expertise may actually distort the perception of the local reality, as shown by the example of deforestation in West Africa (Fairhead and Leach 1998; cf. Harding 1998). Using developing-country residents as the “experts” for the design phase may well have led to a different system.

⁸ There has been considerable research on this variance (see, e.g., Nisbett and Ross 1980; Kahneman *et al.* 1982). In a recent review of the limited research into how people make important medical decisions, Schneider concluded that patients make decisions quite differently from experts or normative decision theory; however, we have few means to measure the quality of medical decisions. Selecting one procedure or course of treatment usually precludes the selection of alternatives, and so strict comparison of results of alternative decision options at an individual level is impossible. The diversity and complexity of individual circumstances and medical aetiologies make comparisons at an aggregated level also problematic (Schneider 1998, pp. 92, 97).

⁹ Developers of public mobile satellite communications networks such as Iridium and ICO, for example, had lead times of a decade for designing, manufacturing and deploying the innovative satellite technology they required (McBurney and Parsons 2001a). To guide this work, the intending investors sought the advice of market researchers on the size of the potential market for mobile satellite services. In the particular decade concerned (1989-1999) demand for terrestrial cellular services grew much faster than anyone had forecast at the outset of the period, to the detriment of demand for mobile satellite services at the end of the decade. In major part this growth in terrestrial cellular demand was spurred by technological changes and the spread of cellular network coverage unanticipated in

Still further difficulties are raised by the fact that advice is often linked with certain contextual limitations or caveats, which users sometimes ignore.¹⁰ How does one assess the quality of advice in such cases? Finally, how does one assess advice for extreme situations or rare events?¹¹

Beneath the foregoing methods of assessment and their vicissitudes lurks a shared assumption, namely that there is a correct or right or “true” decision to be reached by following a set of context-independent (albeit domain-specific) inference rules. However, if assessments based on this straightforward model run into difficulties for knowledge systems, we should expect even greater difficulties when we attempt to evaluate argumentation systems for public deliberative domains that involve multiple interested parties and highly complex value-laden issues. Decisions over urban planning policies, for example, may have no inherent, or independently accessible truth or correctness, since any viewpoint from which we may judge them will invariably be partial and never completely disinterested; only the process used to reach the decision can tell us whether or not it deserves to stand (Forester 1999; Bohman 1996).

Although the intended applications are different, both Zeno and the Risk Agora seek to support deliberative decision-making by multiple participants in a public policy domain. Given the difficulties we have just sketched, any evaluation of argumentation systems such as Zeno or the Agora must begin with a more sophisticated understanding of the relevant kind of decision-making at issue: specifically, assessment criteria must draw not only upon argumentation theory but also on an understanding of public policy formation. Our analysis thus takes as its point of departure Forester’s deliberative approach (1999, chap. 6) to mediated public policy formation and dispute resolution. The key to this approach lies in its understanding of deliberation as a process that is both transformatory and reasonable, and thus able to generate *legitimate decisions*, which is to say: rationally and publicly acceptable decisions. There are four major

1989. Was the advice given by the market researchers wrong because it did not predict unanticipated events, or, if anticipated, did not sufficiently emphasize unlikely events? The advice may not have been wrong if the world had been otherwise in the subsequent decade. Most management consultancy is not assessable or assessed for these two reasons (an observation based on the second author’s decade of experience as a management consultant).

¹⁰ In the one quality assessment of decision support systems known to us, Groothuis and Svensson (2000) assessed the quality of decisions made with the help of computer systems in the provision of welfare assistance in the Netherlands, comparing the actual determination reached in a sample of welfare assistance cases with the decision which should have been made under the relevant laws. The results of this assessment showed that decision quality varied according to the extent to which the system provided support for the complex administrative tasks involved. For those cases where the system provided full support, few human errors were made; in other cases, where only limited support was provided, decision errors were more frequent. The authors concluded that this was due to the human decision-makers trusting the advice of the system even when such trust was not warranted, thus revealing weaknesses in the human decision-makers rather than in the decision support systems.

¹¹ For example, a system designed to support water-flow management through a dam will only be required to recommend responsive actions for 200-year floods on average once every 200 years. There may be insufficient data to design the system or to predict its performance in these circumstances, and possibly only one case every 200 years on which to base a live assessment of that performance. Moreover, if a system is designed for an entirely new activity how does one assess the adequacy of its advice? The various intelligent agent systems currently being deployed by NASA for control of autonomous spacecraft are examples of systems which undertake completely new activities. How is it possible to rate their performance in any other but crude terms, such as overall mission success versus non-success? (NASA 1999)

criteria considerations or dimensions that issue from a deliberative approach. In this section we elucidate these dimensions, before applying them to Zeno and the Agora in subsequent sections.

3.1

A deliberative model of policy formation contrasts with the conventional pluralist model that conceived public policy questions as matters for negotiation and bargaining.¹² Deliberative approaches deny that all public issues and choices reduce to bargaining (Bohman and Rehg 1997). Although decisions over the division of scarce resources, or involving the conflict of particular (non-generalizable) interests, are typically resolved by resort to negotiations or bargaining, decisions over what actions to take in some circumstance require deliberations (cf. Walton and Krabbe 1995; Habermas 1996, chap. 4). According to Forester, mediators should approach public planning and policy formation as deliberations. This approach has important implications for what one expects of both the process and its possible outcomes. As a number of democratic theorists have argued (e.g., Elster 1986; Michelman 1988), deliberative political processes, unlike negotiations, require participants to adopt a civic standpoint oriented toward the transformation of individual preferences and interests in the direction of reaching agreement on a common good or general interest. Forester (1999, p. 184) calls this “*the self-transformative condition.*” Whereas participants in negotiations enter the decision-making process with their preferences fully-formed and aim to achieve a compromise position that balances competing preferences, participants in deliberations may learn from each other and even from the very fact of interacting. As Michelman elsewhere 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.” (Michelman 1989, p. 293)

Because of this openness to change, participants to a deliberation – assuming they are rational – should be willing to share information, a strategy that may not be in the self-interests of participants to a negotiation. To the extent that participants are unwilling to share their knowledge and preferences, we may consider the decision-process to be a negotiation rather than a deliberation.¹³

Within this deliberative approach we identify three further dimensions of analysis. To begin with, we want self-transformation to be reasonable: if participants change their views in response to others' input, then these changes should lead toward a substantively better outcome. Thus we must examine the properties of topic-specific considerations stemming from the particular question or dispute, which determine the adequacy conditions for high quality outcomes (3.2). However, as we shall see, substantive quality cannot be assessed independently of considerations bearing on the participants and their roles (3.3), which in turn point to difficult questions

¹² For a well-known statement of the pluralist model, which draws on James Madison's Federalist Paper no. 10, see Dahl 1956.

