A Dialogue Game to Offer an Agreement to Disagree

Henk-Jan Lebbink[1,2,3], Cilia Witteman[2], and John-Jules Ch. Meyer[1]

[1] Institute of Information and Computing Sciences,
Utrecht University, The Netherlands
henkjan@cs.uu.nl
[2] Social Sciences, Radboud University Nijmegen, The Netherlands
[3] Emotional Brain BV, Almere, The Netherlands

Abstract This paper proposes a dialogue game in which coherent conversational sequences at the speech act level are described of agents that become aware they have an irresolvable disagreement and settle the dispute by agreeing to disagree. A disagreement is irresolvable from an agent's perspective if the agent is aware that both parties have ran out of options to resolve the dispute, and that both parties are aware of this. A dialogue game is formulated in which agents can offer information that may unintentionally result in irreconcilable, mutually inconsistent belief states. Based on the agents' cognitive states, dialogue rules and cognitive rules are defined that allow agents to come to an agreement to disagree. These rules are implemented in the programming language Prolog, resulting in an intuitive design for multi-agent systems.

1 Introduction

If our goal is to understand human conversations, we may need to model carefully their underlying principles, but for generating communication in multi-agent systems, we may be satisfied if we can build computational models that generate efficient and useful conversations. In conversations in general, participating agents have autonomy over their cognitive states, but they may also have desires to change those of others. In trying to do so, these agents may find themselves stuck in impasses over irreconcilable beliefs. This paper addresses the questions how to cope with these impasses and how to devise a computational model to identify irreconcilable beliefs from an agent's local perspective. We will use dialogue games to define reasoning and communication rules to overcome such situations.

Dialogue games have recently received more attention in the field of computer science, and, especially, in the community of multi-agent systems [1, 2]. In multi-agent systems, autonomous software agents communicate and cooperate to reach private and collective goals. We will not address issues related to cooperation, but we will focus primarily on agents engaging in communication. A dialogue typology by Walton and Krabbe [3] identifies four different categories of dialogues by distinguishing the agents' initial situations and goals. The categories are:

R.H. Bordini et al. (Eds.): PROMAS 2004, LNAI 3346, pp. 199–223, 2005.

Persuasion dialogues, in which agents seeks to convince other agents to believe propositions [4, 5, 6]; Negotiation dialogues, in which participants seek to agree on how to divide a resource [7, 8]; Deliberation dialogues, in which participants make plans by discussing which actions to perform in which situations [9] or which beliefs to accept to minimize uncertainty [10] and which result in collective intentions [11] or group-plans and teams [5]; And information seeking dialogues, in which agents seeks to find truth-values of propositions by asking others who may have answers [12, 13, 14]. The current work contributes to the category of persuasion dialogues, and, especially, when persuasion dialogues terminate.

Beun [13] and Lebbink et al. [14] describe communicative acts and communication rules that form dialogue games that agents play to balance their desires and belief states. Such a dialogue game consists of pre-conditions for uttering communicative acts to convey information to other agents, and post-conditions that state the agents' cognitive states after incoming and outgoing information is processed. To describe inconsistent and biased information states, a multi-valued logic [14] is used with which agent can have inconsistent belief states without being forced to perform belief revision. We formulate a semantics for communicative acts to offer information to agree to disagree in the same vein as the FIPA work on agent communication languages [15, 16].

What is lacking in Beun [13] and Lebbink et al. [14] is the possibility for agents to recognize irresolvable disagreements and, based on this recognition, to utter an agreement to disagree, thereby making the disagreement common belief. This common belief may motivate dialogues to redefine the meaning of formulae that resulted in the disagreement.

Agents may be motivated to persuade others to decide to believe certain propositions, for example, when agents participate in group-plans that require cooperation; they may need to agree on certain propositions for the plan to succeed. Consequently, agents may need to offer propositions to others, and, in response, agree to believe propositions, or, on the other hand, reject to believe these propositions when accepting them would result in inconsistent belief states. Our objective is to present a dialogue game in which cognitive agents become aware of irreconcilable beliefs and show this awareness to others while preserving their private beliefs.

In Section 2, our agent architecture is presented in which agents can have inconsistent beliefs and desires; in addition, a reasoning game is defined allowing agents to decide to believe propositions. A dialogue game is presented in Section 3 enabling agents to offer information. In Section 4, the reasoning game is extended to allow agents to agree to disagree. The resulting formalism permits embedded dialogues, verification of existing dialogues, and a straightforward implementation due to its computational nature (Section 5).

2 Agent Architecture

2.1 Dialogue Games

Early approaches to the semantics of propositions in the philosophy of language centred on the view that semantics are truth-conditional. Searle [17] and Austin

[18] –and many others dating back to Aristotle– adhere to this view, explicitly embracing the idea that truth consists in a relation to reality: a belief is true when there is a corresponding fact, and false when there is no corresponding fact. This is what is called the 'correspondence theory of truth': the truth of a proposition is inseparable from reality [19].

However, demanding strict correspondence between truth and reality makes it rather difficult –if not impossible– for agents to have true beliefs, for then they need to know some segment of reality. A practical and obvious problem with this semantics is that some propositions have an inherent uncertainty: most past and future propositions are uncertain; not to mention the epistemological problem how an agent is to *know* what reality really is. (See [20] on the question "do you believe in reality?")

The correspondence theory has been criticised by Dummett [21] and Wright [22], among others, who advocate a verificationist semantics, lifting the burden for agents from knowing reality to verifying evidence. Agents can derive knowledge from a process of inquiry in which a chain of mental and physical intermediaries connect. However, agents can never be sure whether propositions can be treated as knowledge or as possible false beliefs that need further investigation. The underlying semantics of the correspondence theory, proposed by mathematicians Frege and Tarski [23], has also been criticised by the later Wittgenstein. However, Ellenbogen [24] shows that Dummett's notion that certain propositions which we treat as uncertain also rests upon a realist conception of truth and that "his argument ultimately rests on a refusal to recognize an alternative account of what it is to determine the truth value of a sentence..." [24–p.26]. We will not use the verificationist semantics because it presumes that agents have absolute knowledge of the world; what we will use is a semantics based on "use". In Philosophical Investigations [25], Wittgenstein proposes "use" as an alternative to construct the semantics of propositions.

