Dialogue Game Protocols

Peter McBurney[1] and Simon Parsons[2]

[1] Department of Computer Science
University of Liverpool
Liverpool L69 7ZF U. K.
p.j.mcburney@csc.liv.ac.uk
[2] Center for Co-ordination Science
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02142, USA
sparsons@mit.edu

Abstract. Formal dialogue games have been studied in philosophy since at least the time of Aristotle. Recently they have been used as the basis for agent interaction protocols. We review these applications and examine some of the protocols which have been proposed for agent interactions. We discuss the many open questions and challenges, including issues of automatability of agent dialogues, and the semantics of protocols.

1 Introduction

Dialogue game protocols for agent interactions lie between auction mechanisms and generic agent communications languages. Dialogue game protocols are more expressive than auction and game-theoretic mechanisms, typically allowing participants to question and contest assertions, to advance supporting arguments and counter-arguments, and to retract prior assertions, none of which are normally possible under auction mechanisms. Because such questions and arguments are likely to increase understanding by the participants of each other's positions, we would expect dialogue game protocols both to increase the chance of a successful resolution of an interaction, and to hasten it along, relative to auction mechanisms. It is for just these reasons that argumentation approaches were proposed for agent communications protocols in [43]. Moreover, in many decision contexts, determination of an agent's preferences between alternative decision-options may well depend upon the preferences of other participants [31], as for example when participants to a public policy decision consider the totality of consequences of each option, not merely those consequences which are of direct personal concern. In such circumstances, agents do not necessarily enter the interaction with their utilities and preferences fully formed; rather, their utilities and preferences are generated in the very act of interaction itself [47]. Auction mechanisms, because of their impoverished expressiveness, are particularly unsuitable for such "on-the-fly" preference-generation. Marketing theorists, modeling real-world consumer decision-making, have known this for a long time

[28]; political theorists, now viewing democratic decision-making as a process of deliberation rather than one of economic choice, have also realized it [9].

At the other extreme, dialogue game protocols usually do not permit absolutely any contributions whatsoever, and participants may be required to respond to particular utterances in particular ways. Dialogues conducted according to such protocols are therefore more constrained than interactions using generic agent communications languages, such as FIPA's ACL [18]. In addition to allowing belligerent, malicious or badly-coded participants to disrupt conversations, the freedom of generic languages such as FIPA ACL also complicates the task of analysis of utterances and conversations. After all, any given sequence of utterances in the FIPA ACL may be followed by any one of the 22 locutions, and each of them in turn followed by one of 22 locutions, and so on. This creates a state-explosion problem for participants analysing a sequence of utterances, for example, to decide what locution to utter next, or seeking to infer the future utterances of other participants. One attempt to deal with this problem has sought to identify short sequences of utterances, called *conversation policies*, which have some common intention, and with rules connecting the ordering of locutions within them [20, 54]. Thus, a request for information may be followed only by the provision of that information (via an *inform* locution in FIPA ACL) or a statement that the information is not known (also via an *inform* locution), and not any other statement. Conversation policies therefore sit between individual utterances and entire conversations; in particular, each participant may maintain several different conversation policies concurrently, implementing them as and when required, and these policies may differ from the policies of other participants in the same dialogue. Such an approach creates problems of global coherence across the entire conversation [16]. It also distinguishes conversational policies from formal dialogue games, whose rules apply to all participants in the interaction, and to all dialogues conducted under the protocol.

Formal dialogue games are interactions between two or more players, where each player "moves" by making utterances, according to a defined set of rules. Although their study dates at least from Aristotle [5], they have found recent application in philosophy, computational linguistics and Artificial Intelligence (AI). In philosophy, dialogue games have been used to study fallacious reasoning [21, 30] and to develop a game-theoretic semantics for intuitionistic and classical logic [29] and quantum logic [41]. In linguistics, they have been used to explain sequences of human utterances [27], with subsequent application to machine-based natural language processing and generation [24], and to human-computer interaction [7]. Within computer science and AI, they have been applied to modeling complex human reasoning, for example in legal domains [8, 45], and to requirements specification for complex software projects [17]. Dialogue games differ from the games of economic game theory [42] in that payoffs for winning or losing a game are not considered, and because there is no use of uncertainty measures, such as probabilities, to model the possible moves of opponents. They also differ from the abstract games recently used as a semantics for interactive computation [1], since these latter games do not share the rich rule structure of

dialogue games, nor are they intended to have themselves a semantic interpretation involving the co-ordination of actions among a group of agents.

