# Chapter 13
# Dialogue Games for Agent Argumentation

Peter McBurney and Simon Parsons

## 1 Introduction

The rise of the Internet and the growth of distributed computing have led to a major paradigm shift in software engineering and computer science. Until recently, the notion of *computation* has been variously construed as numerical calculation, as information processing, or as intelligent symbol analysis, but increasingly, it is now viewed as distributed cognition and interaction between intelligent entities [60]. This new view has major implications for the conceptualization, design, engineering and control of software systems, most profoundly expressed in the concept of systems of intelligent software agents, or multi-agent systems [99]. Agents are software entities with control over their own execution; the design of such agents, and of multi-agent systems of them, presents major research and software engineering challenges to computer scientists.

One key challenge is the design of means of communication between intelligent agents. Considerable research effort has been expended on the design of artificial languages for agent communications, such as DARPA's *Knowledge Query and Manipulation Language (KQML)* [33] and the Foundation for Intelligent Physical Agents' (now IEEE FIPA) *Agent Communications Language (FIPA ACL)* [35]. These languages, and languages like them, have been designed to be widely applicable. As well as being a strength, this feature can also be a weakness: agents participating in conversations have too many choices of what to utter at each turn, and thus agent dialogues may endure a state-space explosion.

Allowing sufficient flexibility of expression while avoiding state-space explosion had led agent communications researchers to the study of formal dialogue games; these are rule-governed interactions between two or more players (or agents), where

---

Department of Computer Science, University of Liverpool, UK
e-mail: `mcburney@liverpool.ac.uk`
Department of Computer and Information Science, Brooklyn College, New York, USA
e-mail: `parsons@sci.brooklyn.cuny.edu`

I. Rahwan, G. R. Simari (eds.), *Argumentation in Artificial Intelligence,*
DOI 10.1007/978-0-387-98197-0_13, © Springer Science+Business Media, LLC 2009

each player "moves" by making utterances, according to a defined set of rules. Although their study dates to at least the time of Aristotle [5], dialogue games have found recent application in philosophy, computational linguistics and Artificial Intelligence (AI). In philosophy, dialogue games have been used to study fallacious reasoning [41, 62] and to develop a game-theoretic semantics for various logics, e.g., intuitionistic and for classical logics [59]. In linguistics, they have been used to explain sequences of human utterances [57], with subsequent application to machine-based natural language processing and generation [49], and to human-computer interaction [9]. Within computer science and AI, they have been applied to modeling complex human reasoning, for example in legal domains [81], and to requirements specification for complex software systems [34]. Dialogue games differ from the games of economic game theory in that payoffs for winning or losing a game are not considered, and, indeed, the notions of *winning* and *losing* are not always applicable to dialogue games. 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 these latter intended to themselves have a semantic interpretation involving the co-ordination of actions among a group of agents.

This chapter considers the application of formal dialogue games for agent communication and interaction using argumentation. We begin, in the next subsection, with a brief overview of an influential typology of human dialogues, which have proven useful in classifying agent interactions. Because the design of artificial languages for communication between software agents shares much with the study of natural human languages, we structure this chapter according to the standard division within linguistic theory between syntax, semantics and pragmatics; we do this despite this division being imprecise and contested within linguistics (e.g., [58]). Very broadly (following [58]), we may view: *syntax* as being concerned with the surface form and combinatorial properties of utterances, words and their components; *semantics* as being concerned with the truth or falsity of utterances; and *pragmatics* as being concerned with those aspects of the meaning of utterances other than their truth or falsity.[1] Section 2 thus presents a model of a formal dialogue game protocol, focusing primarily on the syntax of such dialogues. We follow this in Section 3 with a discussion of the semantics and the pragmatics of agent dialogues. Section 4 then presents an illustrative example, taken from [68], while Section 5 considers protocol design and assessment. The chapter ends with a brief conclusion in Section 6.

### 1.1 Types of dialogues

An influential model of human dialogues is the typology of primary dialogue types of argumentation theorists Douglas Walton and Erik Krabbe [96]. This categorization is based upon the information the participants have at the commencement of a

---

[1] Note that the word *semantics* is used differently here than in the study of argumentation frameworks, as in Chapter 2 of this volume.

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. If a negotiation dialogue terminates with an agreement, then the resource has been divided in a manner acceptable to all 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. Participants may have only partial or conflicting information, and conflicting preferences. As with negotiation dialogues, if a deliberation dialogue terminates with an agreement, then the participants have decided on a mutually-acceptable course of action. In **Eristic Dialogues**, participants quarrel verbally as a substitute for physical fighting, aiming to vent perceived grievances.

Several comments are important to make here. The first is that although Walton and Krabbe talk about the *goal* of a dialogue and the *goal* of a dialogue type,[2] only participants can have goals since only they are sentient. Participants may believe that a dialogue interaction they enter has an ostensible purpose, but their own goals or the goals of the other participants may not be consistent with this purpose. For example, participants may enter a negotiation dialogue in order to reach an agreement (a *deal*) over the allocation of some resource; or they may enter it to prevent any such agreement being reached, or to delay agreement [24], or to prove that no such agreement is possible, or to gather information from the other participants, or even to signal something to some third party, not in the dialogue. Participants in dialogues may also seek to hide their true goals from the other participants [25, 64]. Instead of dialogue goals it makes sense only to speak of participant goals and dialogue outcomes [74].

Secondly, 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* [96]. Several formalisms have been suggested for computational representation of combinations of dialogue: first, the *Dialogue Frames* of Reed [84], which enable iterated, sequential and embedded dialogues to be represented; sec-

---

[2] as do others, e.g., [80].

ond, the *Agent Dialogue Frameworks* of McBurney and Parsons, based on PDL [66], which permit iterated, sequential, parallel and embedded dialogues to be represented; and third, the more abstract $\mathcal{RASA}$ frameworks of Miller and McBurney [73], which permit iterated, sequential, parallel and embedded combinations of any types of agent interaction protocols. All 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. Walton and Krabbe do not claim their typology is comprehensive, and some recent research has explored other types and sub-types, e.g., [14].