> *"According to the dictum "meaning is use", what makes it correct to call a statement "true" is not its correspondence with how things are, but our criterion for determining its truth. What it means for us to call a statement "true" is that we currently judge it true, knowing that we may some day revise the criteria whereby we do so."* [24–cover page]

Wittgenstein proposes "language games" with several different however related uses. Our focus is not on how language games can be learned, nor how they refer to a multiplicity of language practices in our ordinary languages. We use its reference to models of primitive language that Wittgenstein has invented to clarify the working of language in general. In his view, communicative acts only have meaning within a particular language game: acts outside a game are just meaningless and useless syntactic structures. Language games provide rules of usage: they define when agents are allowed to pose and answer communicative acts; and, additionally, language games provide rules how the agent's cognitive state changes due to communication. Instead of presupposing that agents need to know reality or other agents' cognitive states to verify factual statements, language games are formulated that describe inquiry with reality or agents. In such inquiry games, currently, results obtained from, for example, microscopes or lie detectors are considered truth-bearers; results obtained from witchcraft are not. However, this may change as science progresses.

We propose two language games, one game that defines the meaning of communicative acts to handle information offers like in contemporary agent communication languages [15, 16], and one game that defines the meaning of decisions to believe propositions [26]. The former will be called a *dialogue game*, the latter a *reasoning game*. Remember that meaning set forward in games is understood as rules of usage. The games provide sets of pre-conditions that define which communicative acts or decisions agents may perform given their current cognitive states. In addition, games provide sets of post-conditions that define the contents of the agent's cognitive state after communicative acts are uttered or decision are made. Pre and post-conditions are combined to form dialogue or reasoning rules.

We assume that information can only accumulate in the participants' cognitive states, and cannot be retracted. In these information-monotonic games, additions may introduce inconsistent beliefs. The reasoning game for deciding to believe propositions stipulates what it means to have inconsistent beliefs: the agent is aware of at least two sources that have sufficient and equal persuasive powers, thereby rendering her belief state inconsistent. To have an inconsistent belief state does not mean there is a segment of reality that corresponds to this belief, but that an agent is convinced by two equally persuasive sources. Although we present a reasoning game that prevents agents from deciding to believe propositions that would result in inconsistent belief states, agents use the possibility of future inconsistencies in a look-ahead fashion in their decisions to believe proposition.

Agents can only speak to one agent at a time via an ideal half-duplex communication channel, which means that no information is lost and that information can only flow in one direction at a time. No restrictions are made on the number of participants in the dialogue, agents are assumed to be omniscient, and aware they use the same dialogue and reasoning games.

2.2 Example Dialogue

Consider the following fictitious dialogue between two Sesame Street puppets. Tv wants to insure his Ferrari. To achieve this, he rings an insurance company and explains the situation by stating his desire to an insurance agent (Ia for short). The Ia wants to sell Tv an expensive insurance policy, because Sesame Street puppets are notoriously prone to fast and dangerous driving, especially in Ferraris. What the Ia wants is that Tv accepts that his car is not safe, which justifies an expensive policy. The dialogue consists of offerings of propositions that are either accepted or declined. For the sake of argument, both agents are rather stubborn and will stick to their first beliefs, and do not accept to believe information that would render their belief state inconsistent.

Example 1 (Dialogue about car insurance in Sesame Street).

1.	Tv to Ia	'My car is a Ferrari.'
2.	Ia to Tv	'Ok.'
3.	Tv to Ia	'(and) My car is safe.'
4.	Ia to Tv	'I don't believe that.'
5.	Ia to Tv	'(actually) I think your car is not safe.'
6.	Tv to Ia	'I don't accept that my car is not safe.'
7.	Ia to Tv	'Do you accept that if a car is a Ferrari then it's not a safe car?'
8.	Tv to Ia	'(no) I don't want to accept that.'
9.	Ia to Tv	'Lets agree to disagree whether your car is safe or not.'
10.	Tv to Ia	'Ok.'

In the remainder of this paper, the example dialogue will be shown to be a valid sequence in our dialogue game starting from initial cognitive states (from Example 3). To achieve this result, the example dialogue is translated to a sequence of communicative acts (Section 3) with propositions taken from an ontology (from Example 2).

2.3 Multi-valued Logics

Whereas in classical logic propositions are assigned truth-values *true* or *false*, later, in philosophy 'deviant' logics were developed that are capable of representing uncertain, non-determined states or epistemic attitudes [27]. More recently, other truth-qualities were required by computer scientists for efficient software implementations, for example, in the verification of circuit board design [28]. In the field of multi-agent systems, computer scientists use non-classical truth-values and semantics, predominantly modal logics, to represent the agents' cognitive states [29]. Our approach is to model the agent's cognitive state with multi-valued logics (MVL). In these logics, propositions are primitive formulae that are assigned multiple truth-values from a bilattice structure [30, 31].

Bilattice Structures. Different modalities are needed to represent the agent's cognitive state. In the MVL introduced in Lebbink et al. [14], propositions are constructed in a fashion that is considered truth-value bearing, capable of being the object of belief or ignorance. This MVL can represent a lack of information (*unknown*) as well as over-informative states (*inconsistent*); the truth-values are taken from a bilattice structure.

A bilattice is an algebraic structure that formalises a space of generalized truth-values with two lattice orderings [30, 31]. The bilattice for a four-valued logic, proposed by Belnap [32], is graphically depicted in Figure 1. Truth-values t and f stand for the classical truth-values *true* and *false* respectively; non-orthodox truth-values u and i represent a complete lack of information (*unknown*) and the inconsistent information state (*inconsistent*). Truth-values are ordered by the amount of truth $\leq_t$ and the amount of information $\leq_k$; currently, only the latter order is of interest. For instance, *unknown* has less information than *true* and *false*, denoted by $\mathsf{u} \leq_k \mathsf{t}$ and $\mathsf{u} \leq_k \mathsf{f}$. Truth-values *true* and *false* are

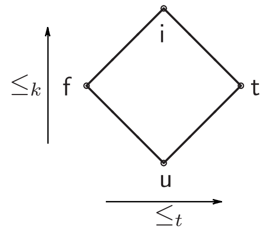

Fig. 1. Smallest complete bilattice for a four-valued logic proposed by Belnap

unrelated to one another in the k-order, that is, $\mathsf{t} \not\leq_k \mathsf{f}$ and $\mathsf{f} \not\leq_k \mathsf{t}$. Bilattices with more truth-values and even a continuum of truth-values can be used to represent biased information or probabilities [30]; we use only the truth-values from Figure 1.

The greatest k-lower bound $\otimes_k$ can be thought of as the truth-value representing the information that is shared by the two truth-values, that is, the mutual information of the two truth-values, for example, $\mathsf{f} \otimes_k \mathsf{t} = \mathsf{u}$. Likewise, the least k-upper bound $\oplus_k$ is thought of as the information that results after combining the two truth-values, for example, $\mathsf{f} \oplus_k \mathsf{t} = \mathsf{i}$. See Ginsberg [30] and Fitting [31] for a formal treatment of the bilattice operators.