This chapter reviews the application of formal dialogue games to the design of agent interaction protocols. We begin, in Section 2, with a brief overview of an influential typology of human dialogues, which will be useful in classifying agent interactions. Section 3 then presents a model of a formal dialogue game protocol, following which we consider, in Section 4, several recent proposals for agent interaction protocols based on dialogue games. In Section 5, we discuss some of the many open issues and the research and development challenges current in this domain.

2 Types of dialogues

An influential model of human dialogues is the typology of primary dialogue types of argumentation theorists Doug Walton and Erik Krabbe [55]. This categorization is based upon the information the participants have at the commencement of a dialogue (of relevance to the topic of discussion), their individual goals for the dialogue, and the goals they share. **Information-Seeking Dialogues** are those where one participant seeks the answer to some question(s) from another participant, who is believed by the first to know the answer(s). In **Inquiry Dialogues** the participants collaborate to answer some question or questions whose answers are not known to any one participant. **Persuasion Dialogues** involve one participant seeking to persuade another to accept a proposition he or she does not currently endorse. In **Negotiation Dialogues**, the participants bargain over the division of some scarce resource. Here, the goal of the dialogue — a division of the resource acceptable to all — may be in conflict with the individual goals of the participants. Participants of **Deliberation Dialogues** collaborate to decide what action or course of action should be adopted in some situation. Here, participants share a responsibility to decide the course of action, or, at least, they share a willingness to discuss whether they have such a shared responsibility. Note that the best course of action for a group may conflict with the preferences or intentions of each individual member of the group; moreover, no one participant may have all the information required to decide what is best for the group. In **Eristic Dialogues**, participants quarrel verbally as a substitute for physical fighting, aiming to vent perceived grievances.

Most actual dialogue occurrences — both human and agent — involve mixtures of these dialogue types. A purchase transaction, for example, may commence with a request from a potential buyer for information from a seller, proceed to a persuasion dialogue, where the seller seeks to persuade the potential buyer of the importance of some feature of the product, and then transition to a negotiation, where each party offers to give up something he or she desires in return for something else. The two parties may or may not be aware of the different nature of their discussions at each phase, or of the transitions between phases. Instances of individual dialogue types contained entirely within other dialogue types are said to be *embedded* [55].

3 Formal Dialogue Games

We now present a model of a generic formal dialogue game in terms of the components of its specification, taken from [36]. We first assume that the topics of discussion between the agents can be represented in some logical language, whose well-formed formulae are denoted by the lower-case Roman letters, p, q, r, etc. A dialogue game specification then consists of the following elements:

Commencement Rules: Rules which define the circumstances under which the dialogue commences.

Locutions: Rules which indicate what utterances are permitted. Typically, legal locutions permit participants to assert propositions, permit others to question or contest prior assertions, and permit those asserting propositions which are subsequently questioned or contested to justify their assertions. Justifications may involve the presentation of a proof of the proposition or an argument for it. The dialogue game rules may also permit participants to utter propositions to which they assign differing degrees of commitment, for example: one may merely *propose* a proposition, a speech act which entails less commitment than would an *assertion* of the same proposition.

Combination Rules: Rules which define the dialogical contexts under which particular locutions are permitted or not, or obligatory or not. For instance, it may not be permitted for a participant to assert a proposition p and subsequently the proposition $\neg p$ in the same dialogue, without in the interim having retracted the former assertion.

Commitments: Rules which define the circumstances under which participants express commitment to a proposition. Typically, the assertion of a claim p in the debate is defined as indicating to the other participants some level of commitment to, or support for, the claim. Since the work of philosopher Charles Hamblin [21], formal dialogue systems typically establish and maintain public sets of commitments, called *commitment stores*, for each participant; these stores are usually non-monotonic, in the sense that participants can also retract committed claims, although possibly only under defined circumstances.

Termination Rules: Rules that define the circumstances under which the dialogue ends.

