

# When is it okay to lie? A simple model of contradiction in agent-based dialogues

Elizabeth Sklar<sup>1</sup>, Simon Parsons<sup>2</sup>, Mathew Davies<sup>1</sup>

<sup>1</sup> Dept of Computer Science, Columbia University,  
1214 Amsterdam Avenue, Mailcode 1140  
New York, NY 10027 USA  
sklar@cs.columbia.edu

<sup>2</sup> Department of Computer and Information Science, Brooklyn College,  
City University of New York, 2900 Bedford Avenue, Brooklyn,  
New York, NY 11210, USA  
parsons@sci.brooklyn.cuny.edu

**Abstract.** When is it okay to lie? And what constitutes a lie, anyway? This paper examines the notion of lying in agent-based systems, focusing on dialogues and situations where it is acceptable for agents to utter locutions that contradict their beliefs. We examine situations in human and animal behavior where lying — acting or making statements that contradict one’s set of beliefs — is considered to be socially acceptable or even necessary for survival.

## 1 Introduction

When is it okay to lie? And what constitutes a lie, anyway? This paper examines the notion of lying in agent-based systems, focusing on dialogues and situations where it is acceptable for agents to utter locutions that contradict their beliefs. We examine situations in human and animal behavior where lying — acting or making statements that contradict one’s set of beliefs — is considered to be socially acceptable or even necessary for survival.

Consider the following examples:

- a teacher presents a contradictory example to her students in order to motivate them
- a parent uses “reverse psychology” to convince his child to finish eating her vegetables
- an opossum pretends that it is dead so that a predator will not attack it
- a wife tells a “white lie” in order to hide from her husband her plans for giving him a surprise birthday party
- a buyer in an art auction hides his “private value” so that he can make bids that are lower than he would truthfully be willing to pay
- a chameleon changes its color as a camouflage mechanism

In each of these situations, one actor is lying but with good reason. Webster defines the verb “lie” as follows: “to make an untrue statement with intent to deceive; to create

a false or misleading impression” [7]. The level or purpose of the deception is what makes these types of untruths socially acceptable.

We have been examining the use of *dialogues* as interaction mechanisms in agent-based systems. In earlier work of Parsons and colleagues [9, 10], the semantics of the dialogue framework restrict an agent from uttering locutions that contradict its belief set. The reason for this restriction is as follows. Following Singh [11], we wish to provide agents using our dialogue framework with a form of *social semantics* in which other agents can contest any assertion, and refuse to accept it until it has been proved truthful to their satisfaction. The simplest way to achieve this is to restrict agents to only assert things that are, as far as they know, true. In this paper we develop the notion that, as illustrated above, there exist socially acceptable situations in which it may be necessary for agents to contradict their own beliefs in a dialogue. As we shall see, doing this while maintaining the social semantics is considerably trickier than when agents have to tell the truth.

We begin by reviewing previous work on dialogues, highlighting terminology and describing the theoretical framework in which we are working. Next, we present a structure for expanding this dialogue framework in order to be able to model contradiction. Then we outline some examples of how we might apply this contradictory behavior to two of the domains we are actively modelling: classroom management in an education setting and negotiation in a car market.

## 2 Background

A *dialogue game* is structured in terms of *moves* made by two players. An influential model devised by Walton and Krabbe [15] defines six basic types of argumentation that can be combined to create complex dialogues:

- *Information-Seeking Dialogues* —where one participant seeks the answer to some question(s) from another participant, who is believed by the first to know the answer(s);
- *Inquiry Dialogues* —where the participants collaborate to answer some question or questions whose answers are not known to any one participant;
- *Persuasion Dialogues* —where one party seeks to persuade another party to adopt a belief or point-of-view he or she does not currently hold;
- *Negotiation Dialogues* —where the participants bargain over the division of some scarce resource in a way acceptable to all, with each individual party aiming to maximize his or her share;
- *Deliberation Dialogues* —where participants collaborate to decide what course of action to take in some situation. Participants share a responsibility to decide the course of action, and either share a common set of intentions or a willingness to discuss rationally whether they have shared intentions;
- *Eristic Dialogues* —where participants quarrel verbally as a substitute for physical fighting, with each aiming to win the exchange.