Researchers in multi-agent systems and in argumentation have articulated dialogue game protocols for many of the types in the Walton and Krabbe typology. For example, the two-party protocol of Amgoud, Maudet and Parsons [3], which is based on MacKenzie's philosophical dialogue game *DC* [62], supports persuasion, inquiry and information-seeking dialogues; a subsequent extension of this protocol with additional locutions supports negotiation dialogues [4]. Information-seeking dialogues have been considered by Hulstijn [49], and analyzed by Cogan, Parsons and McBurney [14]; indeed this latter work, which examines the pre- and post-conditions of dialogues over beliefs in fine detail, identifies several new types of dialogues not explicitly included in the Walton and Krabbe typology. A study of different persuasion protocols can be found in the review paper by Prakken [80]; other protocols for persuasion dialogues include the PADUA protocol for arguments from experience by Wardeh, Bench-Capon and Coenen [97] and a protocol for arguments over access to information by Doutre and colleagues [22, 23, 78].

Protocols for multi-agent inquiry dialogues have been proposed and studied by McBurney and Parsons [65], who consider the circumstances under which an inquiry dialogue may converge to the truth, and by Black and Hunter [11], whose agent reasoning architecture enables generative inquiry dialogues, i.e., those where new proposals may emerge for consideration and possible endorsement by the agents participating. For dialogues over beliefs (information-seeking, inquiry and persuasion dialogues), Parsons and Sklar consider the question of convergence of beliefs of agents engaged in repeated dialogues with one another [77]. In addition to [4] cited above, protocols for negotiation dialogues include those of Sadri, Toni and Torroni [87], McBurney, van Eijk, Parsons and Amgoud [63], and Karunatillake [54]. Regarding dialogues over action which are not negotiations: McBurney, Hitchcock and Parsons [64] and Tang and Parsons [92] have presented protocols for deliberation dialogues; Atkinson, Bench-Capon and McBurney have given a representation for proposals for actions and a dialogue game protocol to discuss these proposals [6]; and Atkinson, Girle, McBurney and Parsons have presented a dialogue game protocol for dialogues over commands [7]. Finally, the dialogue-game protocols presented in the work of Dignum, Dunin-Kęplicz and Verbrugge [20, 21] are intended to enable agents to form teams and to agree joint intentions, respectively.

## 2 Syntax

The syntax of a language concerns the surface form of words and phrases, and how these may be combined. Accordingly, defining the syntax of an agent dialogue game protocol usually involves the specification of the possible utterances which agents can make (the locutions) and the rules which govern the order in which utterances can be made. Since the work of Hamblin [41], it has become standard to talk of speakers in a dialogue incurring *commitments*: a speaker who asserts a statement as being true, for example, may be committed to justifying this assertion when challenged by another participant, or else allowed (or even forced) to retract the assertion. Although such dialogical commitments may be viewed as aspects of the semantics (the meaning) of utterances, the rules regarding commitments are typically included in the specification of dialogue syntax because these rules often influence the order of utterances. The various commitments of the participants are usually tracked in a publicly-readable database, called a *commitment store*.

Within the agents communications community, it has become standard to view utterances as composed of two layers: an inner layer comprising the topics of discussion, and an outer (or wrapper) layer, comprising the locutions. An utterance can thus be seen as an instantiated locution, with one variable of instantiation being the topic. This structure, adopted for both *KQML* and *FIPA ACL*, provides great flexibility, since agents encoded appropriately may use the same wrappers to undertake dialogues over different topics.

We now present a generic framework for specification of a dialogue game protocol in terms of its key components, adapted from [66].[3] We first assume that the topics of discussion between the agents (the inner layer) 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 comprises the following elements, each of which concern the wrapper layer of communications:

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.

Rules for Combination of Locutions: Rules which define the dialogical contexts under which particular locutions are permitted or not, or obligatory or not. For

---

[3] We are also informed by [80]; note, however, that work defines a mathematical model for analyzing multi-party dialogues, rather than defining a framework for specification of dialogue protocols for agent communications.

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 incur dialogical commitments by their utterances, and thus alter the contents of the participants' associated commitment stores. For example, a question posed by one agent to another may impose a commitment on the second to provide a response; until provided, this commitment remains undischarged.

Rules for Combination of Commitments: Rules which define how commitments are combined or manipulated when utterances incurring conflicting or complementary commitments are made. For example, the rules may allow a speaker to assert the truth of a proposition and then to assert its negation, with the commitment store holding only the most recent asserted proposition, or the store may hold the earlier proposition until explicitly retracted. These rules become particularly important when multiple dialogues are involved, as when one dialogue is embedded within another; in such a case, the commitments incurred in the inner dialogue may take priority over those of the outer dialogue, or vice versa [66].

Rules for Speaker Order: Rules which define the order in which speakers may make utterances. It may be that any speaker may speak at any time, as in *FIPA ACL*, or that there are rules regarding turn-taking.

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 . . ."* [41, p. 257]. In contrast, Walton and Krabbe [96, 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 [96, p. 23]. As with Hamblin's treatment, such dialogical commitments to propositions may not necessarily represent a participant's true beliefs. In contrast, Singh's social semantics [90], 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.[4] 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

---

[4] 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 [99, p. 205].

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.

# 3 Semantics and Pragmatics

## *3.1 Purposes of Semantics*

We begin this section by discussing the concept of semantics for agent communications languages and dialogue protocols. These languages and protocols are clearly media for communication (between software entities and/or their human principals) and so researchers have naturally looked to theories developed in human linguistics to understand them. But, unlike human languages, agent communications languages and dialogue protocols are also *formal* constructs, usually defined explicitly and often computationally; thus, understanding their properties can also usefully draw on notions from logic and mathematics. Moreover, because these communications languages and dialogue protocols are usually intended to be used by autonomous software entities, they are also programming languages, since software agents will use them to construct sequences of utterances — commands — with which to interact with one another. The theory of programming language semantics is therefore also relevant to their study.

It is thus important to keep in mind the different functions which a semantics for an agent communications language or dialogue protocol may be required to serve:

- To provide a shared understanding to participants in a communicative interaction of the meaning of individual utterances, of sequences of utterances, and of dialogues.
- To provide a shared understanding to designers of agent protocols and to the (possibly distinct) designers of agents using those protocols of the meaning of individual utterances, of sequences of utterances, and of dialogues.
- To provide a means by which the properties of individual agent communications languages and protocols may be studied formally and with rigor.
- To provide a means by which different agent communications languages and protocols may be compared with one another formally and with rigor.
- To provide a means by which languages and protocols may be readily implemented in production systems.
- To help ensure that implementation of agent communications in open, distributed agent systems is undertaken uniformly.