It is worth noting here that more than one notion of *commitment* is present in the literature on dialogue games. For example, Hamblin treats commitments in a purely dialogical sense: *"A speaker who is obliged to maintain consistency needs to keep a store of statements representing his previous commitments, and require of each new statement he makes that it may be added without inconsistency to this store. The store represents a kind of persona of beliefs; it need not correspond with his real beliefs . . ."* [21, p. 257]. In contrast, Walton and Krabbe [55, Chapter 1] treat commitments as obligations to (execute, incur or maintain) a course of action, which they term action commitments. These actions may be utterances in a dialogue, as when a speaker is forced to defend a proposition

he has asserted against attack from others; so Walton and Krabbe also consider propositional commitment as a special case of action commitment [55, p. 23]. As with Hamblin's treatment, such dialogical commitments to propositions may not necessarily represent a participant's true beliefs. In contrast, Munindar Singh's social semantics [52], requires participants in an interaction to express publicly their beliefs and intentions, and these expressions are called *social commitments*. These include both expressions of belief in some propositions and expressions of intent to execute or incur some future actions.[1] Our primary motivation is the use of dialogue games as the basis for interaction protocols between autonomous agents. Because such agents will typically enter into these interactions in order to achieve some wider objectives, and not just for the enjoyment of the interaction itself, we believe it is reasonable to define commitments in terms of future actions or propositions external to the dialogue. In a commercial negotiation dialogue, for instance, the utterance of an offer may express a willingness by the speaker to undertake a subsequent transaction on the terms contained in the offer. For this reason, we can view commitments as semantic mappings between locutions and subsets of some set of statements expressing actions or beliefs external to the dialogue.

Dialogue game protocols have been articulated for each of the rule-governed primary types of dialogues in the typology of Walton and Krabbe: information-seeking dialogues [24, 55]; inquiries [35]; persuasion dialogues [2, 11]; negotiation dialogues [4, 24, 33, 49]; and deliberations [23]. Some of these proposals are discussed in the next section. There has even been some work on non-co-operation in dialogues [19], work which may yield dialogue game protocols for eristic dialogues. However, as mentioned earlier, most real-world dialogues (whether human or agent) involve aspects of more than one of these primary types. Two formalisms have been suggested for computational representation of combinations of dialogue: the *Dialogue Frames* of Chris Reed [46], which enable iterated, sequential and embedded dialogues to be represented; and our own *Agent Dialogue Frameworks* [36], which permit iterated, sequential, parallel and embedded dialogues to be represented. Both these formalisms are neutral with regard to the modeling of the primary dialogue types themselves, allowing the primary types to be represented in any convenient form, and allowing for types other than the six of the Walton and Krabbe typology to be included.

4 Examples of dialogue game protocols

We now briefly review several of the proposals for dialogue game protocols published in the agent literature. The first of these is the protocol of Leila Amgoud,

[1] It is worth noting that all these notions of *commitment* differ from that commonly used in discussion of agent's internal states, namely the idea of the persistence of a belief or an intention [58, p. 205]. As Singh [51] argues, there is a qualitative difference between social commitments of the kind discussed here, and personal commitments of the kind encoded in beliefs, desires, and intentions. He further argues that one kind of commitment cannot be derived from another.

Nicolas Maudet and Simon Parsons [2]. This protocol is based on James MacKenzie's philosophical dialogue game *DC* [30], a game for two players, both subject to the same rules. *DC* enables the participants to argue about the truth of a proposition and was designed to study the fallacy of begging the question (*petitio principii*, or circular reasoning). The agent interaction protocol of [2] based on *DC* allows four distinct locutions: *assert, accept, question* and *challenge*; these can be instantiated with a single proposition, and also, for the locutions *assert* and *accept*, with a set of propositions which together constitute an argument for a proposition. Thus the participants may communicate both propositional statements and arguments about these statements, where arguments may be considered as tentative proofs (i.e., logical inferences from assumptions which may not all be confirmed). The locutions of this protocol are similar to those of *DC* except that they do not include a locution for retraction of assertions (called *withdrawal* in *DC*). As with MacKenzie's game, the protocol of [2] establishes commitment stores which record, in full public view, the statements each participant has asserted. The syntax for this protocol has only been provided for dialogues between two participants, but could be readily extended to more agents, as the same authors did subsequently in [3].

Amgoud *et al.* demonstrate that their system enables persuasion, inquiry and information-seeking dialogues [2]. However, as the authors note, to permit negotiation dialogues, the protocol requires additional locutions.[2] These are proposed in a subsequent paper by the same authors [4], in which three additional locutions are proposed, *request, promise* and *refuse*, making seven in all. In addition to instantiation with propositions and with arguments for propositions, several of these locution can also be instantiated with a two-valued function expressing a relationship between two resources. For example, the locution $promise(p \Rightarrow q)$ indicates a promise by the speaker to provide resource q in return for receiving resource p.