¹³ In Hitchcock *et al.* (2001), the characteristic of sharing information is taken to be one that distinguishes deliberation from negotiation dialogues. In contrast, because much of the research focus in the area of intelligent multi-agent software systems has been on automated negotiation (Jennings *et al.* 2001), this self-transformative condition is not satisfied by all agent systems.

concerning the process itself and its “vulnerability to insinuations of power” (Forester 1999, p. 177) (3.4).

3.2

For any particular dispute resolution or policy question, one can identify specific considerations tied to the particular issue or dispute at stake, and an adequate deliberative outcome must, presumably, take all these relevant considerations—or at the least, those which are most pressing or salient—into account. An environmental dispute, for example, will typically turn on particular scientific facts about the ecosystem in question, fallible prognostications about the impact of different alternative actions on that ecosystem, and economic assessments about the costs and benefits of different options for different affected parties. In addition, one can also expect that the parties will disagree over quality-of-life issues, fundamental goods and values (aesthetic beauty vs. economic growth, for example), and even, at a deeper level, over basic ontologies and worldviews (Kriesberg *et al.* 1989; Rehg 1999). How well a deliberation addresses the range of relevant considerations determines the quality of its outcome. In effect, the reasonable processing of all relevant considerations—the pertinent questions and objections—determines what counts as an adequate outcome.

The idea of relevance obviously plays a crucial role in this context—but what does it involve? Although the conception of argumentative relevance is very much an open research question (Johnson 2000, pp. 199—204), the factual dimensions of policy argument suggest, on the one hand, a process-independent conception: a consideration is relevant if failure to take it into account in a deliberation can lead to an unsuccessful policy. On the other hand, relevance is also partly internal to the deliberative process: at least some kinds of considerations, such as interests and values, must be relevant *for the participants* if they are to be relevant at all.

This twofold conception of relevance points toward a more precise characterization of the substantive quality of an outcome or decision. Note, to begin with, that substantive quality should probably not be reduced to “truth.” Rather, given the complexity of the various relevant considerations, one does better to characterize substantive adequacy in terms of the rational legitimacy and subsequent success of a policy choice (cf. Habermas 1996; Rehg 1997b). The standard of success stems from the external aspect of relevance. Because policy decisions partly rest on factual assumptions, policies can, in some sense, fail to be “correct,” given the way the world actually is. An environmental dispute resolution may rest on false scientific assumptions, or it may depend on an economic forecast that proves mistaken. To this extent, the renewed interest of argumentation theorists in truth is not entirely misplaced (Goldman 1999, chap. 5; Johnson 2000). But in policy-making contexts, factual claims intertwine with other sorts of reasons, for which the predicates of “true” versus “false” may be less appropriate than something like “justified”—and thus reasonable—versus “arbitrary.” Consequently, what matters is that, in taking all the relevant considerations into account, participants construct and reconstruct their viewpoints in a manner that is not only logically consistent but, more importantly, *dialectically responsible, or “responsive”* (Goldman 1994). To put this idea in a nutshell, participants hold a dialectically responsible position insofar as they have addressed all the relevant questions and objections and thereby reached the most plausible outcome relative to the alternatives.

One may, to be sure, link dialectically responsible positions with a broadened notion of truth, for example along the lines of Rescher’s dialectical account of plausible reasoning or Hintikka’s game-theoretic semantics (Rescher 1976; 1977; Hintikka, 1968). According to such

approaches, claims that hold up in a process of dialectical reasoning as more plausible, or that have the support of a winning strategy for the given argumentation system, enjoy a (defeasible) presumption of truth, relative to the appropriate burden of proof. Arguably, this concept of “truth” is applicable to many, if not most, decisions in the public policy domain, including the application domains for both Zeno and the Risk Agora.

In any case, the internal aspect of relevance is closely linked with dialectical responsibility, and thereby with rational *legitimacy*. For unless the relevant considerations are addressed to the satisfaction of the affected parties, policy decisions are not likely to be considered legitimate. Inasmuch as deficits in legitimacy tend to have a negative effect on compliance with decisions, the dialectical quality of a deliberative outcome is also crucial for its successful implementation (see Lind and Tyler 1988). Thus, substantive quality of the deliberation—its adequacy in taking all the relevant considerations into account—naturally leads into the question of participant inclusiveness, how well a process has given the various affected parties the chance to voice their concerns and affect the outcome. This takes us into the third and fourth dimensions of evaluation.

3.3

To address the issue of inclusive participation, we must characterize the involved parties and their roles. Here, we may first ask how many people are involved in the decision-process being supported by the computer system; the number involved may influence the model of argumentation that is appropriate. As Forester’s case studies (1999) makes clear, public policy deliberations may involve numerous parties with quite different perspectives and interests. Including all the affected parties and giving them voice in the deliberation is crucial to the legitimacy and consequent success of such deliberations. But evaluating argumentation systems for overt or “formal” exclusions—exclusions that are explicit in the distribution of roles and entitlements--only goes part way. More important are those subtle forms of exclusion and coercion that might be built into the design itself as a system of rules. This issue brings us to the role of the system in deliberation, in particular the power of the system in fostering or impeding the expression and reasonable processing of information and viewpoints—an issue we take up under the fourth evaluative dimension.¹⁴

3.4

A legitimate deliberative process must not only grant all the affected parties entrance but also give them effective opportunities to voice their opinions and, still further, foster the participants’ capacity for learning from one another. Only in this way will the outcome represent a generally acceptable resolution of the dispute or question at issue. This dimension of evaluation thus requires us to examine how well a deliberative design, *precisely as an interaction*, fosters a collective, rational learning process, such that reasonable and generally acceptable—and thus legitimate—outcomes are more likely. In particular, we must scrutinize deliberation for subtle forms of *exclusion* and for potential obstructions in the participants’

¹⁴ See Bohman 1996, esp. chap. 3. Even a technology as apparently benign as literacy can be used by those possessing it to establish and maintain political power, as Bledsoe and Robey (1986) show in their study of Arabic literacy among the Mende of Sierra Leone.

capacity for learning from one another. This dimension of analysis thus complements and goes beyond the dimensions of substantive quality and overt (or formal) participant inclusiveness.