Different types of semantics may serve these various purposes to varying degrees, and so it may be useful to develop more than one semantics for a communications language or protocol. [5] In addition, an articulation of semantics could be undertaken at one or more different levels: for each individual utterance, or speech act; for specified short sequences of utterances,[6] such as a question-and-answer sequence; for complete sequences of utterances, or *dialogues*; and for dialogue *protocols*. Most current published work on agent dialogue protocols presents a semantics defined in terms of individual utterances. In the terms defined below, these semantics are most often axiomatic or operational, and are much less often denotational.

## *3.2 Types of Semantics*

We have inherited two conflicting notions of semantics, one deriving from linguistics and the other from mathematical logic. As linguists normally understand these terms, the syntax of a language is *"the formal relation of signs to one another"* and the semantics of the language *"the relations of signs to the objects to which the signs are applicable"* (Morris [75], cited in [58, p. 1]). Thus, it makes sense to speak of the *truth* of a sign (or of an utterance in a language using such signs), since this indicates that the sign has a relationship to external objects in the world. Within mathematics and mathematical logic, a different understanding of semantics has arisen, beginning with Pieri [79] and Hilbert [43] and first articulated formally by Tarski [94]. In this tradition, a semantics for a formal language is a relationship between that language and a space $\mathcal{M}$ of mathematical structures, called *models*. A statement $S$ in the language specifies a subset $\mathcal{M}(S)$ of $\mathcal{M}$. Such a statement is said to be *true* in a particular model $M_0$ if $M_0 \in \mathcal{M}(S)$. A statement is said to be logically true if it is true in every model, i.e., if $\mathcal{M}(S) = \mathcal{M}$.[7] These two notions of semantics — one linguistic, one mathematical — collide.[8] In particular, in the mathematical framework, benefit may be gained from defining different semantic mathematical structures for the one language; Tarski himself, for instance, defined topological [93] and discrete lattice [71] semantics for propositional logic. The benefits of this are that different semantic frameworks may enable different properties of the language to be studied and may provide different insights. Insight may also be gained by comparing the structures with each other, a subject known as *model theory* or *metamathematics* [48]. But, defining and comparing alternative structures

---

[5] Traditional mathematical communications theory, due to Shannon and Weaver [89], explicitly ignores the semantics of messages, and so provides little guidance to designers, developers or users of agent communications languages and protocols.

[6] These are known as *conversations* in the agent communications literature, e.g., [39].

[7] Note that Tarski only applied his framework to formal, or mathematical, languages, and was skeptical about its applicability to natural language [94, pp. 163–165].

[8] Their first skirmish was the argument between Hilbert and Frege over the meaning of Hilbert's axioms for geometry: Frege took what we are calling a linguistic approach, Hilbert a model-theoretic approach; see [95, pp. 408–412] and [46, pp. 7–10].

in this way makes no sense in the linguistic understanding of semantics: how could a language admit more than one set of relationships to *the truth*?

Agent communications languages and dialogue protocols straddle this divide. Because they are formal languages, insight into their properties can be gained by defining semantic relationships to mathematical structures, and studying these structures. However, because they are also intended as media for communication, just as natural language is, each agent using a particular communications language or protocol will wish to ensure that all users share a common understanding of utterances.[9] To verify that agents have the same understanding — the same semantics — for a communications language ultimately requires some form of inspection of their internal states or, equivalently, their program code. This is a challenging, and perhaps conceptually impossible, undertaking since a sufficiently-clever agent can always simulate insincerely any required internal state.[10] Rather, in this chapter, our use of semantic frameworks differs from that in linguistics: first, as in model theory, semantic structures are a means to understand the properties of a formal agent communications language, and, second, because our focus is on computer systems, these structures are a means to support the engineering of multi-agent systems software and to aid uniformity of implementation when software engineering is undertaken by different development teams.

It is therefore helpful to consider several different types of semantic frameworks for formal languages. In doing so, we draw on the summary of the literature on programming language semantics presented by van Eijk in [29, Section 1.2.2]; however, we make no claims that the typology is comprehensive. One type of semantics defines each locution of a communications language in terms of the pre-conditions which must exist before the locution can be uttered, and possibly also the post-conditions which apply following its utterance, in a STRIPS-like fashion [32]. This is called an **axiomatic** semantics [29, 72]. For agent communications languages and dialogue protocols we distinguish between public and private axiomatic approaches. In **public** axiomatic approaches, the pre-conditions and post-conditions all describe states or conditions of the dialogue which can be observed by all participants. In **private** axiomatic approaches, at least some of the pre- or post-conditions describe states or conditions which are internal to one or more of the participants, and thus not directly observable by the others. For example, the semantic language, *SL*, for the locutions of the Agent Communications Language, *FIPA ACL*, of the Foundation for Intelligent Physical Agents (FIPA), is a private axiomatic semantics of the speech acts of the language, defined in terms of the beliefs, desires and intentions of the participating agents [35]. For example, the *inform* locution in the *FIPA ACL* language, allows one agent, say agent $A$, to tell another agent, say $B$, some proposition $p$. The *FIPA ACL* semantics of *inform* only permits agent $A$ to do this if [35, p. 10]: (a) agent $A$ believes $p$ to be true, (b) agent $A$ intends that agent $B$ believes $p$ to be true, and (c) agent $A$ believes that agent $B$ does not already have a belief about

---

[9] Friedrich Dürrenmatt's novel, *Die Panne*, shows what tragic consequences may follow when participants assign very different meanings to the same conversation [28].

[10] For more on this, see [98].

the truth of $p$.[11] Similarly, the semantics defined for many dialogue game protocols for agent interaction, e.g., [3], are also private axiomatic semantics. In contrast, the semantics provided for dialogue games used for modeling legal reasoning in [10] is a public axiomatic semantics.

