Talking about doing

Peter McBurney
Department of Informatics
King's College London
Strand, London WC2R 2LS, UK
peter.mcburney@kcl.ac.uk

Simon Parsons
Department of Computer & Information Science
Brooklyn College, City University of New York
Brooklyn NY 11210 USA
parsons@sci.brooklyn.cuny.edu

June 16, 2013

Abstract

Utterances regarding actions have several aspects that make them different to utterances regarding beliefs or other mental states. First, they often require uptake — agreement from another party — for a commitment to be created on the part of the intended executor of the action. Second, the power to revoke this commitment-to-action may lie with a different agent than the one making the initial utterance or the one (possibly the same) executing the action. In this paper we explore these different possible combinations of executor, uptaker and revoker for utterances over action, and present a game-theoretic semantics for such utterances.

1 Introduction

Philosophers have studied propositions since Aristotle. They have devoted far less attention to other types of utterance. Speech act theory, due primarily to John Austin [Austin, 1962] and John Searle [Searle, 1969], has remedied that bias in modern times, but there were at least two earlier precursors to this work. The Scottish philosopher, Thomas Reid (1710–1796) [Schuhmann and Smith, 1990], and the German philosopher, Adolf Reinach (1883–1917) [Reinach, 1913], both looked at the nature of promises and similar utterances. Each observed that most utterances about action require an audience: one cannot issue a command to oneself or meaningfully make a request to oneself or pray to oneself. To that extent, such utterances about action are *social acts*. Moreover, these utterances may only produce binding commitments to undertake a particular action if someone *accepts* the utterance: a promise by

me to you to do some action only becomes a commitment on me to do that action when you accept the promise; likewise, a request by me to you to do some action only becomes a commitment on you to execute the action if you accept the request. Reinach, whose explorations were motivated by an attempt to understand the law of contractual obligations, called this acceptance event *uptake* [Reinach, 1913].

Uptake is not a property of utterances about truth or knowledge, or utterances about the speaker's desires and preferences. In this respect, utterances about action are fundamentally different to those about beliefs. These two types of utterances differ in another important respect: for utterances over knowledge or beliefs, the original speaker may normally retract his or her prior statement of knowledge or belief and issue a new statement, asserting new knowledge or belief. Such retraction may revoke a commitment to defend a prior statement asserting knowledge or belief. In the case of utterances over actions, however, this power to revoke or retract the utterance or to annul any resulting commitment may not lie with the original speaker, but may be vested — through cultural tradition or linguistic practice — in someone else. For instance, it is not normal (at least in English-speaking cultures) for the maker of a promise that has already been accepted by the receiver of that promise to be able to un-make it: Only the receiver of the promise may — normally — annul it[1].

In previous work [McBurney and Parsons, 2007], we identified revocation as an aspect of utterances over action which had not been given any attention in the agent communications literature. In our earlier paper, we also proposed a syntax, where a speaker may indicate who is intended to execute the action, and who has power to revoke the action, and gave these utterances a denotational semantics in terms of manipulations of the contents of specified tuple spaces.

An example may make this clear. Suppose two agents, named Alice (A) and Bob (B), are engaged in conversation about the performances of some actions. Suppose that at different times in the interaction, A makes the following utterances to B:

Sentence 1: *I command you to wash the car.*

Sentence 2: *I promise you to wash the car.*

The syntactical form of these two statements is identical, despite the fact that almost everything else is different: the identity of the individual that the speaker intended to do the action (Bob for sentence 1, Alice for sentence 2); the illocutionary force of the two utterances (many people would view commands creating stronger commitments than promises, for example); the nature of the commitment created in each case; the identity of the agent who would be obligated under any commitment to execute the action (Bob and Alice, respectively); the identity of the agent to whom the commitment would be due (Alice and Bob, respectively); and if a commitment was created, the identity of the agent with power to revoke it or accept that it had been fulfilled (Alice and Bob, respectively).