In other words, we are concerned here above all with the *freedom* that participants have to express their views and learn from one another in a reasonable manner. Thus, the fourth dimension of analysis scrutinizes deliberative processes for their freedom from domination or coercion, whether the coercion issues from the external pressures that inhibit the exchange of views or from internal psychological mechanisms. Forester (1999, p. 184) terms this the “noncoercive condition” on deliberative participation. We shall thus have to examine the two argumentation systems for the kinds of power they exert on the public deliberation they assist. If such systems in some sense “mediate” or at least facilitate deliberation, then they wield considerable power over the process and its outcome. As Forester (1999, p. 180) points out: “If parties to public dispute-resolution processes not only construct agreements but reconstruct themselves—in part as a result of being exposed to new information, in part as a result of the constellation of participants—then the *political significance and power* of the mediator-facilitator’s role is more important to understand than ever before” (Forester 1999, p. 180).

4. Deliberative Self-Transformation

We now turn to a brief assessment of each of the two systems, Zeno and the Risk Agora, in terms of the four categories sketched above in Section 3. Our intention in these remaining sections is not to provide an exhaustive assessment, but simply to sketch some of the more obvious moves, as an illustration of the basic idea. The first dimension of evaluation concerned the self-transformative character of deliberative processes. Here we note briefly that, at least on the surface, or in terms of their formal design features, both argumentation systems are designed to allow for deliberative self-transformation and not for bargaining or negotiation.¹⁵ Because the Risk Agora is not intended to support real-time debates, its “participants” may in fact be representatives of positions, rather than real persons. However, the system does allow retraction of claims previously asserted or accepted, and thus permits self-transformation. For Zeno, in the paper that describes the model of decision-making used in the domain, revision of constraints is explicitly permitted so as to eliminate any inconsistencies in these (Gordon et al. 1997). However, because precise dialogue rules have not been presented, it is not clear how this revision process is undertaken. It is also unclear if participants may revise their statements of other preferences, positions or arguments, as there do not appear to be specific locutions for retraction of previous statements.

5. Substantive Quality

The analysis of substantive inclusiveness developed in Section 3.2 suggests that, to begin with, one assess the substantive quality of argumentation systems for public deliberation by their capacity to include all the relevant considerations, that is, the relevant factual considerations, empirical prognostications, and values. On the one hand, the external aspect of relevance suggests that the system should be able to incorporate all the factual information that would be relevant to a successful policy, insofar as success depends on the truth or accuracy of certain assumptions about the world and the persons affected by the policy. Concretely, this means that

¹⁵ For a review of automated negotiation systems, see Jennings et al. 2001.

the system should be able to draw upon the relevant expertise and make this available to participants. On the other hand, as we have seen such expertise is not the only source of relevant information, nor is it usually decisive for a *legitimate* outcome. Because relevance also has an internal (i.e., participant-relative) sense, substantive quality requires that the value-orientations, interests, and other considerations that the participants themselves consider relevant should be taken into account by the argumentation system, and thus opened up to public discussion. Again, the formal design features of both systems, Zeno and the Risk Agora, seem to allow for such substantive inclusiveness. Zeno does not seem to restrict the kind of information that can count as a “position,” which can include both facts and preferences, and thus it can incorporate, at least indirectly, value-orientations. To be sure, Zeno may prove to be limited in its ability to distinguish between mere preferences and deeper values—a limitation that may be a problem in some contexts. Here the Risk Agora goes further, allowing for assertions of various types of speech acts (factual claims, claims about what is right, or valuable, and so on; see McBurney and Parsons 2001b).

Secondly, substantive quality also requires that the relevant information and arguments be processed in a dialectically responsible manner. For systems engaged in supporting a human participant or participants to construct, evaluate, contest and defend arguments, we may propose the following list of evaluative questions for an argumentation system (building on Verheij 1999, p. 43): Does the system track the issues raised? Does it track the assumptions made and the reasons adduced? Does it track the conclusions drawn and the counterarguments adduced? Does it track the justification status of the statements made and the commitments incurred? Does it verify that users obey the pertaining rules of argument? Does it identify omissions and weaknesses in arguments? Does it identify counter-arguments? Both Zeno and the Agora meet most of these requirements. The exceptions are: it is not clear whether and how Zeno handles retraction or revision, as was mentioned above; Zeno tracks but does not prohibit violations of its rules, so as to permit the greatest possible flexibility to participants (Gordon *et al.* 1997); and neither system identifies weaknesses or omissions in arguments.

6. Formal Inclusivity

The third dimension of evaluation is concerned with the extent to which all the stakeholders are represented in the world of the computer-aided decision-process. The formal design features of neither Zeno nor the Agora appear to limit participation in any overt way, so that all those people interested in the discussion topic and willing to accept the rules of the dialogue are permitted to participate (in the case of Zeno) or able to have their views represented (in the case of the Agora). In other words, each system is formally inclusive: it gives anyone an equal right to enter the process. This criterion of formal inclusivity, however, is quite limited. In particular, it leaves two problems unresolved, which reveal the importance of the context-of-use in assessing such systems.