Building on the protocol of Amgoud *et al.* [4], Fariba Sadri, Francesca Toni and Paolo Torroni [49] propose a similar protocol but with fewer locutions. The legal locutions proposed here are: *request, promise, accept, refuse, challenge,* and *justify.* The contents of the locutions *request* and *promise* are resources, while the contents for the other four locutions are any of the six locutions. In addition, the locution *justify* allows the utterance of some support for a previous locution.

The dialogue-game protocols presented in the work of Frank Dignum, Barbara Dunin-Kęplicz and Rineke Verbrugge [11, 12] are intended to enable agents to form teams and to agree joint intentions, respectively. For both protocols, the authors assume that one agent, an *initiator* or *proponent*, seeks to persuade others (*opponents*) to join a team, and that another initiator (possibly the same agent) seeks to persuade team members to adopt a group belief or intention. The team-formation dialogue is modeled as an information-seeking dialogue followed by a persuasion, while the joint-intentions-formation dialogue is modeled as a persuasion dialogue, which may include embedded negotiation

[2] It may also require additional locutions for deliberation dialogues, although the authors suggest otherwise.

dialogues. For the persuasion dialogues, the authors adapt the rigorous persuasion dialogue game of Walton and Krabbe [55]. This game is a formalization of a critical discussion — i.e., a rigorous persuasion — in philosophy. Such dialogues involve two parties, one seeking to prove a proposition, and one seeking to disprove it.[3] Unfortunately, the authors do not specify their dialogue game models completely, for example, nowhere stating the set of locutions available to the participating agents. The protocol for joint intention formation dialogues [12] includes seven locutions: *statement, question, challenge, challenge-with-statement, question-with-statement* and *final remarks*; these last include: *"quit"* and *"won."* The statements associated with challenges and questions may be concessions made by the speaker. The protocol for team formation dialogues [11] may also use the same set of locutions, although this is not absolutely clear.

Finally, we briefly mention some of our own proposals for dialogue game agent interaction protocols. Firstly, in joint work with Rogier van Eijk and Leila Amgoud [33], we articulated a dialogue game protocol for negotiation dialogues between potential buyers and sellers of consumer durables; this work drew on a standard model of consumer decision-making. Secondly, together with David Hitchcock [23], we presented a dialogue game protocol for deliberation dialogues, drawing on a theory of deliberative argument from the philosophy of argumentation. Thirdly, in [35], we articulated a dialogue game protocol for agents engaged in an inquiry dialogue; this protocol enables the participants to express uncertain beliefs about claims and to resolve these on the basis of the arguments for and against the claims presented in the dialogue. Inquiries involve a disinterested search for the truthful answer to some question. In many instances, however, we may desire *particular* answers to a question, such as when we seek to identify the possible risks of a new technology. In these cases our search is overlaid with values we impose on the answer-space; in [34], we proposed a dialogue game protocol for agents engaged in such a search.

5 Issues and challenges

The use of formal dialogue games as the basis for agent interaction protocols has only just begun, and there are many challenging issues still open. In this section we consider some of these.

5.1 Protocol semantics

One of the reasons for the popularity of the FIPA ACL is the fact that it has been given a well-defined semantics [26]. This semantics, based on speech act theory from the philosophy of language [6, 50, 10], is defined in terms of the certain beliefs, uncertain beliefs and intentions of the participating agents. Having defined such a semantics means that participants know precisely what other speakers

[3] Note that the persuasion dialogues of Walton and Krabbe [55] deal only with beliefs, not intentions.

mean and intend by their utterances, *assuming those others are conforming to the semantics*. However, verifying conformance to the semantics is a conceptually challenging task [57], since it is always possible for a sufficiently-clever agent to simulate insincerely any required semantic state.