Walton and Krabbe do not claim that these are the only possible kinds of dialogue, and indeed others have introduced additional types. Girle [3] discusses a *command dialogue*

in which one agent tells another what to do. McBurney [4] presents *chance discovery dialogue* where two agents arrive at an idea that neither one had prior to the exchange; instead, the idea arises from or is realized by the agents' discussion. For example, the chance discovery would be acknowledged by a phrase such as "Oh, I never thought of that!"<sup>1</sup> Gabbay and Woods [2] have even analysed *non-cooperation* dialogues in which the participants, who may be hostile to one another, do not share the goal of necessarily completing the dialogue.

Within these types of dialogues, in particular *information seeking*, *inquiry* and *negotiation*, Parsons and colleagues have defined six locutions [9, 10, 13]:

- *assert*( $p$ ) —this locution is used in any dialogue where the agent making the assertion has knowledge of the proposal  $p$  from its belief set and wants the other agent to accept it.
- *accept*( $p$ ) —this locution is uttered in response to an assertion and indicates that the agent making proposition  $p$  is deciding to agree with the assertion.
- *question*( $p$ ) —an agent that does not know whether  $p$  is true or not uses *question* to request this information from another agent.
- *challenge*( $p$ ) —this is when an agent is unsure of proposition  $p$  and so questions the agent who uttered it – it is a way of forcing the utterer to reveal their arguments in support of the proposition. An agent has to respond to this by stating its reasons for having asserted  $p$ .
- *quiz*( $p$ ) —this type of locution belongs to a class of *education dialogues* between a tutor and a learner; a tutor asks a question ( $p$ ) of the learner, but the tutor already knows the answer to the question and is interested in determining whether or not the learner knows the answer.
- *answer*( $p$ ) —this locution also belongs to *education dialogues* and is used by a learner in response to a quiz.

Associated with each of these locutions is a set of rules or *axiomatic semantics* [14] which describe the pre-conditions under which an agent may utter a locution and the post-conditions or changes in the agent's belief state that occur as a result of the utterance.

We follow the notational conventions developed previously (see [10], or [8] in this volume) and highlight the elements pertinent to the work discussed herein:

- $\Sigma_i$  represents the *knowledge base*, or beliefs, of each agent  $i$ . If the dialogue takes place between two agents  $M$  (me) and  $U$  (you), then their corresponding knowledge bases are referred to as  $\Sigma_M$  and  $\Sigma_U$ , respectively. This term loosely refers to all the beliefs of the agent.
- An argument  $(S, p)$  is a pair, where  $p$  is a conclusive proposition and  $S$  is its support.  $p$  is a logical consequence of  $S$ , and  $S$  is a minimal subset of  $\Sigma_i$  from which it can be inferred.
- $\mathcal{A}(\Sigma)$  is the set of all arguments that can be made from  $\Sigma$ .
- $\underline{\mathcal{S}}(\Sigma)$  is the set of all *acceptable* arguments in  $\Sigma$  —that is, arguments that an agent has no reason to doubt (i.e., there are either no arguments that *undercut* them, or all the arguments that undercut them are themselves undercut).

---

<sup>1</sup> This example, of course, is only if the exchanges were taking place as natural language.