A second type of semantics, an **operational semantics**, considers the dialogue locutions as computational instructions which operate successively on the states of some abstract machine. Under this approach, the participating agents and their shared dialogue are viewed conceptually as parts of a large abstract or virtual computer, whose overall state may be altered by the utterance of valid locutions or by internal decision processes of the participants; it is as if these locutions or decisions were commands in some computer programme language acting on the virtual machine.[12] The utterances and agent decision-mechanisms are thus seen as state transition operators, and the operational semantics defines these transitions precisely [29]. This approach to the semantics of agent communications languages makes explicit any link between the internal decision mechanisms of the participating agents and their public utterances to one another. The semantics therefore enables the relationships between the mental states of the participants and the public state of the dialogue to be seen explicitly, and shows how these relationships change as a result of utterances and internal agent decisions. Thus, an operational semantics will typically make assumptions about the internal decision-mechanisms of the agents participating in the interaction; the actual agents engaged in a communicative interaction may not necessarily use the decision-process or realize the mental states assumed. Operational semantics have recently been defined for some agent communications languages, for example, in [30, 44] and for some dialogue protocols, e.g., information-passing interactions [19], negotiation dialogue protocols [54, 63], and a general argumentation protocol [68].

Third, in **denotational semantics**, each element of the language syntax is assigned a relationship to an abstract mathematical entity, its denotation. The possible worlds, or Kripkean, semantics defined for modal logic syntax is an example of such a semantics for a logical language [56]. However, two decades before Kripke's work, a denotational semantics mapping logical formulae to subsets of a topological space was given for the modal logic system S4 [91]. For argumentation systems, three denotational semantics have been provided for the ICRF's Logic of Argumentation *LA* [55]. In the first of these, Ambler [2] articulated a category-theoretic semantics [61] for *LA*, by extending to arguments the Curry-Howard isomorphism, which connects proofs in a deductive logic to the morphisms of a free cartesian closed category. In this semantics, propositions (i.e., premises or claims) correspond to objects in a particular enriched category, and arguments linking propositions to morphisms between the associated objects. A second denotational semantics for *LA*, due to Parsons [76], connects argumentation systems to qualitative probabilistic networks (QPNs). In this semantics, propositions correspond to nodes in a QPN, and arguments linking propositions to edges between the associated nodes. Das [17]

---

[11] Note that condition (a) enforces sincerity on the speaker, which is not necessarily desirable. Also, condition (c) precludes the use of *inform* in authentication dialogues.

[12] This virtual machine is purely a conceptual construct and does not need to exist in reality.

articulated a third denotational semantics for logics of argumentation, based on a Kripkean possible-worlds structure. In this semantics, different arguments are assumed classified according to the degree of support they provide for propositions; these differential degrees of support are translated into separate hyper-relations over the accessibility relations of the Kripke structure.[13]

Perhaps the first example of a denotational semantics for a dialogue protocol was the possible-worlds semantics for question-response interactions defined by Hamblin in 1956 [40]. Although possible-worlds and category-theoretic denotational semantics have a long subsequent history in mathematical linguistics, only recently have denotational semantics been defined for agent dialogue protocols. In [67], McBurney and Parsons articulated a category-theoretic semantics, called a *Trace Semantics*, for a broad class of deliberation dialogue protocols. In this semantics, articulation of proposals for action by participants correspond to the creation of objects in certain categories, while participant preferences between these proposals correspond to the existence of arrows (morphisms) between the corresponding objects. Thus, the semantics is constructed jointly and incrementally by the dialogue participants as the dialogue proceeds, in a manner similar to the natural language semantics of Discourse Representation Theory [53] (which uses possible worlds), or the argumentation graph of Gordon's Pleadings Game [38]. A similar denotational semantics, constructed jointly and incrementally by the participants, is outlined by Atkinson and colleagues in [6], for a dialogue protocol for arguments over proposals for action. In this semantics, the mathematical entities constructed are topoi and maps between them, rather than simply categories.[14]

For the denotational semantics approach to be useful, we must be able to derive the semantic mapping of a compound statement in the language from the semantic mappings of its elements, a property called *compositionality*. This property is not always present; for example, it may be absent if the language contains compound statements with infinite combinations of elements or if compound statements have denotations which differ from the composition of those of their elements, as in Hintikka's Independence-Friendly (IF) Logic [47]. In these cases, a specific type of denotational semantics, **game-theoretic semantics**, has sometimes proven useful [45]. In this semantics, each well-formed statement in the language is associated with a conceptual game between two players, a protagonist and an antagonist. A statement in the language is considered to be *true* when and only when a winning strategy exists for the protagonist in the associated game; a winning strategy for a player is a rule giving that player moves for the game such that executing these moves guarantees the player can win the game, no matter what moves are made by the opposing player. Game semantics have been articulated for propositional and predicate logics [59], linear logic [1], and for probability theory [18], among others.

---

[13] This semantics may be viewed as a form of quantification over possible worlds, of which a more general formalism is that developed subsequently (and independently) by van Eijk and his colleagues to compare network topologies [31].

[14] Topoi are generalizations of the category of sets, and incorporate a categorial analogue of the notion of set membership [37].

What value do these different types of semantics have? Axiomatic semantics show the pre- and post-conditions of individual utterances in a communications interaction. They may also be used to show the pre- and post-conditions of sequences of utterances, or even entire dialogues [14, 74]. Thus, they provide a set of rules to regulate participation in rule-governed interactions. Operational semantics, by showing the state transitions effected by utterances, may be used to identify dialogue states which are not reachable or from which no legal utterance may be made. These semantics can be used, therefore, to demonstrate that termination of dialogues between participants using a particular protocol is or is not possible. Operational semantics also identify which internal agent decision-mechanisms are needed by agents in order to issue and comprehend received utterances. Properties of dialogue protocols may also be demonstrated using denotational semantics. In [67], we used the Deliberation Trace Semantics to generalize a result of Harsanyi [42] regarding the pareto-optimality of deals achieved using Zeuthen's Monotonic Concession Protocol (MCP) [100]. Game semantics have also been used to study the properties of formal argumentation systems and dialogue protocols, such as their computational complexity [26], or the extent of truth-convergence under an inquiry dialogue protocol [65], and to identify acceptable sets of arguments in argument frameworks [13, 51].