The development of appropriate semantics for dialogue game protocols is still very immature, although several different types of semantics have been defined for these protocols. The first of these involves defining locutions in terms of the pre-conditions necessary for, and the post-conditions arising from, their utterance, i.e., what is termed an *axiomatic* semantics in computer programming language theory [53]. We can distinguish, as in [32], between two types of axiomatic semantics. In a *public* axiomatic semantics, the pre-conditions and post-conditions of each locution are given only in terms of observable linguistic behaviour, and so conformity with the protocol can be verified by anyone observing the dialogue. The protocols in [23, 33–35] have been given a public axiomatic semantics. A *private* axiomatic semantics, on the other hand, is one in which some or all locutions have pre- or post-conditions defined in terms of the internal states of the participants. The protocols of Amgoud *et al.* [2, 4], of Sadri *et al.* [49] and of Dignum *et al.* [12] have been given such a semantics. For the protocols of Amgoud *et al.* [2, 4], each participating agent is assumed to be vested with a private argumentation-based reasoning mechanism, which generates defeasible arguments from premises according to defined procedures, as in [13]. The mechanism also permits a preference ordering over the arguments. An agent can only utter a locution $assert(p)$, for p a proposition, if that agent has an acceptable argument for p in its own knowledge base, or in its knowledge base combined with the public commitment stores of the other participants. Here, *acceptable arguments* are those which survive a specified process of attack and defence within an argumentation framework [13].

In the case of the protocol of Sadri *et al.* [49], which is designed for dialogues over scarce resources (i.e, negotiation dialogues), utterances are linked to a first-order logic describing resources. In this semantics, the knowledge of an agent is described [49, Section 3.1] as an abductive logic program consisting of if-then rules and of statements regarding the resources owned by the agent. Integrity constraints are placed on this knowledge in the form of rules[4] which provide an agent with possible responses to utterances it receives. An example of such a rule is: *Accept a request*. The abducibles of this logic program are then the possible locutions which the agent may utter in response to a message it receives. For the protocol of Dignum *et al.* [12], the authors assume that the participating agents have a Belief-Desire-Intention (BDI) mental architecture [56] and vest the locutions with a private axiomatic semantics: as with the speech-act semantics of the FIPA ACL [18], locutions are defined in terms of their impacts on, and pre-conditions of, these private mental states.

One may also view a dialogue under a given protocol as a sequence of commands in a programming language which operates on the states of a virtual machine comprising the internal states of the agents participating in the dialogue.

[4] Confusingly called *dialogue protocols* by these authors.

This view leads to the definition of an *operational* semantics for the protocol [53] in which locutions are defined in terms of their state transitions. Operational semantics have recently been defined for various agent communications languages, e.g., [15, 22]. To our knowledge, the only dialogue game protocol provided with an operational semantics is the consumer negotiation protocol of [33].

Programming language theory also entertains the concept of a *denotational* semantics [53], in which a translation is given between the commands in a programming language and the objects in some mathematical domain. In [37], we defined a denotational semantics for protocols in terms of sub-spaces of n-dimensional euclidean space; in this semantics, dialogues conducted according to a given protocol are mapped to directed paths in the associated sub-space. Another denotational approach arises from viewing agents engaged in dialogue as jointly constructing meaning as the dialogue proceeds, in the same way that humans using natural language may be thought to do. Thus, there may be value in defining a denotational semantics which is constructed incrementally in the course of a dialogue, in a manner analogous to the Discourse Representation Structures of natural language semantics in theoretical linguistics [25]. This is a line of research we are pursuing.

5.2 Formal properties

Why define a semantics for a protocol? One reason, mentioned above, is so that the participants share common meanings for utterances. Another reason is to enable a better understanding of the formal properties of protocols and of the dialogues conducted under them. The study of the formal properties of dialogues and protocols is, like the development of formal semantics, still very immature, and considerable scope exists for further research in this area.

One property of great interest is *termination*: under what circumstances will dialogues conducted under a given protocol terminate? For instance, Sadri and her colleagues [49] demonstrate that a dialogue under their protocol requesting the use of some resource will always terminate in a finite number of steps, if conducted between agents vested with the abductive logic program mechanisms described in the paper.[5] Similarly, Parsons *et al.* [44] consider the termination properties of the protocols of [2, 4] for agents vested with argumentation-based decision architectures.

A related property is *computational complexity*: how many dialogue utterances are required for normal termination of a dialogue under a given protocol? To our knowledge, the only work considering computational complexity of dialogue game protocols is that of Paul Dunne and Trevor Bench-Capon [14], who consider a particular two-party, persuasion dialogue protocol. The computational complexity of general negotiation mechanisms (not only those involving dialogue game protocols) were presented by one of us in [59].