---

**assert**

LOCUTION:

- $M \rightarrow U : \text{assert}(p)$

PRE-CONDITIONS:

1.  $(S, p) \in \underline{S}(\Sigma_M \cup CS_U)$

POST-CONDITIONS:

1.  $CS_{M,i} = CS_{M,i-1} \cup \{p\}$  (update)
  2.  $CS_{U,i} = CS_{U,i-1}$  (no change)
- 

**Table 1.** Axiomatic semantics for *assert*, uttered by  $M$  as the  $i$ th locution of a dialogue.

- We can partition an agent’s belief set  $\Sigma$  by identifying relevant portions of it. The agent’s *commitment store* ( $CS$ ) refers to statements that have been made in the dialogue and which the agents are prepared to defend. We think of  $\Sigma$  as the agent’s private knowledge base—all of the agent’s beliefs—whereas  $CS$  is the agent’s public knowledge base—all the beliefs that the agent has discussed in public (i.e., with other agents), and hence are known to all other agents.

[10] shows how these simple elements can be used to construct information-seeking, inquiry, and persuasion dialogues.

Table 1 shows the axiomatic semantics associated with the locution **assert**. In order for agent  $M$  to be able to assert a proposal,  $p$ , agent  $M$  has to either:

1. have direct knowledge about that assertion in its set of beliefs;
2. contain an argument that will support the assertion in its set of beliefs; or
3. contain an argument that will support the assertion either in its set of beliefs or in the set of utterances made by the other agent(s) involved in the dialogue.

We summarize these three conditions as

$$(S, p) \in \underline{S}(\Sigma_M \cup CS_U)$$

meaning that  $M$  can assert a proposition if there is an argument to support it in its belief set or in the commitment store of agent  $U$ , the other agent engaged in the dialogue.

There is an additional precondition, which refines the three conditions given above, and is not stated in Table 1 since it varies depending on the agent’s *attitude* [10]. The idea of attitude captures the fact that different agents may be more or less strict about the things it asserts. In particular in [10], an agent may adopt one of three *assertion* attitudes. If agent  $M$  is engaged in a dialogue with agent  $U$ , then:

- if  $M$  is *confident*, then it can assert any proposition  $p$  for which  $(S, p) \in \mathcal{A}(\Sigma_M, CS_U)$

- if  $M$  is *careful*, then it can assert any proposition  $p$  for which there is an argument  $(S, p)$  and no stronger argument  $(S, \neg p)$  exists in  $\mathcal{A}(\Sigma_M, CS_U)$
- if  $M$  is *thoughtful*, then it can assert any proposition  $p$  for which there is an *acceptable* argument  $(S, p) \in \mathcal{A}(\Sigma_M, CS_U)$

These constraints were designed under the assumption that uttering a false proposal, or at least one that cannot be backed up in some way—is considered to be socially unacceptable and ruled out by the social semantics. However, as indicated by the examples in the opening paragraphs of this paper, there are nontrivial circumstances in which an agent may need to utter locutions which contradict its beliefs. How, then, can we allow our agents to lie when they need to without sacrificing the social semantics? The next section explains how this may be done.

### 3 Contradiction in dialogues

While the dialogue game can serve as the mechanism for a wide variety of interactions among agents, the axiomatic semantics of **assert** within a dialogue do not permit an agent  $M$  to make an assertion that contradicts its own beliefs, which we define as a *lie*, or a false proposal. To assert a truth  $p$ , an agent  $M$  must have an *acceptable* argument for  $(S, p)$ . Given the semantics of argumentation, as described in [10] for example, this implies two things<sup>2</sup>. The first is that  $M$  has no argument  $(S', \neg p)$  that is as strong as the argument for  $p$ . The second is that there is no  $r \in S$ , such that  $M$  believes  $\neg r$  more strongly than  $r$ .

More formally, we mean that an assertion  $q$  ( $\neg p$  above) is a *direct lie* if  $M$  knows of a stronger argument supporting  $\neg q$  ( $p$ ) in  $\underline{\mathcal{S}}(\Sigma_M \cup CS_U)$ . A direct lie, then, is the assertion of a fact that is believed to be false. This is a violation of the first condition on assertion. We can also distinguish an *indirect lie*, where  $M$  asserts some  $q$  for which it has an argument  $(S'', q)$  even though there is some  $r \in S''$  which  $M$  believes less strongly than  $\neg r$ .  $M$  is therefore asserting something that it does not believe to be supported by what it believes to be true. This is a violation of the second condition of assertion. (A particular assertion can be both a direct and an indirect lie, as when  $\neg q \in \Sigma_M$  but  $M$  asserts  $q$  anyway.) For the remainder of this paper we will only consider direct lies, but a similar analysis can be carried out for indirect lies (which will require a direct lie if the indirect lie is challenged).