[1] In contemporary society there is nothing to prevent the maker of the promise breaking it, but it is clear that a social norm has been violated. It was not so long ago, however, that if a man refused to marry a woman that he had agreed to marry, he could be sued for "breach of promise", and a woman who changed her mind about an offer of marriage that she had accepted was in danger of damaging her reputation, as discussed in Trollope [1865]

In this paper, we present in Section 2 a formal syntax for utterances over action which allows for different allocation of the responsibilities for doing, for uptaking and for revoking action commitments. We then provide a game theoretic semantics for such utterances, in terms of formal dialog games involving subsets of the participants; this is presented in Section 3. Section 4 discusses related work, and Section 5 concludes the paper with a brief discussion of potential future work.

2 Utterances over actions: Syntax

We begin by considering a dialog between two agents, identified as A and B respectively, who consider a single, atomic action, labeled α, to be undertaken at some time in the future. We assume that agent A makes an utterance about the potential action to agent B. The action α may be undertaken by agent A, or by agent B, or by a third agent (possibly not present in the interaction), agent C. It may help to think of the agents with first names: Alice (A) is the speaker of the utterance, Bob (B) is the hearer of the utterance, and Carol (C) is a possibly-non-present third-party. We ignore situations where the nature of the action needs two or more agents to execute it successfully.[2] Uptake (acceptance) of the action and thereby creating a commitment on the agent doing the action could be a power given to any one of the three agents. Likewise, once a commitment to execute α has been created, any one of the three agents could have the power to revoke or annul the commitment. Again ignoring cases where uptake or revocation requires joint or several participants, we therefore have $3^3 = 27$ possible cases, as shown in Table 1.

Some of the cases in Table 1 correspond to the everyday usage of certain locutions in English; for instance, case 5 where the speaker, agent A intends to execute an action after acceptance by agent B, and with revocation power also vested in B describes the allocation of these roles in normal usage of *promise* locutions. However, most of the 27 cases do not have corresponding commonly-used English locutions. Even those cases which do have corresponding natural language locutions may be open to interpretation, or their meanings may differ in different cultures or language-groups.

An example is given by cases 10 and 12, *commands* from agent A to agent B. Legitimate commands are those where the commander has the legal or moral right to issue commands to the recipient of the command, the commandee. This right may well be contestable in specific cases, so it is reasonable to say that the recipient (say, agent B) may question or contest the command. Even if agent B accepts A's right to issue commands to B, the command may require clarification or explanation before it can be executed, as we have explored in [Atkinson *et al.*, 2008]. Even organizations with strict hierarchies, such as military organizations, where members are empowered to issue commands to specified others, still permit recipients of commands to contest or refuse to obey unlawful commands.[3] Once uptaken, however, a command may only be

[2]Examples of such actions are: carrying a heavy object, operating a two-man saw (common in the lumber industry before the invention of the chain saw), and launching nuclear missiles [Engel *et al.*, 2003].

[3]The US television series, *Generation Kill*, for instance, based on the experiences of a US Marine Corp in the invasion of Iraq in 2003 provides many examples of military commands being contested or questioned

revoked or annulled by the commander, not the commandee.

We now present a simple syntax for these utterances, which is:

$$actionloc(\alpha, S, \mathcal{H}, A, U, R)$$

where $actionloc$ indicates that this is intended to be an utterance over action, α represents the proposed action, S is the identity of the speaker of the utterance, $\mathcal{H}$ is a finite set containing the identifiers of agents intended to be hearers of the utterance, A is the identifer of an agent intended to be the executor of the proposed action, U is the identifer of the agent empowered to accept (or uptake) the proposed action, thereby creating a commitment on agent A to execute the action, and R is the identifer of an agent empowered to revoke or annul the commitment on agent A to execute the action. We define $\mathcal{H}$ to be a set of agents, rather than only a single agent, to allow for multiple intended auditors to an utterance; the elements of $\mathcal{H}$ may or may not include agents A, U or R. The syntax of the utterance could also include terms to represent the intended starting time, completion time, or duration of the action α, as well as any necessary pre-conditions, concurrent conditions, or desired end-states of the action. For simplicity of presentation, we ignore such elements here.