First, formal inclusivity does not mean that all affected parties will in fact avail themselves of the system. One must ask whether all the important human stakeholders *in fact* involved with and through the computer system. Do all the conversations between them take place through the medium of the system, or do some of them occur off-line? The answers to these questions for any system will depend on the specifics of their implementation. We could imagine, for example, a situation where use of Zeno was mandated for an urban planning decision, with all stakeholders forced to conduct conversations occurring in and through it. The system has been

designed to support this level of use, and the technical design ensures incorporation and integration of information of heterogeneous types (maps, blueprints, reports, email messages, etc). However, without such mandating of use, there is no guarantee that many, and especially many important, interactions between stakeholders would not occur away from the system. Thus, although argumentation systems may assist in clarifying issues, identifying areas of agreement and differences, and deriving the consequences of arguments, so assisting debate, much of the debate may in fact occur off-line. This situation becomes problematic, however, if all the relevant contributions are not represented in the system, for then one cannot expect it to provide an accurate record of the debate or of the status of various positions.¹⁶

Second, in decision-making contexts the particular formalism itself can present a subtle form of exclusion. For example, if a particular viewpoint rests on styles of argument that are excluded—or simply not representable—in the system, then those viewpoints are excluded from the start.¹⁷ This issue becomes important for types of inferences or arguments whose validity is highly context-dependent (e.g., appeals to authority, or to emotions, etc.; see Walton 1995). This danger is lessened in the Agora, which allows participants to challenge and revise the rules. The Zeno system, on the other hand, attempts to avoid a pre-set formalism for generating arguments, and it seems to allow participants to choose from a range of burden-of-proof standards; indeed, the system even allows users to violate system norms (Gordon and Karacapilidis 1997, pp. 16-17). Presumably these features also lessen the danger of inbuilt exclusion, albeit at the cost of efficient decision making.

With these cursory observations, however, one has only scratched the surface. The possibility of still more subtle, context-dependent forms of exclusion and coercion may arise from inequalities among participants in their familiarity with rule systems and their computer skills, which tend to favor more educated participants. Generally speaking, inequalities in “political capacity”—which can arise in a variety of ways—threaten to engender an elitism of the better educated (Bohman 1996, chap. 3). Although we cannot address this issue here in much depth, we can say something about the starting point for such an analysis: what role does the system itself play in the deliberation, and how might this steer the deliberative process, for better or worse?

7. The Role of the System in Deliberative Contexts: Toward a Noncoercive Process

7.1

Assessing the power of argumentation systems presupposes an understanding of their role(s) in deliberation. Like expert systems in general, argumentation systems can play a variety of roles in deliberation and decision-making. Three kinds of role especially interest us here: participant

¹⁶ By contrast, in systems of autonomous software agents using argumentation, such as the team-formation systems of Dignum *et al.* (2000), the entirety of the dialogue occurs inside the system.

¹⁷ For example, Gardenfors 1994 discusses the reasoning processes used by illiterate Uzbeki peasants, and argues that apparent violations of deductive inference rules are, in fact, differential assessments of argument premises according to the perceived experience of the proponent of the premise. A formalism which did not represent this aspect of argument premises would not be able to represent the styles of arguments used by the Uzbekis.

support, record-keeping (what we will call the “orrery” role), and forum-creation. We begin by simply clarifying how each system plays each of these roles.

First, then, argumentation systems may *support participants*. In fact, the Zeno system provides support in a number of ways. By tracking the deliberation as it unfolds—the changing commitments, open questions, and so on—Zeno supports all the human participants, presumably in a neutral fashion. But it can also provide a kind of partisan support insofar as it assists participants in their construction of arguments and counterarguments. Systems that provide this type of support have been called argument-assistance systems in artificial intelligence (Verheij 1999). Finally, it supports a human moderator or mediator, where this person is not him- or herself one of the decision-makers, for example by identifying common assumptions across different arguments. As we have already seen, such mediators play an important role in public policy deliberations. Urban planners, for instance, often assist community groups to reach a consensus in this manner (Forester 1999). Indeed, Zeno is designed with just these sorts of support in mind. Although the Risk Agora has a somewhat different aim, by reconstructing arguments it not only helps all participants track a debate but can also support individual users as they construct their own positions and arguments. And when the Agora system is fully specified, and thus able to support deliberation over the consequences of regulatory options, it may also support those tasked with facilitating such decisions.

Second, the system may play a record-keeping role, or what we call an *orrery role*, on the analogy of mechanical models of the solar system; it is often undertaken so as to achieve an understanding the reasoning used in the decision-process. Each of the two systems records the reasoning used to reach a decision, Zeno in real-time and with complete accuracy, the Agora in a formal reconstruction of the process. Indeed, Agora’s main role is the orrery one.¹⁸

Third, an argumentation system may support the entire process of decision-making by providing *a forum in which to undertake dialogue*, with defined protocols for this discourse. In this role the system provides something like a structured space in which participants interact. The forum-creating role is central to Zeno’s design. In contrast, although the Risk Agora is designed as a forum for discussion, this is only in an ideal sense, as it provides a forum for the reconstruction of arguments rather than for real-time support.

Note that neither system plays the role of participant or decision-maker. There are, to be sure, argumentation systems envisioned for this more active role, for example the StAR system designed for the automated prediction of chemical properties such as toxicity.¹⁹ Such systems may replace a human decision-maker, a prospect that has raised concern on the part of philosophers and computer scientists.²⁰ When the system plays the role of participant, it needs to

¹⁸ Much of the research in the field of AI and Law appears devoted to developing systems which take an orrery role, formalizing the reasoning processes used by decision-makers so as to better understand them (Bench-Capon *et al.* Forthcoming).

¹⁹ Tonnelier *et al.* 1997. Likewise, systems of autonomous software agents assume decision-making is being undertaken by computer decision-makers, and much of the research effort of second-generation electronic commerce, for instance, is directed at such automated decision-making (Jennings *et al.* 2001).

²⁰ The question to what extent humans should delegate decision-making authority to computer systems has been discussed in ethical philosophy, though a well-known early objection was raised by AI researcher Joseph Weizenbaum 1976; see also Moor 1979, Kuflik 1999.

be able to generate, evaluate, contest and defend arguments itself. This is true even for the mediation role, as a mediator may need to find common ground between different positions; for example, he or she may need to argue that two opposing positions share common assumptions, or that one implies another. How well a given system executes these tasks will be important in any evaluation. Neither Zeno nor the Agora appears to have the capability to generate, evaluate, contest or defend arguments. For the Agora, acting as an orrery, these capabilities are not required. For Zeno, because the designers do not seek to automate the role of mediator, the system can also operate without these capabilities (Gordon and Karacapilidis 1997, p. 11). Indeed, from a computer science perspective, the generation of appropriate arguments in a dialogue, even using only the limited sets of locutions of Zeno or the Agora, is a challenging research problem, and one that is not yet solved.