So then, how can  $M$  assert a direct lie  $q$ , since by definition the agent can find no argument in  $\Sigma_M$  or  $CS_U$  supporting  $q$  that wins out over the counterargument? Our solution to this problem is to construct a set of false beliefs, which we call  $J$ , that an agent can use as the logical basis for *justification* of  $q$ , when  $\neg q$  is supported by  $\Sigma_M$  and/or  $CS_U$ . Using the same conventions for notation, we define  $J_M$  informally as the set of all beliefs  $t_i$ , such that  $M$  asserts  $q$  where:

- $(S', \neg q) \in \underline{\mathcal{S}}(\Sigma_M \cup CS_U)$ , and

<sup>2</sup> The precise formal distinction is a little more subtle than this, but without introducing the dialogue system in its full detail—which we do not have room to do here—we have to skate over this subtlety. Suffice it to say that it makes no difference to the validity of the argument we are making here.

- $(S'', q) \in \underline{S}(\Sigma_M \cup CS_U \cup (\cup_i \{t_i\}))$ .

In other words  $J_M$  is exactly that set of propositions necessary to justify the lies that  $M$  has told. Note that this includes the case in which some  $t_i = q$ , that is  $M$  doesn't try to construct a reason why  $q$  is the case, but just claims it is true—a barefaced lie. The agent  $M$  does not *believe* the proposals in  $J_M$ , but in effect holds them for use in passing off the lie, as if they were genuine beliefs in  $\Sigma_M$ .

This, then, provides a way of maintaining the social semantics. If and when another agents questions the lie,  $M$  can respond with the argument that draws on  $J_M$ . This is not guaranteed to be convincing. Depending on how obvious the falsehoods are, the other agent may be able to spot them easily. However, if  $M$  chooses its justifications well, then it may be able to remain undetected. This is, of course, exactly the way that lying works in human society. A lie remains undetected so long as the party that is being lied to has no way to uncover the falsehood on which the lie is based.

Allowing contradiction thus requires a modification of the original semantics of **assert** given in Table 1. In order for an agent  $M$  to utter  $q$  as a lie, two pre-conditions must hold:

- support for an acceptable argument for  $q$  exists in its justification set  $J$ , taken together with the set of utterances made by the other agent(s) involved in the dialogue (i.e.,  $(S, q) \in \underline{S}(\Sigma_M \cup CS_U \cup J_M)$ ), and
- support for an acceptable argument for  $\neg q$  exists either in its set of beliefs or in the set of utterances made by the other agent(s) involved in the dialogue (i.e.,  $(S, \neg p) \in \underline{S}(\Sigma_M \cup CS_U)$ ).

The first condition states that  $q$  *can* be asserted as a contradiction, and the second condition states that  $q$  cannot be asserted as a truth. Taken together, these conditions imply that  $(S, q) \notin \underline{S}(\Sigma_M \cup CS_U)$ ; in other words, the existence of  $J$  as a non-empty set is instrumental to the assertion of  $q$ .

Now that we have defined a way of justifying a lie within our dialogue framework, we need a way of being able to express that lie. We note that we cannot simply create a new locution **lie(p)** because, by definition of our dialogue framework, the type of locution being uttered is actually *included in the utterance*. So for an agent to say **lie(p)**, it would be revealing the fact that it is lying.

In order to get around this, we introduce the notion of a *contradictory* attitude in which the pre-conditions of an assertion are modified in order to allow an agent to utter a proposition that opposes its belief set. A contradictory attitude may also be, at the same time, confident, careful, or thoughtful, as defined earlier (but considering the set of possible arguments whose support includes  $J$ .) The full axiomatic semantics of contradictory assertion are contained in Table 2. Again the additional “attitude” condition applies.