This syntax is given in the two-layer structure now standard in agent communications, as defined, for example, in the Agent Communications Language ACL of FIPA FIPA [2002]. The outer (or wrapper) layer is given by $actionloc$, with the elements $\alpha, S, \mathcal{H}, A, U$ and R all part of the inner layer. Note that in this syntax $actionloc$ is a specific term, and does not stand in for, or represent, some other specific locution or locutions, such as *promise* or *command*. The use of $actionloc$ as a specific term means that we can use this to embody any one of the 27 possible action locutions presented in Table 1, including those for there is no standard natural language term.

3 Game-theoretic semantics

The preceding discussion has shown how aspects of the meaning of a statement about actions may depend on the utterances of others, whether in creating a commitment via uptake, or in annulling a prior commitment via revocation. Insofar as these aspects of meaning do not relate to the truth-status of the utterance, some linguists would regard them as being part of the pragmatics rather than part of the semantics of the utterance [Levinson, 1983]. We can use the social, pragmatic features of utterances over action to define a game-theoretic semantics for these utterances.

Firstly, a few words on the notion of semantics are in order. There are differences in the meaning of the term *semantics* and in the reasons for, and use of, semantics in the different domains of linguistics, formal logic, programming language theory, and agent communications. We have presented these different notions and purposes in greater

in the course of the operation, and urban legend has it that in the POUM militia in the Spanish Civil War:

> You obeyed an order only after it had been explained in detail, you could refuse it if it didn't make sense [Haldeman, 2009, page 221].

detail in [McBurney and Parsons, 2009, Section 3]. Propositions purport to describe some feature of the real-world, and thus may carry truth values. For agent communications involving propositions it is therefore appropriate to view the notion of semantics as something close to standard usage in linguistics: a semantics for well-formed statements about propositions in a dialog is a mapping between the statements and objects in the real world. However, statements over action are usually not intending or purporting to describe some existing reality, and so do not carry truth values. Thus, a different notion of semantics is required for these statements. Following [Johnson *et al.*, 2005], we define the semantics of utterances over action as the external commitments-to-execute-actions referred to by the statements. Valid utterances in properly-undertaken dialogs can be then be viewed as creating, manipulating, assigning, re-assigning and annulling such commitments.

As explained in [McBurney and Parsons, 2009], there are several different ways such semantic mappings could be defined. A game-theoretic semantics is a semantics for statements in some logical language that associates a conceptual game to each well-formed statement of the language. The game is usually imagined to be played by two imaginary players, often named *Protagonist* and *Antagonist*. The statement of the formal language is deemed to be *true* precisely when one of the parties, usually Protagonist, has a strategy to ensure success in the game associated to the statement. By *strategy* is meant a decision-rule telling Protagonist what game moves to play in the game, for each possible prior move of Antagonist, and for each history of moves in the game to that point. Game theoretic semantics were developed by Jaako Hintikka for statements in first-order logic, creating what he called *Independence-Friendly (IF) Logic* [Hintikka and Sandu, 1997]. Game semantics have been articulated for propositional and predicate logics [Lorenzen and Lorenz, 1978], linear logic [Abramsky, 1997], and for probability statements [Dawid and Vovk, 1999], among others. So-called Dutch-book arguments in probability theory (due to Frank Ramsey and Bruno de Finetti) may also be viewed as a game semantics for subjective probability statements, since these arguments involve infinite (and therefore imaginary) gambles between a decision-maker and an imagined book-maker.

Game semantics have also been used to study the properties of formal argumentation systems and dialogue protocols, such as their computational complexity [Dunne and Bench-Capon, 2003], or the extent of truth-convergence under an inquiry dialogue protocol [McBurney and Parsons, 2001], and to identify acceptable sets of arguments in argument frameworks [Cayrol *et al.*, 2003; Jakobovits and Vermeir, 1999]. One could even view the English common-law legal system in game-theoretic terms, as a student of Trevor Bench-Capon, John Henderson, showed in his 2006 PhD thesis [Henderson, 2006]. Conceptual games have also found application in mathematical model theory and in theoretical computer science, e.g., [Hodges, 1985].