Rather, both Zeno and the Agora support responsibility to explain decisions reached through their use only insofar as the participants justify their statements to each forum. The Agora has a stronger claim than Zeno to this capability, as the dialogue rules require those participants asserting claims to provide arguments justifying these claims when questioned or contested. It is not clear that similar requirements are incurred by the participants in Zeno.

7.2

With the roles of each system clarified, we are in a position to sketch the assessment issues raised by the insinuations of power in deliberative processes. Expert systems raise a number of questions in this regard: What is the institutional status of computational-deliberative results? On whose behalf is the system developed and deployed? Who should have rights to its findings?—and so on.²¹ These questions illustrate the issues that may arise from an examination of the social and institutional relationships surrounding the use of any decision-support technology. None of these questions has a straightforward answer, and just resolution may only be possible on a case-by-case basis. Exploring them, one is confronted with those whom in sociology have been called “the locally powerful” (Bell 1978) – people who influence or control the decisions of others, even if only in a particular context. In the present context, one would expect the locally powerful to include users who are especially adept at employing the argumentation system to support their viewpoint, or those users whose style of reasoning and public representation is favored by the system formalism. The danger in such inequality is that the more powerful players—perhaps even despite their good intentions—would fail to give some viewpoints the consideration they deserve on the merits. Viewpoints and arguments would be hastily dismissed, not given an adequate hearing. These non-standard approaches would thus be subtly forced out of the deliberative forum.

In this scenario, the quality of the deliberation is harmed in a number of ways. Substantively, relevant considerations are not duly taken into account. Such truncation can in itself lead to policies that are less than adequately informed, even about certain factual considerations, and are thus less likely to prove successful.²² But truncated deliberations can also

²¹ Such questions arise, for example, in the use of intelligent systems in medical domains. See, for example, Emery *et al.* 1999.

²² A tragic example of this point is the Challenger launch decision; according to Vaughan (1996), the mistaken decision stemmed in part from deliberative inequalities and subtle forms of exclusion and coercion in NASA decision-making, which consequently led managers to underestimate important data.

mean that the interests and values of some parties are overlooked, a deficit that is bound to damage the legitimacy of the outcome, its acceptability to all those expected to comply with the decision at issue. Finally, by not considering some viewpoints sufficiently, possibilities for self-transformation would be missed, both by those with power and those lacking it.

We are thus brought back to the issue broached in Section 3, regarding the freedom of the participants for a reasonable self-transformation. Only now we cannot define such reasonableness by appealing to the argumentative mechanisms *built into* the system, the formal design features treated in Sections 4-6. For now it is precisely such features—the dialectical formalism and formal inclusivity—that are in question. We are asking how such system features perform, how they are actually used, in particular contexts. Thus one must rather have recourse to what Habermas (1990, pp. 88-89), drawing on Alexy (1978/1990), has dubbed “process” criteria. Although Alexy’s first criterion is that of formal inclusivity, we are now concerned above all with those criteria or “rules” that define the freedom of participants to construct and reconstruct their viewpoints within discourse.²³ The decisive critical perspective behind such criteria is this: to what extent does the deliberation—as a social interaction—allow participants *rationally and freely* to adopt and change their viewpoints, that is, solely on the basis of an insight into the better argument? The rules that define such freedom call for equal, coercion-free participation in the process itself. Thus one might summarize the requirement this way: each participant should have an equal freedom to express and revise his or her viewpoints, feelings, interests, questions, and so on. Such a requirement cannot be guaranteed by any concrete set of formal rules; rather it involves an “idealization” that actual discourses and deliberation only rarely satisfy in full. However, if participants are to consider deliberative outcomes as trustworthy or legitimate, then they must suppose that their actual deliberative process has at least approximated this idealization. More precisely, the actual process of argumentation must supply participants with the sense that their positions are ones that could hold up under idealized conditions of discourse (see Habermas 1993, pp. 49-57).

Applying these process criteria takes us beyond the assessment of formal design in abstracto to its performance in concrete deliberative contexts. One can, to be sure, design the formal features of system with something like Alexy’s criteria in mind. In doing so, one sees to it that the system formalism makes provision for the various kinds of argumentative moves that participants should each be able to employ if they are to represent and change their views. Our assessment of substantive quality suggests that the two systems to some extent incorporate, as formal design features, aspects of Alexy’s rules. In fact, McBurney and Parsons (2001c) are quite explicit about this: they show that the Risk Agora implements an applicable subset of Alexy’s rules of discourse ethics (Alexy 1978/1990) and 15 of Hitchcock’s 18 principles of rational mutual enquiry (Hitchcock 1991). Zeno seems to have a lesser capacity to represent different types of speech acts, and its dialogue-game rules are not articulated (at least, in the publications known to us), but it does allow for participants to represent and change their

²³ Habermas (1990, p. 89) highlights the following rules: each participant is allowed to question any assertion, to introduce any assertions whatsoever, and to express his or her attitudes, desires, and needs, and no participant may be prevented from exercising these rights by internal or external coercion.

preferences.²⁴ But formal design mechanisms and their abstract assessment leave open the critical question that is decisive for the process level: how well participants of diverse backgrounds and capacities make *actual* use of this formalism in concrete contexts of use remains an open research question. It is true that the identification and formal demonstration of the discourse properties of a system will provide confidence to its users. In the final analysis, however, such systems will be successfully implemented only insofar as parties to deliberations actually use the systems and consider computer-assisted outcomes as superior.

At the process level, then, the pertinent question is this: how well does the actual use of the argumentation systems for a given deliberation lead to an outcome that is legitimate, that is, reasonable in light of argumentative idealizations? Given the different roles delineated above, specific questions such as the following become important. Again, these questions are more illustrative than exhaustive of the possibilities.

First, *given the particular context of use*, do the support mechanisms favor some participants over others? Here the partisan support role merits particular scrutiny. Is argument-assistance equally available to each participant? Is the system flexible enough to capture all the various styles of argument that might plausibly merit consideration?