## 4 Carrying off a lie

To knowingly assert even a single falsehood may entail some difficulty for an agent, at least if the agent intends that the lie remain undiscovered. First, if the lie is challenged,

---

*contradictory assert*

LOCUTION:

- $M \rightarrow U : \text{assert}(p)$

PRE-CONDITIONS:

1.  $(S, \neg p) \in \underline{S}(\Sigma_M \cup CS_U)$  AND  
 $(S, p) \in \underline{S}(\Sigma_M \cup CS_U \cup J_M)$

POST-CONDITIONS:

1.  $CS_{M,i} = CS_{M,i-1} \cup p$  (update)
  2.  $CS_{U,i} = CS_{U,i-1}$  (no change)
- 

**Table 2.** Axiomatic semantics for *assert*, contradictory, uttered by  $M$  as the  $i$ th locution of a dialogue.

the agent may have to assert other contradictions (possibly members of  $J$ ), which may in turn require commitment to even more false assertions, resulting in a potential cascade of false commitments —with no guarantee that the original lie  $q$  will even be accepted. Second, even if the lie is accepted, with or without challenge, it may turn out to contradict other (possibly true) proposals the agent may wish to assert in the future. Third, agent  $M$  may wish to maintain consistency with regard to lies uttered in dialogues with particular agents, but may not want to carry the lies into dialogues with *all* agents. In any of these cases, uttering lies is problematic, because each lie potentially impacts the present and future consistency of the agent’s commitment store.

As a method for addressing these issues, we put forth the notion that each lie,  $q$ , has a lifetime. Figuratively speaking, the lie is born when it is first uttered; and the lie dies when the agent who uttered the lie retracts it. We can think of this as adding and subtracting elements from  $J$ . When all lies have been retracted,  $J$  is the empty set. As soon as an agent utters a single lie, it is inserted into  $J$ .

Further, we introduce the notion of “personalized”  $J$ s, whereby agent  $M$  maintains a separate set  $J_{M,j}$  which contains all the lies that  $M$  told to agent  $j$  (that have not been retracted).  $J_{M,j}$  could be thought of as a partition of  $J_M$ . There could be multiple partitions within  $J_M$ . For example, suppose that  $M$  engages in separate dialogues with agents  $U$ ,  $V$  and  $W$ . We will assume here that each dialogue is private, i.e.,  $V$  and  $W$  do not “hear” what is said between  $M$  and  $U$ , and so forth. In talking to  $U$ , any true statement  $p$  that is uttered by  $M$  is just part of  $\Sigma_M$  or  $CS_U$ ; just as in talking to  $V$ , any true statement  $r$  that is uttered by  $M$  is part of  $\Sigma_M$  or  $CS_V$ . But if  $M$  tells a lie,  $q$ , in a dialogue with  $U$ , and then tells another lie,  $s$ , in a dialogue with  $V$ , it is important that  $M$  not assume knowledge of  $s$  when talking to  $U$  nor of  $r$  when talking to  $V$ . The crucial aspect is that  $M$  remember which lies it told to which agents; so in our example,  $q \in J_{M,U}$  and  $s \in J_{M,V}$ , and  $q \notin J_{M,V}$  and  $s \notin J_{M,U}$ . It is also possible that  $M$  wants

to maintain a lie amongst all agents it interacts with, in which case that lie would be a member of each  $J_{M,*}$ . This latter case would circumvent the problem that  $M$  has told  $q$  to  $U$  and  $\neg q$  to  $V$ , but because  $U$  and  $V$  are in contact, these two agents discover the contradiction.