We now present the semantics of the statements defined in Section 2 in terms of two multi-agent dialog games, as follows. For each utterance with the syntactic form:

$$actionloc(\alpha, S, \mathcal{H}, A, U, R)$$

we first associate two specific dialog games, as follows:

Uptake Dialog Game, denoted $UD(\alpha, S, U)$, is a dialog game between agents S and

U, played with the rules given in Table 2, and

Revocation Dialog Game, denoted $RD(\alpha, S, U, R)$, is a dialog game between agents S, U and R, played with the rules given in Table 3.

Table 2: **Uptake Dialog Game: Outline Rules**

Uptake Dialog Game: $UD(\alpha, S, U)$
Dialog Preconditions: A prior utterance of $actionloc(\alpha, S, \mathcal{H}, A, U, R)$. *Dialog Participants:* Agents S and U. *Valid locutions:* Question, Justify, Accept, Reject, End-dialog *Combination rules:* Only agent U may utter Accept or Reject. An agent may only utter Justify following a prior utterance of Question by another agent. *Commitment rules:* When and only when U utters Accept, a commitment on agent A to execute action α is created. *Termination rules:* The dialog ends upon utterance by U of Accept or Reject, or the utterance of End-dialog by any agent. The dialog is said to terminate-with-uptake upon utterance by agent U of Accept.

Tables 2 and 3 present outline rules for these two dialogs, respectively. The dialog rules are given here in accordance with the formal specification structure for agent dialog games first articulated in [McBurney and Parsons, 2002]. This structure defines the valid locutions, any rules for utterance and combination of these locutions, any rules for the creation and manipulation of commitments potentially incurred through the utterance of specified locutions, and any rules for termination of the dialogs. For reasons of space, we only present these rules in outline form here. The locutions *Question* and *Justify* have the obvious intended meanings, which are as defined for the Fatio Protocol [McBurney and Parsons, 2005].

We now define the semantics of action statements as follows. The utterance of the statement,

$$actionloc(\alpha, S, \mathcal{H}, A, U, R)$$

creates a commitment on agent A to execute action α precisely when the associated Uptake Dialog Game $UD(\alpha, S, U)$ terminates-with-uptake. This commitment on agent

A to execute action α remains in force unless and until the action α is executed by agent A or the associated Revocation Dialog Game, $RD(\alpha, S, U, R)$, ends in termination-with-revocation.

Table 3: **Revoke Dialog Game: Outline Rules**

Revoke Dialog Game: $RD(\alpha, S, U, R)$
Dialog Preconditions: 1. A prior utterance of $actionloc(\alpha, S, \mathcal{H}, A, U, R)$ 2. A prior instance of the associated Uptake Dialog Game $UD(\alpha, S, U)$ which terminated-with-uptake and 3. The commitment on agent A to execute action α remains in force. *Dialog Participants:* Agents S, U and R. *Valid locutions:* Propose-revoke, Question, Justify, Revoke, End-dialog *Combination rules:* Only R may utter Revoke. An agent may only utter Justify following a prior utterance of Question by another agent. *Commitment rules:* When and only when R utters Revoke, the commitment on agent A to execute action α is voided. *Termination rules:* The dialog ends upon utterance by R of Revoke, or the utterance of End-dialog by any agent. The dialog is said to terminate-with-revocation upon utterance by agent R of Revoke.

The definition here of the semantics of utterances over action in terms of specified agent dialog games means that we have provided explicit procedures for creating and revoking commitments to act, shown by the outline rules of Tables 2 and 3. These procedures are defined in terms only of publicly-observable behaviours, not in terms of any decision or computational processes internal to the agents involved, which means the processes are semantically verifiable [Wooldridge, 2000], and are in accordance with the first of Hitchcock's *Principles of Rational Mutual Enquiry* [Hitchcock, 1991].[4] The explicit definition of the processes also aids the software engineering of these utterances into agent communications protocols, as shown, for example, [Doutre *et al.*, 2005; Moschoyiannis *et al.*, 2009].

[4]Namely: **Externalization:** The rules should be formulated in terms of verifiable linguistic behavior.