## *3.3 Pragmatics*

Following Levinson [58], we view the study of language pragmatics as dealing with those aspects of linguistic meaning not covered by considerations of truth and falsity. Chief among these aspects are the desires and intentions of speakers, and these are usually communicated by means of speech acts, non-propositional utterances intended to or perceived to change the state of the world. Examples of speech acts are utterances in which a speaker proposes that some action be undertaken, or promises to undertake it, or commands another to perform it. Modern speech act theory was initially due to Austin [8] and Searle [88], who classified spoken utterances by their intended and actual effects on the world (including the internal mental states of those hearing the utterances), and developed pre-conditions for those effects to be realized. Drawing on this theory, Bretier, Cohen, Levesque, Perrault and Sadek were able to present pre- and post-conditions for agent utterances in terms of the mental states of the participants [12, 15, 16]. This work formed the basis for the axiomatic Semantic Language *SL* of the *FIPA Agent Communications Language ACL* mentioned above [35]. One of the criticisms made of *FIPA ACL* is that the language does not support argumentation [70]; accordingly, McBurney and Parsons [68] extended this language by defining five additional locutions to enable agents to assert, question, challenge, justify and retract statements with one another. A set of locution-combination rules are given (although any such rules are absent from the specification of *FIPA ACL* itself), along with an axiomatic semantics in the style of

*SL* and an operational semantics. This protocol is discussed in the next Section, as an example of these ideas.

Long before Austin and Searle, Reinach [85] had noted that speech acts typically require endorsement, or *uptake*, from the hearer before changing the state of the world; a speaker may promise a hearer to perform some action, but the speaker is only obligated to act once the promise is accepted by the hearer. Speech acts are thus essentially social activities and cannot, normally, be executed by a lone reasoner: their natural home is a multi-party dialogue. This observation is particularly true for those speech acts for which the speaker does not have power of retraction or revocation. In [69], McBurney and Parsons presented an analysis of the differences in meaning between, for instance, *commands* and *promises*. Once uptaken (i.e., once in effect), a command may only be revoked by the original speaker, whereas a promise may only be revoked by the agent to which the promise was made, not the speaker.

However, capturing such differences in the syntax of utterances can be difficult. For example, the syntactical form of the two utterances:

*I command you to wash the car.*

*I promise you to wash the car.*

is identical, but the illocutionary force, the effects on the world, the nature of the obligation incurred, the identity of the agent with revocation powers, and even the identity of the agent intended by the speaker to perform the action are different. Although formal agent communications languages should be less ambiguous than natural language, an interpretation of the syntax of utterances is required for elimination of any ambiguity in meaning. McBurney and Parsons [69] dealt with this problem by modifying the denotational trace semantics of [67] to map action-utterances to tuples in a partitioned tuple space [36]. The different dialogical powers that participating agents have of issuance, endorsement and revocation for particular types of utterance then correspond to different permissions to write, copy and delete (respectively) tuples from associated sub-spaces of the tuple space.

## 4 Example

As an example of the ideas in this chapter, we present the *Fatio* protocol of [68]. This protocol comprises five locutions which may be added to the 22 locutions of *FIPA ACL*, in order to enable communicating agents to engage in rational argument: *assert, question, challenge, justify* and *retract*. These five locutions are subject to six locution-combination rules, which together encode a particular dialectical argumentation theory. Because these locutions are intended to be complementary to *FIPA ACL*, there are no locutions for commencing or terminating a dialogue.[15] For reasons of space, we only give examples of two of the five legal locutions of *Fatio*:

---

[15] It would be easy to take these from another protocol, such as [63].

F1: assert($P_i, \phi$): A speaker $P_i$ asserts a statement $\phi \in \mathcal{C}$ (a belief, an intention, a social connection, an external commitment, etc). In doing so, $P_i$ creates a dialectical obligation within the dialogue to provide justification for $\phi$ if required subsequently by another participant.

F3: challenge($P_j, P_i, \phi$): A speaker $P_j$ challenges a prior utterance of *assert*($P_i, \phi$) by another participant $P_i$, and seeks a justification for $\phi$. In contrast to a question, with this locution, $P_j$ also creates a dialectical obligation on himself to provide a justification for not asserting $\phi$, for example an argument against $\phi$, if questioned or challenged. Thus, *challenge*($P_j, P_i, \phi$) is a stronger utterance than *question*($P_j, P_i, \phi$).

For illustration, we present two of the six *Fatio* locution-combination rules. Here, $\boldsymbol{\Phi} \vdash^{+} \phi$ indicates that $\boldsymbol{\Phi}$ is an argument in support of $\phi$.

CR2: The utterances *question*($P_j, P_i, \phi$) and *challenge*($P_j, P_i, \phi$) may be made at any time following an utterance of *assert*($P_i, \phi$). Similarly, the utterances *question*($P_j, P_i, \boldsymbol{\Phi}$) and *challenge*($P_j, P_i, \boldsymbol{\Phi}$) may be made at any time following an utterance of *justify*($P_i, \boldsymbol{\Phi} \vdash^{+} \phi$).

CR3: Immediately following an utterance of *question*($P_j, P_i, \phi$) or *challenge*($P_j, P_i, \phi$), the speaker $P_i$ of *assert*($P_i, \phi$) must reply with *justify*($P_i, \boldsymbol{\Phi} \vdash^{+} \phi$), for some $\boldsymbol{\Phi} \in \mathcal{A}$.

In [68], both an axiomatic and an operational semantics for this protocol are articulated. The axiomatic semantics is defined in terms of the beliefs, desires and intentions of the participating agents consistent with the axiomatic semantics *SL* of *FIPA ACL*. For example, the semantics of the locution *assert(.)* is defined as follows, with $B_i\phi$ indicating that "Agent *i* believes that $\phi$ is true", and $D_i\phi$ that "Agent *i* desires that $\phi$ be true."

- *assert*($P_i, \phi$)
  *Pre-conditions:* A speaker $P_i$ desires that each participant $P_j (j \neq i)$, believes that $P_i$ believes the proposition $\phi \in \mathcal{C}$.
  $((P_i, \phi, +) \notin DOS(P_i)) \wedge (\forall j \neq i)(D_i B_j B_i \phi)$.
  *Post-conditions:* Each participant $P_k (k \neq i)$, believes that participant $P_i$ desires that each participant $P_j (j \neq i)$, believe that $P_i$ believes $\phi$.
  $(P_i, \phi, +) \in DOS(P_i) \wedge (\forall k \neq i)(\forall j \neq i)(B_k D_i B_j B_i \phi)$.
  *Dialectical Obligations:* $(P_i, \phi, +)$ is added to $DOS(P_i)$, the Dialectical Obligations Store of speaker $P_i$.