Second, *given the particular context of use*, is the tracking or record-keeping genuinely “neutral”? Marxist and postmodern critics of power have cast serious doubt on ideals of neutrality (e.g., Harding 1998). The crucial question here may perhaps be formulated in a more concrete manner as follows: Can each affected party perceive that the system has represented its position, interests, values, and arguments accurately? Again, the flexibility of the system formalism becomes an important factor in system quality.

The third area for questions involves the system as a forum. Here the basic question follows directly from the idealizing requirement above: *given the particular context of use*, is the forum one in which each party has an equal opportunity to present and argue for its viewpoint—and is equally free to reconstruct its position in the light of the counterarguments?

8. Conclusions

This paper has explored a number of evaluative issues associated with the use of computer decision support systems for public deliberation, in particular the Zeno system of Gordon and Karacapilidis and their colleagues at GMD and the Risk Agora of McBurney and Parsons. We have argued that assessments of the quality of such systems present challenges to system designers and users, and we have suggested a conceptual structure within which to undertake such assessments. This structure draws on deliberative democratic theory as developed in Forester’s (1999) framework for assessment of deliberative planning processes. Our analysis indicates that the formal design features of the two argumentation systems capture many of the features of reasonable argumentation—mechanisms for ensuring substantive inclusiveness of relevant considerations, dialectically responsible argumentation, and an overtly open forum of participation. But these properties do not ensure high quality performance in the actual

²⁴ Although similar procedural criteria have been proposed for assessment of automatic electronic auction systems (Sandholm 1999), to our knowledge this is the first time such proposals have been made for argumentation-based computer systems.

deliberative contexts of use. For that, all the parties affected by a policy decision must actually avail themselves of the system and, in addition, find that the system increases their understanding of the issue, their ability to voice their concerns, and their confidence that outcomes and decisions are thereby superior to what they would otherwise have been.

Acknowledgments

This work was partly funded by the British Engineering and Physical Sciences Research Council (EPSRC) through a PhD studentship. We gratefully acknowledge this support. We are also grateful for discussions on these topics with Trevor Bench-Capon, Rod Girle, Muki Haklay, David Hitchcock and Bart Verheij.

References

- Alexy, R. (1978/1990): "A theory of practical discourse." In: S. Benhabib and F. Dallmayr (Editors): *The Communicative Ethics Controversy*, pp. 151—190. Studies in Contemporary German Social Thought. Cambridge, MA, USA: MIT Press. (Translation by D. Frisby of: "Eine Theorie des praktischen Diskurses." In: W. Oelmüller, Editor, *Normenbegründung-Normendurchsetzung*. Paderborn, Germany: Schöningh, 1978).
- Alty, J. L. and M. J. Coombs (1984): *Expert Systems: Concepts and Examples*. Manchester, UK: National Computer Centre Publications.
- Amgoud, L., N. Maudet and S. Parsons (2000): "Modelling dialogues using argumentation." In: E. Durfee (Editor): *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, pp. 31—38. Boston, MA, USA: IEEE Press.
- Bell, C. (1978): "Studying the locally powerful: Personal reflections on a research career." In: C. Bell and S. Encel (Editors): *Inside the Whale: Ten Personal Accounts of Social Research*. Sydney, Australia: Pergamon Press.
- Bench-Capon, T. J. M., J. Freeman, H. Hohmann and H. Prakken (forthcoming): "Computational models, argumentation theories and legal practice." In: C. Reed and T. Norman (Editors): *Symposium on Argument and Computation*, Bonskeid, Scotland.
- Benfer, R. A., E.E. Brent, Jr., and L. Furbee (1990). *Expert Systems*. Newbury Park: Sage.
- Bledsoe, C. H. and K. M. Robey (1986): "Arabic literacy and secrecy among the Mende of Sierra Leone." *Man*, 21 (2) (n.s.): 202—226.
- Bohman, J. (1996): *Public Deliberation: Pluralism, Complexity and Democracy*. Cambridge, MA, USA: MIT Press.
- Bohman, J., and W. Rehg, eds. (1997): *Deliberative Democracy: Essays on Reason and Politics*. Cambridge, Mass.: MIT Press.
- Buchanan, B. G. and E. H. Shortliffe (1984): *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley.
- Carbogim, D.V., D. S. Robertson and J. R. Lee (2000): "Argument-based applications to knowledge engineering." *The Knowledge Engineering Review*, 15 (2): 119—149.
- Dahl, R. A. (1956): *A Preface to Democratic Theory*. Chicago: University of Chicago Press.
- Dignum, F., B. Dunin-Kępcicz and R. Verbrugge (2000): "Agent theory for team formation by dialogue." In: C. Castelfranchi and Y. Lesperance (Editors): *Pre-Proceedings of the Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL-2000)*, pp. 141—156. Boston, USA.