[5] To some extent this result is not surprising since these mechanisms require the agents to be co-operative.

A third property of importance in practical applications concerns the degree to which a dialogue protocol may support automated dialogues between suitably-equipped software agents, what may be called *automatability*. In [33] we showed, using an operational semantics, that agents vested with appropriate, domain-specific, internal decision mechanisms could undertake fully automated consumer negotiation dialogues using the dialogue game protocol we defined. The internal decision mechanisms were derived from standard models of buyer and seller decision-making taken from marketing theory [28]. No other dialogue game protocol so far proposed has been proven to be automatable.

5.3 Protocol design and assessment

One reason a more thorough study of protocol properties is needed is to provide assistance to designers and users of agent interaction protocols. At present, designers of dialogue game protocols currently have no guidance for issues such as:

- How many locutions should there be?
- What types of locutions should be included? For example: assertions, questions, etc.
- What are the appropriate rules for the combination of locutions?
- When should behavior be forbidden, e.g., repeated utterance of one locution?
- Under what conditions should dialogues be made to terminate?
- When are dialogues conducted according to a particular protocol guaranteed to terminate?

Similarly, the absence of formal studies of dialogue game protocols means that agents (or their human principals) intending to use such protocols have no guidance for issues such as:

- When are two protocols the same? With the proliferation of protocols, it will become important for agents to decide which protocol should be used for particular interactions.
- How may different protocols be compared and how may differences be measured?
- Which protocols are to be preferred under which circumstances? In other words, what are their advantages and disadvantages?
- What is the formal relationship between dialogue game protocols and general agent communications languages, e.g., FIPA ACL.
- When are dialogue game protocols preferable to other forms of agent interaction, such as auction mechanisms or general agent communications languages?

As part of a longer-term project to develop a formal theory of interaction protocols, we have taken initial steps towards answering some of these questions. In work with Michael Wooldridge [39], we proposed thirteen desirable properties

of interaction protocols using dialogue games. These properties included: *Separation of syntax and semantics*, so that conformance with the protocol could be assessed on the basis only of observable behaviour; *Clarity of argumentation theory*, so that participants would know in advance their dialectical obligations when making assertions, contesting others' assertions, etc; and *Enablement of self-transformation*, so that participants would be able to change their beliefs, preferences or intentions, and readily express such changes, in the course of an interaction. We applied these 13 properties to assess several existing dialogue game protocols and to the FIPA ACL; all were found wanting, to a greater or lesser extent.

In another step towards a formal theory of interaction protocols we have considered when two protocols may be considered the same. Not only is an answer to this question necessary for choosing between two protocols, but it is also essential in order to assess if a protocol is new (in the sense of providing a different functionality to an existing protocol rather than just having equivalent locutions with different names) or to assess if a protocol conforms to some specification (such as that laid down as the standard for interacting within some electronic institution [48]). In work with Mark Johnson [38], we have recently defined several reasonable notions of protocol equivalence and shown that these lead to distinct classes of protocols. Some of these notions were derived from ideas of process equivalence [40], and the links between protocols and processes are worth further exploration.

6 Conclusion

In this chapter, we have provided a brief overview of agent interaction protocols based on formal dialogue games. These games have almost 2500 years of study in philosophy behind them, and have recently found application in the design of autonomous software agent protocols. Only a handful of such protocols have so far been proposed, and only a few of their properties yet studied formally. This area of Artificial Intelligence therefore still has open many questions, research challenges and implementation issues. We believe that dialogue game protocols have great potential for multi-agent systems applications because they represent an effective compromise between the strict rule-governed nature of economic auction mechanisms and the greater expressiveness of generic agent communications languages. Proving this claim conclusively will need to await a formal theory of interaction protocols connecting dialogue game protocols, auction mechanisms and agent communications languages in a single coherent framework.

References

1. S. Abramsky. Semantics of interaction: an introduction to game semantics. In A. M. Pitts and P. Dybjer, editors, *Semantics and Logics of Computation*, pages 1–31. Cambridge University Press, Cambridge, UK, 1997.