Language of MVL. In MVL, atomic and non-atomic MVL propositions are distinguished. The atomic MVL proposition $p{:}\theta$ is a formula p taken from an ontology $\mathcal{O}$ with a truth-value θ from a bilattice structure $\mathcal{B}$. The proposition $p{:}\theta$ is read as “formula p has at least truth-value θ”. The formula is said to have at least the information represented in the truth-value. In the sequel, we will speak of propositions instead of MVL propositions. For our current purposes, an ontology is a set of primitive formulae; for ontologies with more structure, see for example Sowa [33].

The non-atomic proposition $(\psi \rightarrowtail \varphi){:}\theta$ is an inference rule with truth-value θ. Proposition ψ is the antecedent for proposition φ; φ is the consequent of the inference rule. In case θ equals t, the inference rule is written as $\psi \rightarrowtail \varphi$. Remark that $\rightarrowtail$ is not a normal connective but a formula, part of the logical language that codes an inference rule in the object language. Also remark that the truth-values of antecedents and consequents are embedded in the rule. This nesting of sub-sentences is non-standard in MVL, but blurring syntax and semantics will not introduce problems. Other connectives are not defined.

A special purpose proposition $a2d(x, y, \psi, \varphi) : \theta$ states that agent x and y, member of a set of agents $\mathcal{A}$, agree to disagree on propositions ψ and φ, this is further discussed in Section 4. The set of truth-values is denoted C.

Definition 1 (Language of MVL). *Given bilattice* $\mathcal{B} = \langle C, \leq_k, \leq_t \rangle$*, ontology* $\mathcal{O}$*, and agents* $\mathcal{A}$*, the* language of MVL $\mathcal{L}^{\mathcal{B}}$ *is the smallest set satisfying:*

1. *if* $p \in \mathcal{O}$ *and* $\theta \in C$ *then* $p{:}\theta \in \mathcal{L}^{\mathcal{B}}$*,*
2. *if* $\psi, \varphi \in \mathcal{L}^{\mathcal{B}}$ *and* $\theta \in C$ *then* $(\psi \rightarrowtail \varphi){:}\theta \in \mathcal{L}^{\mathcal{B}}$*,*
3. *if* $\psi, \varphi \in \mathcal{L}^{\mathcal{B}}$*,* $x, y \in \mathcal{A}$ *and* $\theta \in C$ *then* $a2d(x, y, \psi, \varphi){:}\theta \in \mathcal{L}^{\mathcal{B}}$*.*

Theories of MVL. Five deduction rules are defined which will be used to construct theories of MVL. Instead of theories of MVL, we will speak of theories.

The complete lack of information associated with truth-value *unknown* always applies to all propositions present in a theory. Remember that $p{:}\theta$ reads that p has at least truth-value θ; therefore, all propositions of the language of MVL have a minimal and unique information state *unknown* in a theory.

$$p{:}\theta \in \mathcal{L}^{\mathcal{B}} \quad \Longrightarrow \quad p{:}\mathsf{u} \in \mathcal{T} \tag{R1}$$

The reading of propositions enforces that if a proposition is part of a theory, then all propositions with less information are also part of the theory. If proposition $p{:}\theta_1$ has at least truth-value θ_1, and θ_2 represents less information than θ_1, then formula p also has at least truth-value θ_2. The information in $p{:}\theta_2$ is said to be subsumed under $p{:}\theta_1$.

$$p{:}\theta_1 \in \mathcal{T} \quad \& \quad \theta_2 \leq_k \theta_1 \quad \Longrightarrow \quad p{:}\theta_2 \in \mathcal{T} \tag{R2}$$

Information is closed in a theory if the least k-upper bound of truth-values of the same formula present in the theory is also present. Remember that the least k-upper bound is thought of as the information that results from combining two truth-values. For example, interpret the theory as an agent's belief state. If the agent believes that p is *true*, and, at the same time, that p is *false*, then the agent also believes that p is *inconsistent*.

$$p{:}\theta_1 \in \mathcal{T} \quad \& \quad p{:}\theta_2 \in \mathcal{T} \quad \Longrightarrow \quad p{:}\theta_1 \oplus_k \theta_2 \in \mathcal{T} \tag{R3}$$

Dual theories are theories with an ordering $\leq_k^\delta$ with $\theta_1 \leq_k^\delta \theta_2 = \theta_2 \leq_k \theta_1$. Due to this reversed order, the least and unique information state from R1 is reversed. All propositions part of the language of MVL have a unique minimal state i in a dual theory.

$$p{:}\theta \in \mathcal{L}^{\mathcal{B}} \quad \Longrightarrow \quad p{:}\mathsf{i} \in \mathcal{T} \tag{R1d}$$

The reading of propositions also enforces subsumed information in dual theories. If a proposition is part of a dual theory, then all propositions with more information are also present in the dual theory.

$$p{:}\theta_1 \in \mathcal{T} \quad \& \quad \theta_2 \leq_k^\delta \theta_1 \quad \Longrightarrow \quad p{:}\theta_2 \in \mathcal{T} \tag{R2d}$$

Theories of MVL are defined as sets of propositions closed under deduction rules. We denote by $Cn^d(\Psi, \mathcal{R})$ the set of propositions that results from $\Psi \subseteq \mathcal{L}^{\mathcal{B}}$ closed under the set of deduction rules $\mathcal{R}$. If ambiguity is unlikely to occur we write $Cn^d(\Psi)$ instead of $Cn^d(\Psi, \mathcal{R})$.

Definition 2 (MVL Theory). *Given a language of MVL $\mathcal{L}^{\mathcal{B}}$, three MVL theories, $\mathcal{T}, \mathcal{T}^c, \mathcal{T}^\partial \subseteq \mathcal{L}^{\mathcal{B}}$ are defined by the closure under deduction rules.*

- *(Normal) theory $\mathcal{T} = Cn^d(\mathcal{T}, \{R1, R2\})$;*
- *Complete theory $\mathcal{T}^c = Cn^d(\mathcal{T}^c, \{R1, R2, R3\})$;*
- *Dual theory $\mathcal{T}^\partial = Cn^d(\mathcal{T}^\partial, \{R1d, R2d\})$.*

Interpretation. An example of a closed theory is an agent's belief state. An agent has at least no information about a proposition, that is, it is not possible to be less informed about a formula p than $p{:}\mathsf{u}$ (Rule R1). If an agent believes that formula p has at least truth-value *inconsistent* then she also believes that p has at least truth-value *true* and *false* (Rule R2). If an agent believes a proposition $p{:}\mathsf{t}$ and she concludes to believe $p{:}\mathsf{f}$, then she also believes $p{:}\mathsf{i}$ which means that she has an inconsistent belief state with regard to formula p (Rule R3).

An example of a dual theory is that of an agent's ignorance state. Consider the situation in which an agent x keeps record of another agent y's ignorance state. If agent x believes that y is ignorant about $p{:}\mathsf{t}$, that is, x believes that y does not believe $p{:}\mathsf{t}$, then x also believes that y is ignorant about $p{:}\mathsf{i}$ (Rule R2d).