Similarly, the operational semantics for *Fatio* defined in [68] articulates the state-transition effected by an *assert(.)* utterance (here labeled **F1**) on the mental states of, firstly, the agent who uttered it and, secondly, on any agent who heard it.

TR2: $\langle P_i, \mathbf{D1}, \textit{utter-assert}(\phi)\rangle \xrightarrow{\mathbf{F1}} \langle P_i, \mathbf{D5}, \textit{listen}\ \rangle$

TR3: $\langle P_i, \mathbf{D1}, \textit{utter-assert}(\phi)\rangle \xrightarrow{\mathbf{F1}} \langle P_j, \mathbf{D5}, \textit{do-mech(D2)}\ \rangle$

These excerpts from the semantics of *Fatio* are intended simply for illustration. Full details are given with the protocol definition [68].

## 5 Protocol design and assessment

The science and software engineering of agent communications and interactions is still in its infancy. Accordingly, designers of multi-agent dialogue game protocols still have little guidance for design questions such as: How many locutions should there be? What types of locutions should be included, e.g., 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? What are the properties of a proposed protocol?

Similarly, the immaturity of the discipline means that software developers and their agents still lack answers to questions such as: How may different protocols be compared and differences measured? Which protocols are to be preferred under which circumstances? In other words, what are their advantages and disadvantages? How should a system developer (or an agent) choose between two protocols which both support the same type of dialogue, for example, two negotiation protocols? When are dialogue game protocols preferable to other forms of agent interaction, such as auction mechanisms or general agent communications languages, such as *FIPA ACL*?

Some work has been undertaken which would assist with such questions. For example, McBurney, Parsons and Wooldridge [70], proposed thirteen desirable properties of agent interaction protocols using dialogue games, and then applied these properties to assess several dialogue game protocols and *FIPA ACL*; all were found wanting, to a greater or lesser extent. From an empirical perspective, Karunatillake [54] undertook an evaluation of various negotiation protocols. This work used simulation studies to compare performance in negotiation interactions for agents using protocols with and using protocols without argumentation, in order to identify the circumstances under which argumentation-based negotiation was beneficial. From a theoretical perspective, Johnson, McBurney and Parsons [52] defined various measures of protocol equivalence, both syntactical and semantic, and showed the relationship of these measures to one another. Knowing that two protocols are equivalent allows inference about their properties (such as termination), and about their compliance with a given specification, such as that laid down as the standard for interacting within some electronic institution [86].

## 6 Conclusion

In this chapter we have given a brief introduction to the theory of dialogue game protocols for agent interaction and argument, a subject which has become important with the rise of multi-agent systems. We have focused on the syntax and the semantics of these protocols because these topics are important, not only for analysis of protocols, but also for the software engineering specification, design and implementation of agent interaction systems. In only a chapter, there are many topics we do not have space to discuss, for example: the computational complexity of decision-making involved in making utterances, and in deciding whether or not these comply with a protocol, e.g., [27]; strategic issues over what utterances an agent should make and when, under a given protocol [82, 83]; properties of specific protocols, e.g. [3, 11, 87]; experiences arising from implementation [23]; and allowing agents to choose protocols themselves, even at run-time [50, 74]. As can be seen, there are many avenues to explore in this rich and exciting subject.

**Acknowledgements** We are grateful for partial financial support received from the EC *Information Society Technologies* (IST) programme, through project ASPIC (IST-FP6-002307). This work was also partially supported by the US Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

## 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. S. J. Ambler. A categorical approach to the semantics of argumentation. *Mathematical Structures in Computer Science*, 6:167–188, 1996.
3. 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.
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. K. Atkinson, T. Bench-Capon, and P. McBurney. A dialogue game protocol for multi-agent argument for proposals over action. *Autonomous Agents and Multi-Agent Systems*, 11(2):153–171, 2005.
7. K. Atkinson, R. Girle, P. McBurney, and S. Parsons. Command dialogues. In I. Rahwan and P. Moraitis, editors, *Proceedings of the Fifth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2008)*, Lisbon, Portugal, 2008. AAMAS 2008.
8. J. L. Austin. *How To Do Things with Words*. Oxford University Press, Oxford, UK, 1962.
9. 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, 1991.
10. 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.
11. E. Black and A. Hunter. A generative inquiry dialogue system. In M. Huhns, O. Shehory, E. H. Durfee, and M. Yokoo, editors, *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2007)*, Honolulu, Hawaii, USA, 2007. IFAAMAS, ACM Press.