## 5 Why lie?

We have not yet addressed the most important question concerning contradictory dialogues: why do agents lie? As described in the foregoing analysis, the task of agent  $M$  who has asserted proposal  $p$  (that it believes to be true) is to find an acceptable argument  $(S, p) \in \underline{S}(\Sigma_M \cup CS_U)$ . The complexity of this task is several steps above the complexity of checking the consistency of  $p$  in  $\Sigma_M \cup CS_U$  in the hierarchy of computational complexity [10]. It is evident, then, that the task of constructing a set  $J$  of justifications to support the lie  $q$  is not any more difficult, generally speaking, than finding an acceptable argument for some  $p$  that is not a lie.

Suppose agent  $M$  is engaged in a dialogue with a particular goal in mind. With experience,  $M$  will be able to judge the relative merits and difficulties entailed in employing truthful assertions in arguing towards the goal, as compared to employing contradictory ones. Let  $p$  and  $q$  be contradictory and non-contradictory proposals, respectively, each of which, if uttered, could move  $M$  closer to its goal. Since  $M$  can at best only estimate the difficulty of justifying either proposal, inventing an acceptable argument for the lie  $q$  may indeed be considerably easier than finding an acceptable argument for  $p$ , at least in the short term. As previous discussion suggests, the goal may ultimately be defeated if  $q$  results in an unforeseen inconsistency that contradicts some other necessary proposal, as discussed above, or if  $q$  is exposed as a lie.

In natural environments, agents may lie or deceive one another without regard to complexity. However, the goal behind such contradictory behavior need not be socially unacceptable, as indicated in the opening of the paper. For instance, a teacher may assert a contradiction in an education dialogue with students (a form of dialogue we have begun formalizing in [13]), either playing the devil's advocate or to present a counter-example or to provoke the students to challenge the teacher and in so doing explore a set of arguments around some topic.

As another example, both humans and animals are known to exhibit feigned behaviors (such as aggression or flight) both in play, and when learning the purpose and meaning of such behaviors through imitation. [1] has suggested that artificial agents may only develop intelligence recognizable to people through human-like social conditioning, which may require assertions of contradiction either in dialogue or in behavior. Such contradictions may be not just socially condoned, but actually constitute a part of the social and economic fabric of a society. In the next section, we explore several well-known examples of contradictory locution and behavior in human and animal societies that help illuminate the role and necessity of contradiction in complex societies.

## 6 Application domains

Two application areas in which we are actively working show more concretely why we believe that it is important to be able to lie. The first is drawn from our work on simulating aspects of the education system [12], while the second comes from an ecommerce application.

### 6.1 SimEd

The SimEd project is constructing simulations of a number of aspects of the educational system in the US [12]. One of these simulations is of interactions at the classroom level—we are building models that simulate the effects on learning outcomes of different teaching strategies. As a result we are interested in student-teacher dialogues, and recently proposed a formal model of such dialogues [13] which focuses on kinds of dialogue that are common in the classroom but which have not been studied formally before now. (These dialogues do not include contradiction.)

Now, while teachers usually tell the truth to their students, there are occasions upon which lying may be an appropriate action. For example, one way to encourage a child to think through a problem is to present them with a problem and a false solution, and, when they object to the solution, asking them to justify their reaction. The reason for doing this, of course, is to get the student to explain the route to obtaining the correct answer.

Such an interaction is precisely what our framework is capable of providing. The teacher **asserts** the wrong answer, the student then **asserts** the contrary, the teacher **challenges** the student's assertion, and the student has to provide their reasoning. When this is complete there is an explicit ("so what I said to begin with was wrong") or implicit ("yes, you're right") retraction of the initial lie.

### 6.2 Car Market

Our second example comes from [6]. Consider a dialogue about the purchase of a car between the agent for a buyer and a sales agent. This may involve a combination of a number of the kinds of dialogue identified by Walton and Krabbe [15] (combined, for example, as discussed in [5]).

The dialogue might open with an information seeking dialogue in which the sales agent attempts to find out how much money the buyer is prepared to spend, and what features the buyer is looking for. It might then pass into an inquiry stage, during which the two agents attempt to identify the best car, then a negotiation to settle the price, and this latter may include some persuasion on the part of the sales agent in order to get the buyer to agree.