2.4 The Agent's Cognitive State

An agent's cognitive state consists of a finite number of mental states, which are theories of MVL. We will not present a full repertoire of all possible mental states agents have regarding themselves and others; only those are identified that are used in the present paper. Remember that set $\mathcal{A}$ denotes the set of agent identifiers.

- Private belief B_x is a complete theory denoting agent x's beliefs. For instance, $p{:}\mathsf{t} \in B_x$ states that x believes that formula p has at least truth-value *true*.
- Private desire to believe D_xB_y is normal theory with $y \in \mathcal{A}$ and *not* $x \neq y$, denoting agent x's desires that agent y is to believe. $\psi \in D_xB_y$ states that agent x desires that agent y is to believe ψ, and $\psi \in D_xB_x$ denotes that x desires to believe ψ.
- Manifested belief B_xB_y is a complete theory with $y \in \mathcal{A}$ and $x \neq y$, denoting the beliefs of y that x keeps record of. For instance, $\psi \in B_xB_y$ states that x is aware that y believes ψ.
- Manifested desires to believe $B_xD_yB_x$ is a normal theory with $y \in \mathcal{A}$ and $x \neq y$, denoting x's awareness of y's desire that x is to believe.
- Manifested ignorance state B_xI_y is a dual theory with $y \in \mathcal{A}$ and $x \neq y$, denoting the propositions that y does not believe that x is aware of. $\psi \in B_xI_y$ states that agent x is aware that agent y is ignorant of ψ.
- In addition, other higher-order manifested mental states are defined likewise. $B_xB_yB_x$ is a complete theory, $B_xB_yI_x$, $B_xI_yB_x$ and $B_xB_yI_xB_y$ are dual theories, $B_xB_yD_xB_x$ and $B_xB_yB_xD_yB_x$ are theories; other mental states like $B_xI_yI_x$ and $B_xI_yD_xB_x$ are not discussed and not used.

The above-mentioned mental states are part of a structure CS which represents the agent's cognitive state: we mean by $CS_x \models \Pi$ the set of set-theoretical propositions Π that hold for agent x's cognitive state. For example, to state that agent x does not believe that agent y believes ψ, and, at the same time does desire that agent y believes ψ, is denoted $CS_x \models \{(\psi \notin B_xB_y),\ (\psi \in D_xB_y)\}$. If Π is a singleton set, it is substituted with its element; for example, we write $CS_x \models \psi \in B_x$ instead of $CS_x \models \{\psi \in B_x\}$. We write $\{CS_x, CS_y\} \models \Pi$ instead of $\forall \pi \in \Pi\ (CS_x \models \pi$ 'or' $CS_y \models \pi)$.

Example 2 (Ontology). To model the dialogue from Example 1, two primitive formulae are needed: $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari})$ and $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})$ which state that the object $\mathsf{this\_car}$ is an instance of the class $\mathsf{ferrari}$, and that it is of class safe, we have: $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe}),\ \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari}) \in \mathcal{O}$.

The cognitive state of all participating agents will be called a collective cognitive state.

Example 3 (Initial collective cognitive state). By default, Tv believes that his car is a Ferrari and a safe one. Next to that, Tv desires that the Ia believes that its car is a Ferrari and safe. On the other hand, the Ia desires by default, that Tv believes its car is not safe; and, in addition, the Ia believes that if a car is a Ferrari then the car is not safe. The agents do not have other desires or beliefs. The superscript 1 denotes the initial cognitive state, consecutive numbers denote cognitive states as the dialogue unfolds. Together with the ontology from Example 2 we have:

$$
\begin{aligned}
CS^1_{\mathsf{ia}} \models \{\ & \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})\text{:}\mathsf{f} \in D_{\mathsf{ia}} B_{\mathsf{tv}}, \\
& \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari})\text{:}\mathsf{t} \rightarrowtail \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})\text{:}\mathsf{f} \in B_{\mathsf{ia}}\ \} \\
CS^1_{\mathsf{tv}} \models \{\ & \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari})\text{:}\mathsf{t} \in D_{\mathsf{tv}} B_{\mathsf{ia}}, \\
& \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari})\text{:}\mathsf{t} \in B_{\mathsf{tv}}, \\
& \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})\text{:}\mathsf{t} \in D_{\mathsf{tv}} B_{\mathsf{ia}}, \\
& \mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})\text{:}\mathsf{t} \in B_{\mathsf{tv}}\ \}
\end{aligned}
$$

2.5 A Reasoning Game to Decide to Believe Propositions

Different definitions when agents are to decide to believe propositions are possible: one could state that agents are allowed to decide to believe a proposition if they themselves believe the criteria to deduce that proposition with an inference rule. We add to this capability the situation in which agents conform to other agents' beliefs.

Cognitive processes are prescribed with reasoning rules that define when agents are allowed to make decisions, and the effects these decisions have on the agents' cognitive state. This is done by specifying sets of pre and post-conditions. Currently, only decisions to add propositions to belief and desire states are possible. Three cognitive processes for making a decision to believe proposition are distinguished: (1) deducing consequences of private beliefs with inference rules; (2) deciding to believe propositions based on other agents' beliefs; and, (3) deducing that one has an irresolvable disagreement with another agent. The latter cognitive process is described in Section 4.

Reasoning Game. A reasoning game is a finite set of reasoning rules that allow agents to make decisions according to the pre and post-conditions of specific decisions. A decision's pre and post-conditions are combined in a reasoning rule, providing the semantics of the decision in a "meaning is use" fashion.

A generic reasoning rule for an abstract decision $\lambda(x, \psi, ms)$ is defined, which can be instantiated with concrete reasoning rules. The decision $\lambda(x, \psi, ms)$ regarding some proposition ψ and mental state ms is allowed for agent x if the set of pre-conditions of $\lambda(x, \psi, ms)$ holds in her cognitive state, that is, $CS_x \models pre(\lambda(x, \psi, ms))$. After the decision is made, the agent's cognitive state is updated, resulting in a new cognitive state in which the post-condition of $\lambda(x, \psi, ms)$ holds, that is, $CS'_x \models post(\lambda(x, \psi, ms))$. Note that the post-conditions never require propositions *not* to be part of mental states; this holds for both reasoning rules and dialogue rules (Section 3.2) both resulting in information monotonic theory updates.

$$CS_x \models pre(\lambda(x, \psi, ms)) \quad \Longrightarrow \quad CS'_x \models post(\lambda(x, \psi, ms)) \tag{RR}$$

We are interested in cognitive states closed under sets of reasoning rules. By $Cn^r(CS_x, \mathcal{R})$ we denote agent x's cognitive state that results from the closure under the set of reasoning rules $\mathcal{R}$. If ambiguity is unlikely to occur, we write $Cn^r(CS_x)$ instead of $Cn^r(CS_x, \mathcal{R})$, and if $\mathcal{R}$ is a singleton set, it is substituted with its element. If the sets of pre and post-conditions are confined to propositions with regard to one mental state, for example, the agent x's belief state, one may want to write $Cn^r(B_x, \mathcal{R})$ instead of $Cn^r(CS_x, \mathcal{R})$. Note that the set $Cn^r(\mathcal{T}, \mathcal{R})$ yields a theory instead of an entire cognitive state.