12. P. Bretier and D. Sadek. A rational agent as the kernel of a cooperative spoken dialogue system: Implementing a logical theory of interaction. In J. P. M. *et al.*, editor, *Intelligent Agents III*, Lecture Notes in Artificial Intelligence 1193, pages 189–204. Springer, Berlin, Germany, 1997.
13. C. Cayrol, S. Doutre, and J. Mengin. On decision problems related to the preferred semantics for argumentation frameworks. *Journal of Logic and Computation*, 13(3):377–403, 2003.
14. E. Cogan, S. Parsons, and P. McBurney. New types of inter-agent dialogs. In S. P. *et al.*, editor, *Argumentation in Multi-Agent Systems: Second International Workshop (ArgMAS 2005)*, Lecture Notes in Computer Science 4049, pages 154–168. Springer, Berlin, Germany, 2006.
15. P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. C. *et al.*, editor, *Intentions in Communication*, pages 221–255. MIT Press, Cambridge, MA, USA, 1990.
16. P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.
17. S. Das. How much does an agent believe: An extension of modal epistemic logic. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, Lecture Notes in Artificial Intelligence 1455, pages 415–426. Springer, Berlin, Germany, 1998.
18. A. P. Dawid and V. G. Vovk. Prequential probability: principles and properties. *Bernoulli*, 5:125–162, 1999.
19. F. S. de Boer, R. Eijk, W. v. Hoek, and J.-J. C. Meyer. A fully abstract model for the exchange of information in multi-agent systems. *Theoretical Computer Science*, 290(3):1753–1773, 2003.
20. 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*, Lecture Notes in Artificial Intelligence 1986, pages 150–166, Berlin, Germany, 2000. Springer.
21. F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Creating collective intention through dialogue. *Logic Journal of the IGPL*, 9(2):305–319, 2001.
22. S. Doutre, P. McBurney, and M. Wooldridge. Law-governed Linda as a semantics for agent interaction protocols. In F. D. *et al.*, editor, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, pages 1257–1258, New York City, NY, USA, 2005. ACM Press.
23. S. Doutre, P. McBurney, M. Wooldridge, and W. Barden. Information-seeking agent dialogs with permissions and arguments. Technical Report ULCS-05-010, Department of Computer Science, University of Liverpool, Liverpool, UK, 2005.
24. P. E. Dunne. Prevarication in dispute protocols. In G. Sartor, editor, *Proceedings of the Ninth International Conference on AI and Law (ICAIL-03)*, pages 12–21, New York, NY, USA, 2003. ACM Press.
25. P. E. Dunne. Suspicion of hidden agenda in persuasive argument. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Computational Models of Argument: Proceedings of COMMA 2006*, pages 329–340, Amsterdam, The Netherlands, 2006. IOS Press.
26. P. E. Dunne and T. J. M. Bench-Capon. Two party immediate response disputes: Properties and efficiency. *Artificial Intelligence*, 149(2):221–250, 2003.
27. P. E. Dunne and P. McBurney. Optimal utterances in dialogue protocols. In J. S. R. *et al.*, editor, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, pages 608–615, New York City, NY, USA, 2003. ACM Press.
28. F. Dürrenmatt. *A Dangerous Game*. Jonathan Cape, London, UK, 1960. (Translation by R. and C. Winston of *Die Panne*, published in German in 1956.).
29. R. Eijk. *Programming Languages for Agent Communications*. PhD thesis, Department of Computer Science, Utrecht University, Utrecht, The Netherlands, 2000.
30. R. Eijk, F. S. de Boer, W. v. Hoek, and J.-J. C. Meyer. Operational semantics for agent communications. In F. Dignum and M. Greaves, editors, *Issues in Agent Communications*, Lecture Notes in Artificial Intelligence 1916, pages 80–95. Springer, Berlin, Germany, 2000.
31. R. Eijk, F. S. de Boer, W. v. Hoek, and J.-J. C. Meyer. Modal logic with bounded quantification over worlds. *Journal of Logic and Computation*, 11(5):701–715, 2001.

32. R. E. Fikes and N. J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
33. T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. Bradshaw, editor, *Software Agents*, pages 291–316. MIT Press, Cambridge, USA, 1997.
34. 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.
35. FIPA. Communicative Act Library Specification. Standard SC00037J, IEEE Foundation for Intelligent Physical Agents, 3 December 2002.
36. D. Gelernter. Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.
37. R. Goldblatt. *Topoi: The Categorial Analysis of Logic*. North-Holland, Amsterdam, The Netherlands, 1979.
38. T. F. Gordon. The Pleadings Game: An exercise in computational dialectics. *Artificial Intelligence and Law*, 2:239–292, 1994.
39. M. Greaves, H. Holmback, and J. Bradshaw. What is a conversation policy? In F. Dignum and M. Greaves, editors, *Issues in Agent Communication*, Lecture Notes in Artificial Intelligence 1916, pages 118–131. Springer, Berlin, Germany, 2000.
40. C. L. Hamblin. *Language and the Theory of Information*. Ph.D. thesis, Logic and Scientific Method Programme, University of London, London, UK, 1957.
41. C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.
42. J. C. Harsanyi. Approaches to the bargaining problem before and after the theory of games: a critical discussion of Zeuthen's, Hicks' and Nash's theories. *Econometrica*, 24:144–157, 1956.
43. D. Hilbert. Grundlagen der Geometrie. In *Festschrift zur Feier der Enthüllung des Gauss-Weber-Denkmals in Göttingen*, pages 3–92. Teubner, Leipzig, Germany, 1899.
44. K. V. Hindriks, F. S. de Boer, W. Hoek, and J.-J. C. Meyer. Formal semantics for an abstract agent progamming language. In M. P. S. *et al.*, editor, *Intelligent Agents IV*, Lecture Notes in Artificial Intelligence 1365, pages 215–229. Springer, Berlin, Germany, 1998.
45. J. Hintikka. Language-games for quantifiers. *Americal Philosophical Quarterly Monograph Series 2: Studies in Logical Theory*, pages 46–72, 1968. Blackwell, Oxford, UK.
46. J. Hintikka. On the development of the model-theoretic viewpoint in logical theory. *Synthese*, 77(1):1–36, 1988.
47. J. Hintikka and G. Sandu. Game-theoretical semantics. In J. Benthem and A. Meulen, editors, *Handbook of Logic and Language*, pages 361–410. Elsevier, Amsterdam, The Netherlands, 1997.
48. W. Hodges. *A Shorter Model Theory*. Cambridge University Press, Cambridge, UK, 1997.
49. J. Hulstijn. *Dialogue Models for Inquiry and Transaction*. PhD thesis, Universiteit Twente, Enschede, The Netherlands, 2000.
50. J. Hulstijn, M. Dastani, and L. Torre. Negotiation protocols and dialogue games. In *Proceedings of the Belgian-Dutch AI Conference (BNAIC-2000), ADDRESS =*, 2000.
51. H. Jakobovits and D. Vermeir. Dialectic semantics for argumentation frameworks. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, pages 63–72, New York, NY, USA, 1999. ACM Press.
52. M. W. Johnson, P. McBurney, and S. Parsons. When are two protocols the same? In M.-P. Huget, editor, *Communication in Multi-Agent Systems: Agent Communication Languages and Conversation Policies*, Lecture Notes in Artificial Intelligence 2650, pages 253–268. Springer, Berlin, Germany, 2003.
53. H. Kamp and U. Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht, 1993.
54. N. C. Karunatillake. *Argumentation-Based Negotiation in a Social Context*. Ph.D. thesis, School of Electronics and Computer Science, University of Southampton, UK, 2006.
55. P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11 (1):113–131, 1995.