- Dung, P. M. (1995): "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-persons games." *Artificial Intelligence*, 77: 321—357.
- Elster, J. (1986). "The market and the forum." In: J. Elster and A. Aunund (Editors): *The Foundations of Social Choice Theory*. Cambridge: Cambridge University Press, pp. 103—132. (Reprinted in Bohman and Rehg 1997, pp. 3—33.)
- Emery, J., R. Walton, A. Coulson, D. W. Glasspool, S. Ziebland and J. Fox (1999): "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.
- Ess, C. (1996): "The political computer: Democracy, CMC, and Habermas." In: C. Ess (Editor): *Philosophical Perspectives on Computer-Mediated Communication*, pp. 197—230. Albany, NY, USA: SUNY Press.
- Fairhead, J. and M. Leach (1998): *Reframing Deforestation: Global Analysis and Local Realities: Studies in West Africa*. London, UK: Routledge.
- Feenberg, A. (1999): *Questioning Technology*. New York, NY, USA: Routledge.
- Forester, J. (1999): *The Deliberative Practitioner: Encouraging Participatory Planning Processes*. Cambridge, MA, USA: MIT Press.
- Fox, J. and S. Das (2000): *Safe and Sound: Artificial Intelligence in Hazardous Applications*. Cambridge, USA: MIT Press.
- Gardenfors, P. G. (1994): "The role of expectations in reasoning." In: M. Masuch and L. Polos (Editors): *Knowledge Representation and Reasoning under Uncertainty: Logic at Work*, pp. 1—16. Berlin, Germany: Springer. (Lecture Notes in Artificial Intelligence 808).
- Girlé, R., D. Hitchcock, P. McBurney and B. Verheij (Forthcoming): "Decision Support for Practical Reasoning: a theoretical and computational perspective." In: C. Reed and T. Norman (Editors): *Symposium on Argument and Computation*, Bonskeid, Scotland.
- Goldman, A. I. (1994): "Argumentation and social epistemology." *Journal of Philosophy*, 91: 27—49.
- Goldman, A. I. (1999): *Knowledge in a Social World*. Oxford, UK: Oxford University Press.
- Gordon, T. F. (1994): "The Pleadings Game: An exercise in computational dialectics." *Artificial Intelligence and Law*, 2: 239—292.
- Gordon, T. F., and N. Karacapilidis (1997). "The Zeno argumentation framework." *Proceedings of the Sixth International Conference on AI and Law*, pp. 10—18.. New York, NY, USA: ACM Press.
- Gordon, T. F., N. Karacapilidis, H. Voss and A. Zauke (1997): "Computer-mediated cooperative spatial planning." In: H. Timmermans (Editor): *Decision Support Systems in Urban Planning*, pp. 299—309. E and F. N. Spon.
- Greenberg, H. J. (1987): "Validation of Decision Support Systems." In: G. Mitra (Editor): *Mathematical Models for Decision Support*, pp. 641-657. Berlin: Springer (NATO ASI Series).
- Groothuis, M. M., and J. S. Svensson (2000): "Expert system support and juridical quality." In: J. Breuker (Editor): *Proceedings of the Thirteenth International Conference on Legal Knowledge-Based Systems (JURIX-2000)*. Enschede, The Netherlands.
- Habermas, J. (1981/1984): *The Theory of Communicative Action: Volume I: Reason and the Rationalization of Society*. London, UK: Heinemann. (Translation by T. McCarthy of: *Theorie des Kommunikativen Handelns, Band I, Handlungsrationality und gesellschaftliche Rationalisierung*. Frankfurt-am-Main: Suhrkamp, 1981.)
- Habermas, J. (1983/1991): *Moral Consciousness and Communicative Action*. Cambridge, MA, USA: MIT Press. (Translation by C. Lenhardt and S. W. Nicholsen of: *Moralbewusstsein und kommunikatives Handeln*. Frankfurt-am-Main: Suhrkamp, 1983.)

- Habermas J. (1991/1993): *Justification and Application: Remarks on Discourse Ethics*. Cambridge, MA, USA: MIT Press. (Translation by C. P. Cronin of *Erläuterung zur Diskursethik*. Frankfurt-am-Main: Suhrkamp, 1991.)
- Habermas, J. (1992/1996): *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Cambridge, MA, USA: MIT Press. (Translation by W. Rehg of *Faktizität und Geltung*, Frankfurt-am-Main: Suhrkamp, 1992.)
- Haklay, M. (Forthcoming): *Public Environmental Information Systems: Challenges and Perspectives*. PhD Thesis. Department of Geography, University College, London, UK.
- Hamblin, C. L. (1971): "Mathematical models of dialogue." *Theoria*, 37: 130—155.
- Harding, S. (1998): *Is Science Multicultural?* Bloomington: Indiana University Press.
- Hintikka, K. J. J. (1968): "Language-games for Quantifiers." *Americal Philosophical Quarterly Monograph Series 2: Studies in Logical Theory*, pp. 46—72. Oxford, UK: Blackwell.
- Hitchcock, D. (1991): "Some principles of rational mutual inquiry." In: F. van Eemeren et al. (Editors): *Proceedings of the Second International Conference On Argumentation*, pp. 236—243. Amsterdam: SICSAT.
- Hitchcock, D., P. McBurney and S. Parsons (2001): "A Framework for Deliberation Dialogues." In: H. V. Hansen, C. W. Tindale, J. A. Blair and R. H. Johnson (Editors): *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*. Windsor, Ontario, Canada (this volume).
- Jackson, P. (1986): *Introduction to Expert Systems*. Wokingham, UK: Addison-Wesley.
- Jennings, N. R., P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra and M. Wooldridge (2001): "Automated negotiation: prospects, methods and challenges." *International Journal of Group Decision and Negotiation*. (In press).
- Johnson, R. (2000): *Manifest Rationality: A Pragmatic Theory of Argument*. Mahwah, NJ, USA: Erlbaum.
- Kahneman, D., P. Slovic and A. Tversky (1982) (Editors): *Judgment under Uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Karacapilidis, N., D. Papadias, T. Gordon and H. Voss (1997): "Collaborative environmental planning with GeoMed." *European Journal of Operational Research*, 102(2): 335—346.
- Kirwan, B. and L. K. Ainsworth (1993): *A Guide to Task Analysis*. London, UK: Taylor and Francis.
- Krause, P., S. Ambler, M. Elvang-Goransson and J. Fox (1995): "A logic of argumentation for reasoning under uncertainty." *Computational Intelligence*, 11(1): 113—131.
- Kriesberg, L., T. A. Northrup, and S. J. Thorson (Editors) (1989): *Intractable Conflicts and Their Transformation*. Syracuse, USA: Syracuse University Press.
- Kuflik, A. (1999): "Computers in control: rational transfer of authority or irresponsible abdication of autonomy." *Ethics and Information Technology*, 1(3): 173—184.
- Lind, E. A. and T. R. Tyler (1988): *The Social Psychology of Procedural Justice*. New York, USA: Plenum.
- Lindley, D. V. (1985): *Making Decisions*. London, UK: Prentice-Hall. Second Edition.
- MacKenzie, J. D. (1979): "Question-begging in non-cumulative systems." *Journal of Philosophical Logic*, 8: 117—133.
- McBurney, P. and S. Parsons (2000): "Risk Agoras: Dialectical argumentation for scientific reasoning." In: C. Boutilier and M. Goldszmidt (Editors): *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pp. 371—379. Stanford, CA, USA: Morgan Kaufmann.