Deducing Consequences. Agents may deduce new beliefs that are based on their current beliefs. If an agent holds the belief that an inference rule linking two propositions has a designated truth-value *true*, and she believes the antecedent of the inference rule, then the agent may deduce the consequent and add this inferred proposition to her belief state.

If agent x believes ψ and inference rule $\psi \rightarrowtail \varphi$, then she may deduce φ. That is, x may decide to believe φ. These two pre-conditions are part of the reasoning action $d2a_1(x, \varphi, B_x)$ that denotes that x decides to add φ to its mental state B_x. With Reasoning Rule D2A$_1$ we mean the generic Reasoning Rule RR instantiated for this decision to believe a proposition. Different decisions are indexed to distinguish different sets of pre-conditions.

$$(\psi \in B_x),\ (\psi \rightarrowtail \varphi \in B_x) \ \ \in \ \ pre(d2a_1(x, \varphi, B_x))$$

Note that only inference rules with a truth-value *true* are used and that the blurring of syntax and semantics due to nesting of sub-formulas is minimal. Inference rules with truth-value $\theta \not\geq_k \mathtt{t}$, for example *false*, do not have a straightforward interpretation; *true* states that there is a relation between the consequent and the antecedent, *unknown* states that there is no relation. However, what *false* could denote is not clear. (See Rescher [27] on the notion of designated truth-values and consequence relations in MVLs.)

If the pre-conditions of the decision $d2a_1(x, \psi, B_x)$ hold, the agent is allowed to perform the act of deciding to believe ψ, resulting in the state in which ψ is contained in x's belief state. The set of post-conditions is straightforward.

$$(\psi \in B_x) \ \ \in \ \ post(d2a(x, \psi, B_x))$$

Example 4 (Ia decides to believe). Abbreviate ψ for $\mathsf{is\_a(this\_car, ferrari){:}t}$, φ for $\mathsf{is\_a(this\_car, safe){:}f}$ and assume $CS^2_{\mathsf{ia}} \models \psi \rightarrowtail \varphi \in B_{\mathsf{ia}}$. If the Ia is persuaded to believe ψ, then with Reasoning Rule D2A$_1$ the Ia decides to believe φ. However, if $CS^2_{\mathsf{ia}} \models \mathsf{is\_a(this\_car, safe){:}t} \in B_{\mathsf{ia}}$ already holds, then so does $CS^{2'}_{\mathsf{ia}} \models \mathsf{is\_a(this\_car, safe){:}i} \in B_{\mathsf{ia}}$ with Deduction Rule R3.

Next to agents' deducing new beliefs for themselves, agents may also deduce that other agents deduce new beliefs. An agent x may deduce that y should decide to add proposition ψ to her belief state, if x believes that y believes φ and $\varphi \rightarrowtail \psi$. This decision is denoted $d2a_1(x, \psi, B_xB_y)$, which is read as "agent x decides to add ψ to its belief about agent y's beliefs".

$$(\varphi \in B_xB_y),\ (\varphi \rightarrowtail \psi \in B_xB_y) \ \in \ pre(d2a_1(x, \psi, B_xB_y))$$

Agents may even deduce what other agents can deduce about their beliefs.

$$(\varphi \in B_xB_yB_x),\ (\varphi \rightarrowtail \psi \in B_xB_yB_x) \ \in \ pre(d2a_1(x, \psi, B_xB_yB_x))$$

Conformism. The second possibility for an agent to obtain new beliefs is by conforming to other agents' belief states. If an agent believes that another agent believes a proposition, and the agent does not herself believe the proposition, then she may decide to believe the proposition if the other agent is trustworthy. It must be noted that agents have no choice but to decide to believe a proposition: if the pre-conditions of the decision are met, then the agent needs to act accordingly, in effect deciding to perform the decision. The abstract Reasoning Rule RR is instantiated to form the rule for decision $d2a_2(x, \psi, ms)$, this rule is denoted D2A$_2$.

Agent x may decide to believe proposition ψ if x is aware that another agent y believes ψ and x does not already believe ψ. Agent x's decision to believe ψ in conformity to another agent's belief state is denoted by $d2a_2(x, \psi, B_x)$.

$$(\psi \in B_xB_y),\ (\psi \notin B_x) \ \in \ pre(d2a_2(x, \psi, B_x))$$

An additional pre-condition to conforming to another agent's belief state is that agents may only decide to believe propositions if these do not introduce new inconsistencies. Stated differently, for every inconsistent proposition present in a belief state after addition of ψ, holds that this inconsistent proposition was already present before ψ was added.

$$\big(\forall p{:}\mathsf{i} \in \mathcal{L}^{\mathcal{B}}\ (p{:}\mathsf{i} \in Cn^r(B_x \cup \{\psi\}, \mathrm{D2A}_1) \Rightarrow p{:}\mathsf{i} \in B_x)\big) \ \in \ pre(d2a_2(x, \psi, B_x))$$

As a result of the previous pre-condition, agents decide to believe propositions in a first-come-first-serve basis, making the order of uttering communicative acts of importance for the outcome of the dialogue.

Example 5 (Ia decides to believe). Abbreviate ψ for $\mathsf{is\_a(this\_car, ferrari){:}t}$. The Ia may perform a $d2a_1(\mathsf{ia}, \psi, B_x)$ if she does not believe ψ; however, she does believe that Tv believes ψ, $CS^2_{\mathsf{ia}} \models \{\psi \notin B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}}B_{\mathsf{tv}}\}$. In this cognitive state, Reasoning Rule D2A$_2$ is applicable resulting in state $CS^{2'}_{\mathsf{ia}} \models \psi \in B_{\mathsf{ia}}$.

3 A Dialogue Game to Offer Information

3.1 Communicative Acts

The dialogue game to offer information provides semantics for three syntactically different communicative acts. A communicative act $\lambda(x, y, \psi)$ is uttered by speaker x directed to addressee y regarding proposition ψ.