2. L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, pages 31–38, Boston, MA, USA, 2000. IEEE Press.
3. L. Amgoud and S. Parsons. Agent dialogues with conflicting preferences. In J-J. Meyer and M. Tambe, editors, *Pre-Proceedings of the Eighth International Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*, pages 1–14, Seattle, WA, USA, 2001.
4. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, pages 338–342, Berlin, Germany, 2000. IOS Press.
5. Aristotle. *Topics.* Clarendon Press, Oxford, UK, 1928. W. D. Ross, Editor.
6. J. L. Austin. *How To Do Things with Words.* Oxford University Press, Oxford, UK, 1962.
7. T. J. M. Bench-Capon, P. E. Dunne, and P. H. Leng. Interacting with knowledge-based systems through dialogue games. In *Proceedings of the Eleventh International Conference on Expert Systems and Applications*, pages 123–140, Avignon, France, 1991.
8. T. J. M. Bench-Capon, T. Geldard, and P. H. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8:233–254, 2000.
9. J. Bohman and W. Rehg, editors. *Deliberative Democracy: Essays on Reason and Politics.* MIT Press, Cambridge, MA, USA, 1997.
10. P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.
11. F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Agent theory for team formation by dialogue. In C. Castelfranchi and Y. Lespérance, editors, *Intelligent Agents VII: Proceedings of the Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL 2000)*, LNAI 1986, pages 150–166, Berlin, Germany, 2000. Springer.
12. F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Creating collective intention through dialogue. *Logic Journal of the IGPL*, 9(2):305–319, 2001.
13. P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-persons games. *Artificial Intelligence*, 77:321–357, 1995.
14. P. E. Dunne and T. J. M. Bench-Capon. Two party immediate response disputes: properties and efficiency. Technical Report ULCS-01-005, Department of Computer Science, University of Liverpool, Liverpool, UK, 2001.
15. R. van Eijk, F. S. de Boer, W. van der Hoek, and J-J. Ch. Meyer. A fully abstract model for the exchange of information in multi-agent systems. *Theoretical Computer Science*, April 2002. *In press.*
16. R. Elio, A. Haddadi, and A. Singh. Task models, intentions, and agent conversation policies. In R. Mizoguchi and J. K. Slaney, editors, *Topics in Artificial Intelligence: Proceedings of the Sixth Pacific Rim Conference in Artificial Intelligence (PRICAI-2000)*, LNAI 1886, pages 394–403, Berlin, Germany, 2000. Springer.
17. A. Finkelstein and H. Fuks. Multi-party specification. In *Proceedings of the Fifth International Workshop on Software Specification and Design*, Pittsburgh, PA, USA, 1989. ACM Sigsoft Engineering Notes.
18. FIPA. Communicative Act Library Specification. Technical Report XC00037H, Foundation for Intelligent Physical Agents, 10 August 2001.
19. D. M. Gabbay and J. Woods. Non-cooperation in dialogue logic. *Synthese*, 127(1-2):161–186, 2001.

20. M. Greaves, H. Holmback, and J. Bradshaw. What is a conversation policy? In F. Dignum and M. Greaves, editors, *Issues in Agent Communication*, LNAI 1916, pages 118–131. Springer, Berlin, Germany, 2000.
21. C. L. Hamblin. *Fallacies.* Methuen, London, UK, 1970.
22. K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J-J. Ch. Meyer. Formal semantics for an abstract agent progamming language. In M. P. Singh, A. S. Rao, and M. J. Wooldridge, editors, *Intelligent Agents IV: Proceedings of the Fourth International Workshop on Agent Theories, Architectures and Languages (ATAL-97)*, LNAI 1365, pages 215–229. Springer, Berlin, Germany, 1998.
23. D. Hitchcock, P. McBurney, and S. Parsons. 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, 2001.
24. J. Hulstijn. *Dialogue Models for Inquiry and Transaction.* PhD thesis, Universiteit Twente, Enschede, The Netherlands, 2000.
25. H. Kamp and U. Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory.* Kluwer, Dordrecht, The Netherlands, 1993. Two Volumes.
26. Y. Labrou, T. Finin, and Y. Peng. Agent communication languages: The current landscape. *IEEE Intelligent Systems*, 14(2):45–52, March/April 1999.
27. J. A. Levin and J. A. Moore. Dialogue-games: metacommunications structures for natural language interaction. *Cognitive Science*, 1(4):395–420, 1978.
28. G. L. Lilien, P. Kotler, and K. S. Moorthy. *Marketing Models.* Prentice-Hall, Englewood Cliffs, NJ, USA, 1992.
29. P. Lorenzen and K. Lorenz. *Dialogische Logik.* Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1978.
30. J. D. MacKenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.
31. M. Mandler. A difficult choice in preference theory: rationality implies completeness or transitivity but not both. In Elijah Millgram, editor, *Varieties of Practical Reasoning*, pages 373–402. MIT Press, Cambridge, MA, USA, 2001.
32. P. McBurney. *Rational Interaction.* PhD thesis, Department of Computer Science, University of Liverpool, Liverpool, UK, 2002.
33. P. McBurney, R. M. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems*, 2002. *In press.*
34. P. McBurney and S. Parsons. Chance discovery using dialectical argumentation. In T. Terano, T. Nishida, A. Namatame, S. Tsumoto, Y. Ohsawa, and T. Washio, editors, *New Frontiers in Artificial Intelligence: Joint JSAI 2001 Workshop Post Proceedings*, LNAI 2253, pages 414–424. Springer, Berlin, Germany, 2001.
35. P. McBurney and S. Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, 32(1–4):125–169, 2001.
36. P. McBurney and S. Parsons. Games that agents play: a formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.
37. P. McBurney and S. Parsons. A geometric semantics for dialogue-game protocols for autonomous agent interactions. *Electronic Notes in Theoretical Computer Science*, 52(2), 2002.