- McBurney, P., and S. Parsons (2001a): "Using Belief Functions to forecast demand for Mobile Satellite Services." In: R. P. Srivastava and T. Mock (Editors): *Belief Functions for Business Decisions*. New York, NY, USA: Physica (in press).
- McBurney, P., and S. Parsons (2001b). "Intelligent systems to support deliberative democracy in environmental regulation." *Journal of Information and Communications Technology Law*, 10(1): 33—43.
- McBurney, P. and S. Parsons (2001c): "Representing epistemic uncertainty by means of dialectical argumentation." *Annals of Mathematics and Artificial Intelligence* (in press).
- McGlade, J. (1999). "Indigenous knowledge and the coastal environment: the SimCoast system." *Presentation to ESRC/GEC Seminar on Environmental Knowledge: Uncertainty, Authority and Responsibility*. Centre for the Study of Environmental Change, University of Lancaster, UK. 22 November 1999.
- Michelman, F. I. (1988): "Law's Republic." *Yale Law Review*, 97(8): 1493—1537.
- Michelman, F. I. (1989): "Conceptions of democracy in American Constitutional argument: the case of pornography regulation." *Tennessee Law Review*, 56: 291—319.
- Moor, J. (1979): "Are there decisions computers should never make?" *Nature and System*, 1: 217—229.
- NASA (1999). "Computer program assumes spacecraft command." Press release, 17 May 1999. Pasadena, CA, USA: National Aeronautics and Space Administration. Available from: www.jpl.nasa.gov/releases/99/remote.html
- Nisbett, R. and L. Ross (1980): *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ, USA: Prentice-Hall.
- Parker, C. (2000): "Requirements analysis and evaluation of decision support systems: the question approach." In: J. Reid and M. Lalmas (Editors): *Proceedings of the British Computer Society/Institution of Electrical Engineering Workshop on the Evaluation of Information Management Systems*. Department of Computer Science, Queen Mary College, University of London, London, September 2000.
- Parsons, S., C. Sierra and N. R. Jennings (1998): "Agents that reason and negotiate by arguing." *Journal of Logic and Computation*, 8(3): 261—292.
- Pera, M. (1994): *The Discourses of Science*. Chicago, IL, USA: University of Chicago Press.
- Prakken, H. and G. Sartor (1998): "Modelling reasoning with precedents in a formal dialogue game." *Artificial Intelligence and Law*, 6: 231—287.
- Prakken, H. and G. Vreeswijk (2001): "Logics for Defeasible Argumentation." In: D. Gabbay (Editor): *Handbook of Philosophical Logic*, Second Edition. Dordrecht, The Netherlands: Kluwer Academic (forthcoming).
- Reed, C. and T. Norman (Editors) (Forthcoming): *Symposium on Argument and Computation*, Bonskeid, Scotland.
- Rehg, William (1997a). "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*, pp. 358—377. New Haven, CN, USA: Yale University Press.
- Rehg, William (1997b): "Legitimacy and deliberation in epistemic conceptions of democracy: Between Habermas and Estlund." *The Modern Schoolman*, 74: 355—374.
- Rehg, William (1999): "Intractable conflicts and moral objectivity: A dialogical, problem-based approach." *Inquiry*, 42: 229—258.
- Rittel, H. W. J. and M. M. Webber (1973): "Dilemmas in a general theory of planning." *Policy Sciences*, 4: 155—169.

- Sandholm, T. W. (1999): "Distributed rational decision making." In: G. Weiss (Editor): *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, pp. 201—258. Cambridge, MA, USA: MIT Press.
- Schneider, C.E. (1998). *The Practice of Autonomy: Patients, Doctors and Medical Decisions*. Oxford, UK: Oxford University Press.
- Seffers, G. I. (1998): "Crisis system for President: DARPA's Project Genoa would speed U.S. decision-making", *Defense News*, 30 March, 13(13): 1.
- Sierra, C., N. R. Jennings, P. Noriega and S. Parsons (1998): "A framework for argumentation-based negotiation." In: M. P. Singh, A. Rao and M. J. Wooldridge (Editors): *Intelligent Agents IV: Agent Theories, Architectures, and Languages. Proceedings of the Fourth International ATAL Workshop*, pp. 177—192. Lecture Notes in Artificial Intelligence No. 1365. Berlin, Germany: Springer Verlag.
- Stathis, K. (2000): "A game-based architecture for developing interactive components in computational logic." *Electronic Journal of Functional and Logic Programming*, 2000(5).
- Taylor, A. (1991): "Organizational and social implications." In: T. Bench-Capon (Editor): *Knowledge-Based Systems and Legal Applications*, pp. 309-327. London, UK: Academic.
- Tonnelier, C. A. G., J. Fox, P. N. Judson, P. J. Krause, N. Pappas and M. Patel (1997): "Representation of chemical structures in knowledge-based systems: the StAR system." *Journal of Chemical Information and Computer Sciences*, 37: 117-123.
- Toulmin, S. E. (1958): *The Uses of Argument*. Cambridge, UK: Cambridge University Press.
- Vaughan, D. (1996): *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago, IL, USA: University of Chicago Press.
- Verheij, B. (1998): "ArguMed: a template-based argument mediation system for lawyers." In: J. C. Hage, T. J. M. Bench-Capon, A. W. Koers, C. N. J. de Vey Mestdagh and C. A. F. M. Grutters (Editors): *Proceedings of the Eleventh Conference of Legal Knowledge-Based Systems (JURIX-98)*, pp. 113—130. Nijmegen, The Netherlands: Gerard Noodt Instituut.
- Verheij, B. (1999): "Automated argument assistance for lawyers." In: *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pp. 43—52. New York City, NY, USA: ACM Press.
- Walton, D. N. (1995): *A Pragmatic Theory of Fallacy*. Tuscaloosa: University of Alabama Press.
- Walton, D. N. (2000): "The place of dialogue theory in logic, computer science and communication studies", *Synthese*, 123: 327-346.
- Walton, D. N. and E. C. W. Krabbe (1995): *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Series in Logic and Language. Albany, NY, USA: State University of New York Press.
- Weizenbaum, J. (1976): *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: Freeman.