With the communicative act $oba(x, y, \psi)$ the addressee y is offered a proposition ψ with the request to decide to believe it, the act is read as "Are you (y) willing to decide to believe proposition ψ?". The abbreviation oba is short for offer a belief addition. The communicative act $goba(x, y, \psi)$ is read as "I (x) am willing to decide to believe ψ." The addressee can interpret this act as an affirmative answer to an oba; the offer to decide to believe a proposition is granted, hence the abbreviation goba. The communicative act $doba(x, y, \psi)$ is read as "I (x) am not willing to decide to believe ψ." The addressee can interpret this act as a negative answer to an oba; the offer to decide to believe a proposition is denied, the abbreviation doba stands for denying an oba.

Example 6 (Dialogue about car insurance in Sesame Street). In the first line of the dialogue in Example 1, Tv states that its car is a Ferrari. We consider this expression equal to "Are you, Ia, willing to decide to believe that it is true that my car is a Ferrari?" allowing it to be translated to $oba(\mathsf{tv}, \mathsf{ia}, \mathsf{is\_a(this\_car, ferrari):t})$. In response, the Ia decides Tv's offer in line 2. The Ia utters 'Ok.' which is interpreted to be equal to expression "I am willing to decide to believe that your car is a Ferrari." In line 4, the Ia rejects Tv's offer from line 3, the expression "I don't believe that." is interpreted as "no, I am not willing to decide to believe that it is true that your car is safe." A rendition of the dialogue from Example 1 is presented next and used in the remainder of this paper.

1. $oba(\mathsf{tv}, \mathsf{ia}, \mathsf{is\_a(this\_car, ferrari):t})$
2. $goba(\mathsf{ia}, \mathsf{tv}, \mathsf{is\_a(this\_car, ferrari):t})$
3. $oba(\mathsf{tv}, \mathsf{ia}, \mathsf{is\_a(this\_car, safe):t})$
4. $doba(\mathsf{ia}, \mathsf{tv}, \mathsf{is\_a(this\_car, safe):t})$
5. $oba(\mathsf{ia}, \mathsf{tv}, \mathsf{is\_a(this\_car, safe):f})$
6. $doba(\mathsf{tv}, \mathsf{ia}, \mathsf{is\_a(this\_car, safe):f})$
7. $oba(\mathsf{ia}, \mathsf{tv}, \mathsf{is\_a(this\_car, ferrari):t} \rightarrowtail \mathsf{is\_a(this\_car, safe):f})$
8. $doba(\mathsf{tv}, \mathsf{ia}, \mathsf{is\_a(this\_car, ferrari):t} \rightarrowtail \mathsf{is\_a(this\_car, safe):f})$
9. $oba(\mathsf{ia}, \mathsf{tv}, \mathsf{a2d(ia, tv, is\_a(this\_car, safe):f, is\_a(this\_car, safe):t):t})$
10. $goba(\mathsf{tv}, \mathsf{ia}, \mathsf{a2d(tv, ia, is\_a(this\_car, safe):f, is\_a(this\_car, safe):t):t})$

3.2 Dialogue Game

A dialogue game is a finite set of dialogue rules that define when agents are allowed to communicate and how their cognitive states are to be updated afterwards. Similar to reasoning rules from Section 2.5, the pre and post-conditions of communicative acts are combined in dialogue rules to provide the semantics and the rules of usage of the communicative acts.

A generic dialogue rule of a communicative act $\lambda(x, y, \psi)$ states that if all pre-conditions of $\lambda(x, y, \psi)$ hold according to agent x's cognitive state, then $\lambda(x, y, \psi)$ may be uttered to agent y. The post-conditions define the contents of agent x and y's cognitive state after a communicative act is uttered or received.

$$CS_x \models pre(\lambda(x, y, \psi)) \quad \Longrightarrow \quad \{CS'_x, CS'_y\} \models post(\lambda(x, y, \psi)) \tag{DR}$$

3.3 Semantics of Communicative Acts

Information Offer. An agent's motivation to utter a question can be defined as balancing its belief and desire states [13]. Similarly, the motivation to offer information can be defined as an agent balancing her desire regarding another agent's belief sate and her belief state regarding this other agent's belief state. The dialogue rule for the communicative act of offering information is denoted OBA, which is an instantiation of the generic Dialogue Rule DR.

The motivation to offer information regarding proposition ψ is defined as the situation in which agent x has the desire that y believes ψ, and x is not aware that y already believes ψ. In addition, an agent is not allowed to put forward propositions she does not believe. This motivation is part of the pre-conditions to utter an offer.

$$(\psi \in D_x B_y),\ (\psi \notin B_x B_y),\ (\psi \in B_x) \ \in \ pre(oba(x, y, \psi))$$

An information offer is allowed by a speaker x to an addressee y if x is motivated to do so. Given these pre-conditions, addressee y may deduce the following properties of speaker x's cognitive state: x had the desire that y is to believe ψ, x was not aware that y believed ψ, and that the speaker x believes ψ. After the utterance of the offer, the cognitive state of the addressee y has changed according to the following post-conditions.

$$(\psi \in B_y D_x B_y),\ (\psi \in B_y I_x B_y),\ (\psi \in B_y B_x) \ \in \ post(oba(x, y, \psi))$$

A speaker may assume the addressee derives the same post-condition as she would have done if she had received the communicative act herself. Consequently, after uttering an $oba(x, y, \psi)$, speaker x is aware that y is aware that x desires that y believes ψ. In addition, the speaker x is aware that y is aware that x was not aware that y believed ψ, and that the speaker x is aware that the addressee y is aware that x believes ψ. The cognitive state of the speaker x has changed according to the following post-condition.

$$(\psi \in B_x B_y D_x B_y),\ (\psi \in B_x B_y I_x B_y),\ (\psi \in B_x B_y B_x) \ \in \ post(oba(x, y, \psi))$$

In addition to the motivations to communicate are the Gricean maxims that specify principles of cooperative dialogue [34]. These maxims state that utterances of communicative acts should be informative. For example, a speaker is not allowed to ask anything she already believes. Analogously, a speaker is not allowed to put forward information that the addressee already believes as seen from the speaker's perspective. In addition, agents are not allowed to utter communicative acts more than once because of the ideal communication channel in which no information is lost. To realize the restriction that an offer may not be uttered more than once, at least one of the previous post-conditions must not hold. This negated post-condition is added to the set of pre-conditions to restrict the situations in which offers may be uttered. (See Figure 2 for an overview of the pre and post-conditions.)

$$(\psi \notin B_x B_y D_x B_y) \ \in \ \mathit{pre}(\mathit{oba}(x, y, \psi))$$

Example 7 (Information offer). Abbreviate ψ for $\mathsf{is\_a(this\_car, ferrari){:}t}$. The communicative act $\mathit{oba}(\mathsf{tv}, \mathsf{ia}, \psi)$ is allowed in CS^1 (Example 3) with OBA rule resulting in CS^2.

$$\begin{aligned} \mathit{CS}^1_{\mathsf{tv}} &\models \{\ \psi \in D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{tv}},\ \psi \notin B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}}\ \} \\ \mathit{CS}^2_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}} I_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}}\ \} \\ \mathit{CS}^2_{\mathsf{tv}} &\models \{\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} I_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} B_{\mathsf{tv}}\ \} \end{aligned}$$

Granting of an Offer. Next to giving restrictions, the Gricean maxims provide motivations to answer questions and offers; in this case, offers should always be answered by either granting or declining it. The dialogue rule for the communicative act of goba is denoted GOBA, which is an instantiation of the generic Dialogue Rule DR.