4 Related Work

To the best of our knowledge, there is no prior work directly related to the work of this paper. As mentioned, some of our ideas are inspired by the philosophy of language of Thomas Reid [Schuhmann and Smith, 1990] and the philosophy of legal language of Adolf Reinach [Reinach, 1913], although neither considers retraction or revocation. Elsewhere in speech act theory, Jürgen Habermas explored the different nature of challenges or rebuttals required by different types of utterances, including statements about actions, although he seems not to have considered either uptake or revocation [Habermas, 1984]. Much recent work in computational argumentation has explored argumentation schemes and their associated critical questions (CQs), often proposing formal dialog games to allow participants to argue over the default conclusion of the scheme using the critical questions, for example, [Atkinson *et al.*, 2005; Medellin-Gasque *et al.*, 2013]. Such approaches may also be viewed as providing a game-theoretic semantics for the scheme and CQs, with the default conclusion of the scheme being adopted precisely when the player designated as Protagonist wins the associated dialog game.

Our approach differs from other work in agent communications on commitments. The social semantics of Singh and Colombetti and their respective colleagues [Colombetti and Verdicchio, 2002; Singh, 1999] treats utterances in agent dialogs as devices for manipulating the social relationships between the speakers. Our work, focused only statements about actions, is at a lower level of abstraction than social semantics. We assume that a dialog commences with two or more participants joining together with the shared intention of deciding what action or actions to take in some circumstance. There may already be prior social relationships between the participants, which could thereby allow, for example, commands to be uttered legally by one agent to another. However, once a dialog about action commences, we desire to understand how commitments to execute actions (or not to) are created and manipulated by the participants in the dialog. Our focus is therefore on the short-term effects of utterances inside a dialog on commitments, not their longer-term effects on the social relationships between the participants.

One could ask why the semantic differences of speech acts identified in Section 1 could not be captured by the notion of agent roles, as in a framework such as that of [Wooldridge *et al.*, 2000]. The reason is that the role of uptaker or revoker of an utterance is not usually fixed throughout an interaction; it potentially depends on: the type of the utterance (promise, command, etc); the identities of the agent making the utterance, the agent receiving it, and the agent tasked with the action; and possibly also on the history of the dialog to that point. All of these may change through the course of a dialog, particularly if there are embedded dialogs or other complex combinations of dialogs, e.g., [McBurney and Parsons, 2002; Reed, 1998]; thus, agent roles will usually be too rigid a framework for tracking these abilities to uptake or revoke utterances.

5 Conclusions

In this paper we have presented a formal syntax for utterances over action which allows for different allocations of the responsibilities for doing, for accepting (uptaking), and for revoking action commitments. We then provided a game-theoretic semantics for these utterances, in terms of formal dialog games each involving subsets of the participants specified in the original utterance. The benefits of this approach are several. Firstly, the proposed syntax makes explicit exactly who has responsibility for doing, for accepting, and for revoking an action commitment. Secondly, the syntax allows for every possible allocation of these responsibilities across three different agents, i.e, for all 27 possible combinations. The syntax therefore generalizes from those special cases of utterances over action that human language users have distinguished over the years, namely, promises, commands, requests, prayers, etc. Thirdly, the game-theoretic semantics we have defined gives explicit procedures for agents to accept (or not) a proposed action intention (and thus a procedure for creating a commitment on the part of the intended executor of the action to undertake the action), and for agents to revoke or annul an existing action commitment. These explicit procedures are defined in terms of formal agent dialog games which may readily be implemented in agent communication systems.

In this work we have thus far only considered atomic actions, each undertaken by a single actor. Two avenues of potential future work therefore are to consider utterances involving combinations of actions and to consider actions requiring two or more agents to be executed. Recent research by Rolando Medellin Gasque, for example, has explored dialogs over plans involving multiple actions and actors, work which demonstrates the subtlety of the issues involved when more than one action is considered [Medellin-Gasque *et al.*, 2012, 2013]. Conditional actions, those whose execution depend in some way upon the execution of other actions, and the resulting combinations of action commitments, is another area of potential future work. In [McBurney and Parsons, 2002], we showed that the combination of commitments arising from agent communications is also not necessarily straightforward; as with multiple actions and actors, these issues will require careful treatment.