38. P. McBurney, S. Parsons, and M. W. Johnson. When are two protocols the same? In M. P. Huget, F. Dignum, and J. L. Koning, editors, *Agent Communications Languages and Conversation Policies. Proceedings of a Workshop at the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, Bologna, Italy, 2002.
39. P. McBurney, S. Parsons, and M. Wooldridge. Desiderata for agent argumentation protocols. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, Bologna, Italy, 2002.
40. R. Milner. *Communication and Concurrency.* Prentice Hall, Hemel Hempstead, UK, 1989.
41. P. Mittelstaedt. *Quantum Logic.* D. Reidel, Dordrecht, The Netherlands, 1979.
42. M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, Cambridge, MA, USA, 1994.
43. S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
44. S. Parsons, M. Wooldridge, and L. Amgoud. An analysis of formal interagent dialogues. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, Bologna, Italy, 2002. (*In press*).
45. H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.
46. C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, pages 246–253. IEEE Press, 1998.
47. W. Rehg. The argumentation theorist in deliberative democracy. *Controversia*, 2002. *(In press).*
48. J. A. Rodríguez, F. J. Martin, P. Noriega, P. Garcia, and C. Sierra. Towards a test-bed for trading agents in electronic auction markets. *AI Communications*, 11(1):5–19, 1998.
49. F. Sadri, F. Toni, and P. Torroni. Logic agents, dialogues and negotiation: an abductive approach. In M. Schroeder and K. Stathis, editors, *Proceedings of the Symposium on Information Agents for E-Commerce, Artificial Intelligence and the Simulation of Behaviour Conference (AISB-2001)*, York, UK, 2001. AISB.
50. J. Searle. *Speech Acts: An Essay in the Philosophy of Language.* Cambridge University Press, Cambridge, UK, 1969.
51. M. P. Singh. A conceptual analysis of commitments in multiagent systems. Technical Report 96-09, Department of Computer Science, North Carolina State University, Raleigh, NC, USA, 1996.
52. M. P. Singh. A social semantics for agent communications languages. In F. Dignum, B. Chaib-draa, and H. Weigand, editors, *Proceedings of the Workshop on Agent Communication Languages, International Joint Conference on Artificial Intelligence (IJCAI-99)*, Berlin, Germany, 2000. Springer.
53. R. D. Tennent. *Semantics of Programming Languages.* Prentice-Hall, Hemel Hempstead, UK, 1991.
54. A. Unruh and M. Nodine. Industrial-strength conversations. In F. Dignum and M. Greaves, editors, *Issues in Agent Communication*, LNAI 1916. Springer, Berlin, Germany, 2000.
55. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning.* SUNY Press, Albany, NY, USA, 1995.

56. M. Wooldridge. *Reasoning about Rational Agents.* MIT Press, Cambridge, MA, USA, 2000.
57. M. Wooldridge. Semantic issues in the verification of agent communication languages. *Journal of Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.
58. M. Wooldridge. *Introduction to Multiagent Systems.* John Wiley and Sons, New York, NY, USA, 2002.
59. M. Wooldridge and S. Parsons. Languages for negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, pages 393–397, Berlin, Germany, 2000. IOS Press.