An agent x is motivated to utter an accepting response $\mathit{goba}(x, y, \psi)$ if x is aware the addressee y has the desire to make x believe ψ, and, x believes ψ. This results in the following pre-conditions.

$$(\psi \in B_x D_y B_x),\ (\psi \in B_x) \ \in \ \mathit{pre}(\mathit{goba}(x, y, \psi))$$

Given these pre-conditions, addressee y may deduce the following properties of speaker x's cognitive state: x was aware that y desired that x believes ψ, and that x believes ψ. The set of post-conditions for the addressee's cognitive state are the following.

$$(\psi \in B_y B_x D_y B_x),\ (\psi \in B_y B_x) \ \in \ \mathit{post}(\mathit{goba}(x, y, \psi))$$

After the speaker has uttered a $\mathit{goba}(x, y, \psi)$, she may deduce the following properties of the addressee's cognitive state.

$$(\psi \in B_x B_y B_x D_y B_x),\ (\psi \in B_x B_y B_x) \ \in \ \mathit{post}(\mathit{goba}(x, y, \psi))$$

To prevent a superfluous goba from occurring, the speaker should not be aware that she uttered the communicative act before. Agents can be sure about this if at least on of the previous post-condition does not hold. The extra pre-condition reads as "agent x does not believe that y believes that x believes ψ".

$$(\psi \notin B_x B_y B_x) \ \in \ \mathit{pre}(\mathit{goba}(x, y, \psi))$$

Example 8 (Granting an oba). Abbreviate ψ for $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari}){:}\mathsf{t}$. The communicative act $goba(\mathsf{tv}, \mathsf{ia}, \psi)$ is allowed in CS^2 (from Example 7) with GOBA rule resulting in CS^3. Example 5 shows that the Ia decides to believe ψ.

$$
\begin{aligned}
CS^2_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}}\ \} \\
CS^3_{\mathsf{tv}} &\models \{\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}}\ \} \\
CS^3_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}}\ \}
\end{aligned}
$$

Declining an Offer. The motivation to utter a negative response $doba(x, y, \psi)$ to an information offer is similar to an affirmative response, with the main difference that the speaker in case of the negative response does not believe proposition ψ, while in the affirmative response she does. With Dialogue Rule DOBA we mean the generic Dialogue Rule DR instantiated for communicative act doba.

An agent x is motivated to utter a doba if the speaker x is aware the addressee y has the desire to make x believe ψ, and, x does not believe ψ. The preconditions are the following.

$$(\psi \in B_x D_y B_x),\ (\psi \notin B_x)\ \ \in\ \ pre(doba(x, y, \psi))$$

After the communicative act is uttered, the addressee of a doba may deduce properties of the speaker's cognitive state, and the speaker may deduce properties of the addressee's cognitive state. The set of post-conditions for speakers and addressees is the following.

$$\begin{array}{l} (\psi \in B_y B_x D_y B_x),\ (\psi \in B_y I_x), \\ (\psi \in B_x B_y B_x D_y B_x),\ (\psi \in B_x B_y I_x) \end{array}\ \ \in\ \ post(doba(x, y, \psi))$$

To prevent a decline of an oba from being superfluous, at least one of the previous post-conditions must not hold. This criteria reads as "agent x does not believe that agent y believes that x is ignorant about ψ", which is the case if x has not informed y that she does not believe ψ.

$$(\psi \notin B_x B_y I_x)\ \ \in\ \ pre(doba(x, y, \psi))$$

Example 9 (Decline an oba). Abbreviate ψ for $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe}){:}\mathsf{t}$. The communicative act $oba(\mathsf{tv}, \mathsf{ia}, \psi)$ is allowed in CS^3 (from Example 8) with OBA rule resulting in CS^4, from which $doba(\mathsf{ia}, \mathsf{tv}, \psi)$ is allowed with DOBA rule resulting in CS^5.

$$
\begin{aligned}
CS^4_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{ia}} B_{\mathsf{tv}} I_{\mathsf{ia}}\ \} \\
CS^5_{\mathsf{tv}} &\models \{\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{tv}} I_{\mathsf{ia}}\ \} \\
CS^5_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}} D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} I_{\mathsf{ia}}\ \}
\end{aligned}
$$

Motivations to utter	Update of addressee y	Update of speaker x	Restriction to utter
$pre(oba(x,y,\psi))$ $\psi \in D_x B_y$ $\psi \notin B_x B_y$ $\psi \in B_x$	$post(oba(x,y,\psi))$ $\psi \in B_y D_x B_y$ $\psi \in B_y I_x B_y$ $\psi \in B_y B_x$	$post(oba(x,y,\psi))$ $\psi \in B_x B_y D_x B_y$ $\psi \in B_x B_y I_x B_y$ $\psi \in B_x B_y B_x$	$pre(oba(x,y,\psi))$ $\psi \notin B_x B_y D_x B_y$
$pre(goba(x,y,\psi))$ $\psi \in B_x D_y B_x$ $\psi \in B_x$	$post(goba(x,y,\psi))$ $\psi \in B_y B_x D_y B_x$ $\psi \in B_y B_x$	$post(goba(x,y,\psi))$ $\psi \in B_x B_y B_x D_y B_x$ $\psi \in B_x B_y B_x$	$pre(goba(x,y,\psi))$ $\psi \notin B_x B_y B_x$
$pre(doba(x,y,\psi))$ $\psi \in B_x D_y B_x$ $\psi \notin B_x$	$post(doba(x,y,\psi))$ $\psi \in B_y B_x D_y B_x$ $\psi \in B_y I_x$	$post(doba(x,y,\psi))$ $\psi \in B_x B_y B_x D_y B_x$ $\psi \in B_x B_y I_x$	$pre(doba(x,y,\psi))$ $\psi \notin B_x B_y I_x$

Fig. 2. Overview of the pre and post-conditions of the communicative acts

Auxiliary Offer. An auxiliary offer is an information offer that substantiates some claim to believe another proposition. This offer is syntactically indistinguishable from the offer defined in Section 3.3. However, from a semantic perspective the auxiliary offer is a different communicative act, since it has different pre-conditions. Nevertheless, the post-conditions derived from the pre-conditions are not different from those of the ordinary offer. To distinguish between the two offers, auxiliary offers are indexed 2. The dialogue rule is denoted OBA_2.