Acknowledgments

This paper extends and generalizes the ideas presented by the authors at the *Symposium on Logic and Games in MultiAgent Systems (LoGaMAS)*, held in Liverpool, UK, on 16–17 December 2002, under the title of *"Towards a game semantics for logics of practical reasoning."* We are grateful for comments received from the audience on that occasion, particularly Marc Pauly.

Some Personal Recollections

PM: I first met Trevor Bench-Capon in November 1999, just before moving to University of Liverpool to complete my PhD. I had previously seen references to his work, but had not read any of it. In February 2002, I was appointed to a faculty position at

Liverpool. A few months later, I was approached by Katie Greenwood (now Atkinson) about doing a PhD. Katie then applied for a Departmental scholarship to do a PhD in the area of agent negotiation under my supervision, and her application was successful. Subsequent to the award of the scholarship, but before she started, Trevor approached me to ask if he could "sit in" on our supervisory meetings. I did not know what that meant, and asked him. He said that he wanted to learn more about the topics of agent interaction and negotiation and thought that witnessing our discussions would help him do this. Still unsure what this involved, I hesitated. He responded by saying that he would just sit quietly in the corner, and not ever utter a word in our meetings. At the time, I did not know Trevor very well, and so I accepted this promise, and allowed him to join us as a silent witness. Fortunately, he did not keep his promise very long. The immensely interesting discussions between the three of us led directly to our joint development of the fruitful model for making and arguing about proposals over action [Atkinson *et al.*, 2005], work which built on prior work of both Trevor and myself, as well as others, on qualitative models of practical reasoning involving consequences.[5]

In October 2003, Julian Padget of the University of Bath gave an invited research seminar to the Department of Computer Science at the University of Liverpool. In his talk, Julian explained that something he was saying would be obvious to a computer scientist. Realizing that perhaps not everyone in the audience was from the Department, he then added, *"But, of course, not everyone here is necessarily a computer scientist."* Trevor brought laughter to the room with his witty reply, *"Yes — I am only contingently a computer scientist."*

SP: The two stories that I tell most often about Trevor are the following. The first took place in a pub (no surprise for anyone who knows Trevor) (or, come to that, anyone who knows me). My recollection is that the pub was The Cambridge, close to the University of Liverpool campus, but I may be wrong. In any case, Trevor was standing near the bar, pint in one hand, lit cigarette in the other (this was probably in 2001, back when he was a heavy smoker). It was shortly before Easter, and, hoping to find something to tease him about, I asked Trevor what he had given up for Lent. His reply was that, as always, he had given up smoking and drinking. My reply, naturally, was to ask how that squared with him both drinking and smoking, to which he responded "The point of Lent is to contrast human frailty with divine perfection. Failing to give something up is the point". To which there was no answer, except to buy him another drink.

The second took place in Washington DC, at a workshop on argumentation that Trevor, Henry Prakken and I organised as part of a AAAI Fall Symposium Series. These meetings always include a plenary session in which each symposium is presented to an audience made up of attendees from all the events. Often the talks are rather dry, and AAAI usually encourages presenters to try to make them a little lively. In the case

[5]For example, [Bench-Capon, 2003; Fox and Parsons, 1998; McBurney *et al.*, 1999]. In 1992, Peter McBurney and David Shuker developed a framework for planning and writing large-scale applications for national telecommunications licences which we called the *IPOC Model*. This framework organized materials around the various processes to be undertaken in implementing and launching a public telecommunications network, with each required process described in terms of its required Inputs, the tasks involved in the Process itself, its Outputs when successfully executed, and the anticipated Consequences of its execution. One could view a telecommunications licence application as just a very large and sophisticated proposal for action.

of our symposium, Guillermo Simari had agreed to give the talk, and just as he was giving a detailed technical description of an argument, Trevor, who was sitting in the audience shouted "That isn't an argument". Into the stunned silence I replied "Yes it is", and we alternated "No it isn't", "Yes, it is"[6] until the audience realised that this was planned and not some spontaneous heckling.

References

Abramsky, S. (1997). 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.

Atkinson, K., Bench-Capon, T., and McBurney, P. (2005). A dialogue-game protocol for multi-agent argument over proposals for action. *Journal of Autonomous Agents and Multi-Agent Systems*, **11**(2), 153–171.