There are several points here where the buyer might find it advantageous to lie. It might be beneficial for the buyer to misrepresent the price that she is prepared to pay, mentioning a smaller amount than is really the case (to avoid the inflation of prices, for example, and also to rule out any attempt by the sales agent to present unsuitably expensive vehicles). It might also be beneficial to lie about the features sought—covering up a weakness for small red sporty cars for example—if these might be exploited to the

agent's disutility, or to be able to pretend that a figure mentioned during the negotiation is so high that negotiations should be broken off then and there (in the hope of gaining a concession).

Again, these forms of lying are exactly those provided for in our model.

## 7 Summary

This paper has presented a formal model of lying in agent-based systems. Arguing that lying can be a useful, and under certain circumstances a desirable, feature of agent-based systems, we have adapted a dialogue framework from our previous work to allow the assertion of untruths. We have presented an axiomatic semantics for the new part of this framework, and have discussed some of the consequences of the modification.

Our work on this topic is ongoing, and there are many areas that we need to explore in order to have a comprehensive treatment of lying. What we have provided here is the start of a semantics for lying in the context of argumentation. That formalisation needs to be completed. However, the formal semantics alone is not enough. We also need to develop our understanding of the *pragmatics* of lying as well. When is it acceptable to lie? When is it better (and in what sense) to lie than to tell the truth? If we are going to lie, what basis shall we use for our lies? These and other questions need to be answered. In addition, we are also looking to implement the dialogue framework to allow us to experimentally evaluate the utility of allowing agents to lie.

## 8 Acknowledgements

This work was made possible by funding from NSF #REC-02-19347 and NSF #IIS 0329037. We are grateful to the reviewers for their helpful comments.

## References

1. K. Dautenhahn. Getting to know each other – artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems*, 16:333–356, 1995.
2. D. M. Gabbay and J. Woods. Non-cooperation in dialogue logic. *Synthese*, 127(12):161–186, 2001.
3. R. Girle. Commands in Dialogue Logic. In D. M. Gabbay and H. J. Ohlbach, editors, *Practical Reasoning: Proceedings of the First International Conference on Formal and Applied Practical Reasoning (FAPR 1996), Bonn, Germany*, Lecture Notes in Artificial Intelligence 1085, pages 246–260, Berlin, Germany, 1996. Springer.
4. P. McBurney and S. Parsons. Chance discovery using dialectical argumentation. In *Proceedings of the Workshop on Chance Discovery, Fifteenth Annual Conference of the Japanese Society for Artificial Intelligence*, Matsue, Japan, 2001.
5. P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *J. Logic, Language, and Information*, 11(3):315–334, 2002.
6. 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*, 7(3):235–273, 2003.

7. <http://www.m-w.com>.
8. S. Parsons, P. McBurney, and M. Wooldridge. Some preliminary steps towards a meta-theory for formal inter-agent dialogues. In I. Rahwan, editor, *Proceedings of the First International Workshop on Argumentation in Multi-Agent Systems*. (this volume), 2004.
9. S. Parsons, M. Wooldridge, and L. Amgoud. On the outcomes of formal inter-agent dialogues. In *2nd International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2003.
10. S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.
11. M. P. Singh. Agent communication languages: Rethinking the principles. In *IEEE Computer 31*, pages 40–47, 1998.
12. E. Sklar, M. Davies, and M. Co. Simed: Simulating education as a multi agent system. In C. Sierra and E. Sonenberg, editors, *Proceedings of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems*. IEEE Press, 2004.
13. E. Sklar and S. Parsons. Towards the application of argumentation-based dialogues for education. In C. Sierra and E. Sonenberg, editors, *Proceedings of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems*. IEEE Press, 2004.
14. R. D. Tennent. *Semantics of Programming Languages*. International Series in Computer Science. Prentice Hall, Hemel Hempstead, UK, 1991.
15. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, USA, 1995.