56. S. Kripke. A completeness proof in modal logic. *Journal of Symbolic Logic*, 24:1–14, 1959.
57. J. A. Levin and J. A. Moore. Dialogue-games: metacommunications structures for natural language interaction. *Cognitive Science*, 1(4):395–420, 1978.
58. S. C. Levinson. *Pragmatics*. Cambridge University Press, Cambridge, UK, 1983.
59. P. Lorenzen and K. Lorenz. *Dialogische Logik*. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1978.
60. M. Luck, P. McBurney, O. Shehory, and S. Willmott. *Agent Technology: Computing as Interaction. A Roadmap for Agent Based Computing*. AgentLink III, the European Co-ordination Action for Agent-Based Computing, Southampton, UK, 2005.
61. S. Mac Lane. *Categories for the Working Mathematician*. Springer, New York, USA, 1971.
62. J. D. MacKenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.
63. P. McBurney, R. Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems*, 7(3):235–273, 2003.
64. P. McBurney, D. Hitchcock, and S. Parsons. The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems*, 22(1):95–132, 2007.
65. 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.
66. 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.
67. P. McBurney and S. Parsons. A denotational semantics for deliberation dialogues. In I. Rahwan, P. Moraitis, and C. Reed, editors, *Argumentation in Multi-Agent Systems*, Lecture Notes in Artificial Intelligence 3366, pages 162–175. Springer, Berlin, 2005.
68. P. McBurney and S. Parsons. Locutions for argumentation in agent interaction protocols. In R. M. van Eijk *et al.*, editor, *Developments in Agent Communication*, Lecture Notes in Artificial Intelligence 3396, pages 209–225. Springer, Berlin, Germany, 2005.
69. P. McBurney and S. Parsons. Retraction and revocation in agent deliberation dialogs. *Argumentation*, 21(3):269–289, 2007.
70. 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)*, pages 402–409, New York City, NY, USA, 2002. ACM Press.
71. J. C. C. McKinsey and A. Tarski. Some theorems about the Sentential Calculus of Lewis and Heyting. *Journal of Symbolic Logic*, 13(1):1–15, 1948.
72. B. Meyer. *Introduction to the Theory of Programming Languages*. International Series in Computer Science. Prentice Hall, New York City, NY, USA, 1990.
73. T. Miller and P. McBurney. Using constraints and process algebra for specification of first-class agent interaction protocols. In G. O. *et al.*, editor, *Engineering Societies in the Agents World VII*, Lecture Notes in Artificial Intelligence 4457, pages 245–264, Berlin, Germany, 2007. Springer.
74. T. Miller and P. McBurney. Annotation and matching of first-class agent interaction protocols. In L. P. *et al.*, editor, *Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*, Estoril, Portugal, 2008.
75. C. W. Morris. Foundations of the theory of signs. In O. Neurath, R. Carnap, and C. Morris, editors, *International Encyclopedia of Unified Science*, pages 77–138. Chicago University Press, Chicago, IL, USA, 1938.
76. S. Parsons. Normative argumentation and qualitative probability. In D. M. G. *et al.*, editor, *Qualitative and Quantitative Practical Reasoning*, Lecture Notes in Artificial Intelligence 1244, pages 466–480, Berlin, Germany, 1997. Springer.
77. S. Parsons and E. Sklar. How agents alter their beliefs after an argumentation-based dialogue. In S. P. *et al.*, editor, *Argumentation in Multi-Agent Systems: Second International Workshop (ArgMAS 2005)*, Lecture Notes in Computer Science 4049, pages 297–312. Springer, Berlin, Germany, 2006.

78. L. Perrussel, S. Doutre, J.-M. Thevenin, and P. McBurney. *Argumentation in Multi-Agent Systems*, chapter A persuasion dialog for gaining access to information, pages 63–79. Lecture Notes in Artificial Intelligence 4946. Springer, Berlin, Germany, 2008.
79. M. Pieri. Sui principi che reggiono la geometria di posizione. *Atti della Reale Accademia delle scienze di Torino*, 30:54–108, 1895.
80. H. Prakken. Formal systems for persuasion dialogue. *The Knowledge Engineering Review*, 21(2):163–188, 2006.
81. H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.
82. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and E. Sonenberg. Argumentation-based negotiation. *Knowledge Engineering Review*, 18(4):343–375, 2003.
83. I. Rahwan, E. Sonenberg, N. R. Jennings, and P. McBurney. STRATUM: a methodology for designing automated negotiation strategies. *Applied Artificial Intelligence*, 21(6):489–527, 2007.
84. 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.
85. A. Reinach. Die apriorischen Grundlagen des bürgerlichen Rechtes. *Jahrbuch für Philosophie und phänomenologische Forschung*, 1:685–847, 1913.
86. 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.
87. 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 (AISB-2001)*, York, UK, 2001. AISB.
88. J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, UK, 1969.
89. C. E. Shannon. The mathematical theory of communication. In C. E. Shannon and W. Weaver, editors, *The Mathematical Theory of Communication*, pages 29–125. University of Illinois Press, Chicago, IL, USA, 1963.
90. M. P. Singh. A social semantics for agent communication languages. In F. Dignum and M. Greaves, editors, *Issues in Agent Communication*, Lecture Notes in Artificial Intelligence 1916, pages 31–45. Springer-Verlag: Heidelberg, Germany, 2000.
91. T. C. Tang. Algebraic postulates and a geometric interpretation for the Lewis calculus of strict implication. *Bulletin of the American Mathematical Society*, 44:737–744, 1938.
92. Y. Tang and S. Parsons. Argumentation-based dialogues for deliberation. In F. D. *et al.*, editor, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, pages 552–559, New York City, NY, USA, 2005. ACM Press.
93. A. Tarski. Der Aussagenkalkül und die Topologie. *Fundamenta Mathematicae*, 31:103–134, 1938.
94. A. Tarski. The concept of truth in formalized languages. In *Logic, Semantics, Metamathematics*, pages 152–278. Clarendon Press, Oxford, UK, 1956. (Translated by J. H. Woodger).
95. R. Torretti. *The Philosophy of Physics*. Cambridge University Press, Cambridge, UK, 1999.
96. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany, NY, USA, 1995.
97. M. Wardeh, T. J. M. Bench-Capon, and F. Coenen. Arguments from experience: The PADUA protocol. In P. B. *et al.*, editor, *Computational Models of Argument: Proceedings of COMMA 2008*, pages 405–416, Amsterdam, The Netherlands, 2008. IOS Press.
98. M. J. Wooldridge. Semantic issues in the verification of agent communication languages. *Journal of Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.
99. M. J. Wooldridge. *Introduction to Multiagent Systems*. John Wiley and Sons, New York, NY, USA, 2002.
100. F. Zeuthen. *Problems of Monopoly and Economic Warfare*. Routledge and Sons, London, UK, 1930.