Atkinson, K., Girle, R., McBurney, P., and Parsons, S. (2008). Command dialogues. In I. Rahwan and P. Moraitis, editors, *Argumentation in Multi-Agent Systems. Proceedings of the Fifth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2008)*, Lecture Notes in Artificial Intelligence 5384, pages 93–106. Springer, Berlin, Germany.

Austin, J. L. (1962). *How To Do Things with Words*. Oxford University Press, Oxford, UK. (Originally delivered as the William James Lectures at Harvard University in 1955.).

Bench-Capon, T. J. M. (2003). Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, **13**(3), 429–448.

Cayrol, C., Doutre, S., and Mengin, J. (2003). On decision problems related to the preferred semantics for argumentation frameworks. *Journal of Logic and Computation*, **13**(3), 377–403.

Colombetti, M. and Verdicchio, M. (2002). An analysis of agent speech acts as institutional actions. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), Bologna, Italy*, pages 1157–1164, New York City, NY, USA. ACM Press.

Dawid, A. P. and Vovk, V. G. (1999). Prequential probability: principles and properties. *Bernoulli*, **5**, 125–162.

Doutre, S., McBurney, P., Wooldridge, M., and Barden, W. (2005). Information-seeking agent dialogs with permissions and arguments. Technical Report ULCS-05-010, Department of Computer Science, University of Liverpool, Liverpool, UK.

[6] Part of the script from the Monty Python "Argument Sketch" Python [1972].

Dunne, P. E. and Bench-Capon, T. J. M. (2003). Two party immediate response disputes: Properties and efficiency. *Artificial Intelligence*, **149**(2), 221–250.

Engel, J. A., Slattery, C., Ebeling, M., Pogany, E., and Squitieri, A. R. (2003). *The Missile Plains: Frontline of America's Cold War*. US National Parks Service, Minuteman Missile NHS, SD, USA.

FIPA (2002). Communicative Act Library Specification. Standard SC00037J, Foundation for Intelligent Physical Agents.

Fox, J. and Parsons, S. (1998). Arguing about beliefs and actions. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, volume 1455 of *Lecture Notes in Artificial Intelligence*, pages 266–302. Springer, Berlin, Germany.

Habermas, J. (1984). *The Theory of Communicative Action: Volume 1: Reason and the Rationalization of Society*. Heinemann, London, UK. (Translation by T. McCarthy of: *Theorie des Kommunikativen Handelns, Band I, Handlungsrationalitat und gesellschaftliche Rationalisierung.* Suhrkamp, Frankfurt, Germany. 1981.).

Haldeman, J. (2009). *The Forever War*. Thomas Dunne Books.

Henderson, J. (2006). *A Description of Common Law as a Moving Classification System*. Ph. D., Department of Computer Science, University of Liverpool, Liverpool, UK.

Hintikka, J. and Sandu, G. (1997). Game-theoretical semantics. In J. Benthem and A. Meulen, editors, *Handbook of Logic and Language*, pages 361–410. Elsevier, Amsterdam, The Netherlands.

Hitchcock, D. (1991). Some principles of rational mutual inquiry. In F. Eemeren, R. Grootendorst, J. A. Blair, and C. A. Willard, editors, *Proceedings of the Second International Conference on Argumentation*, pages 236–243, Amsterdam, The Netherlands. SICSAT: International Society for the Study of Argumentation.

Hodges, W. (1985). *Building Models by Games*. Cambridge University Press, Cambridge, UK.

Jakobovits, H. and Vermeir, D. (1999). 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. ACM Press.

Johnson, M. W., McBurney, P., and Parsons, S. (2005). A mathematical model of dialog. *Electronic Notes in Theoretical Computer Science*, **141**(5), 33–48.

Levinson, S. C. (1983). *Pragmatics*. Cambridge Textbooks in Linguistics. Cambridge University Press, Cambridge, UK.

Lorenzen, P. and Lorenz, K. (1978). *Dialogische Logik*. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany.

McBurney, P. and Parsons, S. (2001). Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, **32**(1–4), 125–169.

McBurney, P. and Parsons, S. (2002). Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, **11**(3), 315–334.

McBurney, P. and Parsons, S. (2005). Locutions for argumentation in agent interaction protocols. In R. M. van Eijk, M.-P. Huget, and F. Dignum, editors, *Developments in Agent Communication*, Lecture Notes in Artificial Intelligence 3396, pages 209–225. Springer, Berlin, Germany.

McBurney, P. and Parsons, S. (2007). Retraction and revocation in agent deliberation dialogs. *Argumentation*, **21**(3), 269–289.

McBurney, P. and Parsons, S. (2009). Dialogue games for agent argumentation. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*, chapter 13, pages 261–280. Springer, Berlin, Germany.

McBurney, P., Parsons, S., Fox, J., and Glasspool, D. (1999). Consequence logics for cancer risk counselling systems: a specification framework. Technical report, Department of Electronic Engineering, Queen Mary & Westfield College, University of London, London, U.K.

Medellin-Gasque, R., Atkinson, K., McBurney, P., and Bench-Capon, T. (2012). Arguments over co-operative plans. In S. Modgil, N. Oren, and F. Toni, editors, *Theory and Applications of Formal Argumentation*, Lecture Notes in Artificial Intelligence 7131, pages 50–66, Berlin, Germany. Springer.

Medellin-Gasque, R., Atkinson, K., Bench-Capon, T., and McBurney, P. (2013). Strategies for question selection in argumentation about plans. *Argument and Computation. In press.*

Moschoyiannis, S., Bryant, D., Krause, P., and McBurney, P. (2009). Verifiable protocol design for agent argumentation dialogues. In O. Kaynak and E. Damiani, editors, *Proceeedings of the Third International IEEE Conference on Digital Ecosystems and Technologies (IEEE DEST 2009)*, Istanbul, Turkey. IEEE.

Python, M. (1972). `http://www.youtube.com/watch?v=kQFKtI6gn9Y`.

Reed, C. (1998). 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.

Reinach, A. (1913). Die apriorischen Grundlagen des bürgerlichen Rechtes. *Jahrbuch für Philosophie und phänomenologische Forschung*, **1**, 685–847. English translation by John F. Crosby, published in *Aletheia*, III (1983), pp. 1–141.

Schuhmann, K. and Smith, B. (1990). Elements of speech act theory in the philosophy of Thomas Reid. *History of Philosophy Quarterly*, **7**, 47–66.

Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, UK.

Singh, M. P. (1999). An ontology for commitments in multiagent systems: toward a unification of normative concepts. *Artificial Intelligence and Law*, **7**, 97–113.

Trollope, A. (1865). *Can You Forgive Her?* Chapman & Hall.

Wooldridge, M. J. (2000). Semantic issues in the verification of agent communication languages. *Journal of Autonomous Agents and Multi-Agent Systems*, **3**(1), 9–31.

Wooldridge, M. J., Jennings, N. R., and Kinny, D. (2000). The Gaia methodology for agent-oriented analysis and design. *Journal of Autonomous Agents and Multi-Agent Systems*, **3**(3), 285–312.

Table 1: **Combinations of Actors, Uptakes and Revokers**. A speaks to B. C is not present.

No.	Who does action	Who uptakes	Who revokes	Type of Utterance
1.	A	A	A	Intention to act by A
2.	A	A	B	
3.	A	A	C	
4.	A	B	A	
5.	A	B	B	Promise by A to B
6.	A	B	C	
7.	A	C	A	
8.	A	C	B	
9.	A	C	C	
10.	B	A	A	Command by A to B (valid)
11.	B	A	B	
12.	B	A	C	
13.	B	B	A	Command by A to B (contested)
14.	B	B	B	Request by A to B
15.	B	B	C	
16.	B	C	A	
17.	B	C	B	
18.	B	C	C	
19.	C	A	A	Command by A to C (valid)
20.	C	A	B	
21.	C	A	C	
22.	C	B	A	
23.	C	B	B	
24.	C	B	C	
25.	C	C	A	Request or Prayer by A to C
26.	C	C	B	
27.	C	C	C	Request by A to C