An agent x may utter an auxiliary offer if she has the desire that another agent y believes some proposition φ, and she is not aware that y already believes φ. These pre-conditions are equal to the motivation of the ordinary offer; however, a number of other pre-conditions are added. Agent x is motivated to utter an auxiliary offer regarding another proposition ψ if according to x, agent y would decide to believe φ if y decides to believe ψ. Agents use Reasoning Rule $\mathrm{D2A}_1$ to deduce properties of other agent's cognitive state, and based on these findings justify auxiliary offers. Formally, if ψ is set-theoretically added to mental state $B_x B_y$, and φ is an element of the closure under the agent y's reasoning rules, then the auxiliary offer is allowed, that is, $\varphi \in Cn^r(B_x B_y \cup \{\psi\}, \mathrm{D2A}_1)$. In addition, agent x should believe ψ, and she should not be aware that y believes ψ.

$$\begin{array}{l}(\varphi \in D_x B_y),\ (\varphi \notin B_x B_y),\ (\psi \in B_x),\\ (\psi \notin B_x B_y),\ (\varphi \in Cn^r(B_x B_y \cup \{\psi\}, \mathrm{D2A}_1))\end{array} \in\ pre(oba_2(x,y,\psi))$$

Example 10 (Auxiliary offer). Abbreviate ψ for $\mathsf{is\_a(this\_car, safe){:}f}$ and φ for $\mathsf{is\_a(this\_car, ferrari){:}t} \rightarrowtail \mathsf{is\_a(this\_car, safe){:}f}$. The communicative act $doba(\mathsf{ia}, \mathsf{tv}, \psi)$ is allowed in CS^5_{ia} (from Example 9) with DOBA rule resulting in CS^6_{tv}, from which $oba(\mathsf{ia}, \mathsf{tv}, \varphi)$ is allowed with OBA_2 rule resulting in CS^7.

$$\begin{aligned} CS^6_{\mathsf{tv}} &\models \{\ \psi \in D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{tv}} B_{\mathsf{ia}},\ \varphi \in B_{\mathsf{tv}},\ \varphi \notin B_{\mathsf{tv}} B_{\mathsf{ia}},\\ &\qquad \psi \in Cn(B_{\mathsf{tv}} B_{\mathsf{ia}} \cup \{\varphi\}, \mathrm{D2A}_1)\ \}\\ CS^7_{\mathsf{tv}} &\models \{\ \psi \in B_{\mathsf{tv}} D_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{tv}} I_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}}\ \}\\ CS^7_{\mathsf{ia}} &\models \{\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} D_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} I_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}}\ \} \end{aligned}$$

4 To Agree to Disagree

Agents in conversation may become aware of parts of their communication partner's cognitive states, and while they do, it may happen that they become aware of irresolvable disagreements. If agents participate in some group-plan that requires mutual agreement on certain proposition, agents have a direct incentive to resolve disagreements regarding these propositions. Although we do not provide an explicit incentive, we do assume that agents have one, and, consequently, they will act to resolve disagreements.

A disagreement can be resolved with four different dialogue games: (1) an agent can convince the agent she disagrees with to believe a proposition that resolves the disagreement. (2) An agent can ask others to help to convince her to believe propositions that resolve the disagreement. (3) An agent can request the agents she disagrees with to forget propositions that result in the disagreement, and (4) an agent can ask others to help to convince her to forget propositions that result in the disagreement. We only consider the first situation: agents can resolve disagreements by convincing others to decide to believe propositions, thereby resolving the disagreement.

If all options to resolve the situation have been exhausted, agents are to conclude that they have an irresolvable disagreement about a specific proposition. If agents offer this awareness to the other agents, both can agree on their disagreement and make the disagreement a manifested belief. This agreement to disagree may trigger a new dialogue in which, for example, a coin flipping method is proposed to resolve the situation, or the meaning of the formula in the proposition is debated. A reasoning rule is defined in Section 4.3 to conclude that an agreement to disagree is in order. This rule combined with the dialogue game to offer information enables agents to agree to disagree.

4.1 Disagreements

Two pieces of information are conflicting when they are not subsumed under each other in the information order. That is, truth-values are in conflict, denoted $\not\cong$, when they are unrelated in $\leq_k$. That is, $\theta_1 \not\cong \theta_2$ iff $\theta_1 \not\leq_k \theta_2$ and $\theta_2 \not\leq_k \theta_1$.

A disagreement between agents x and y about formula p exists if and only if x believes $p{:}\theta_1$ and y believes $p{:}\theta_2$, and the truth-values are in conflict. Additionally, it needs to be the case that both propositions are the most informative, that is: for all truth-values θ_3 part of the bilattice hold that $\theta_3 \leq_k \theta_1$ and $\theta_3 \leq_k \theta_2$, for agents x and y respectively. Note that in a four-valued logic only one disagreement exists: *true* disagrees with *false* because $\mathsf{t} \not\leq_k \mathsf{f}$, and $\mathsf{f} \not\leq_k \mathsf{t}$. If an agent believes $p{:}\mathsf{u}$ and another believes $p{:}\mathsf{t}$ (and these propositions are the most informative), then they do not disagree about p, the latter agent is just more informed than the former naïve agent.

If a disagreement exists between two agents, both need not be aware of this. An agent x is aware she has a disagreement with another agent y if and only if she believes a proposition $p{:}\theta_1$ and she believes that y believes $p{:}\theta_2$ and the truth-values θ_1 and θ_2 are in conflict. A second-order disagreement awareness exists when an agent x is aware that another agent y believes a proposition $p{:}\theta_1$ and x is aware that y is aware that x believes proposition $p:\theta_2$ and θ_1 is in conflict with θ_2.

Example 11 (Disagreement awareness). Abbreviate p for $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})$. In CS^3_{ia} (from Example 8, or Example 6, line 3), the Ia believes that Tv believes $p{:}\mathsf{t}$. From this moment, the Ia is aware of a disagreement about p: $CS^3_{\mathsf{ia}} \models \{p{:}\mathsf{f} \in B_{\mathsf{ia}},\ p{:}\mathsf{t} \in B_{\mathsf{ia}}B_{\mathsf{tv}}\}$. In CS^4_{tv} (from Example 9; or Example 6, line 4) that Tv also becomes aware of the disagreement: $CS^4_{\mathsf{tv}} \models \{p{:}\mathsf{t} \in B_{\mathsf{tv}},\ p{:}\mathsf{f} \in B_{\mathsf{tv}}B_{\mathsf{ia}}\}$. In Example 6, line 5, the Ia becomes aware of a second-order disagreement after she stated $p{:}\mathsf{f}$, because $CS^5_{\mathsf{ia}} \models \{p{:}\mathsf{f} \in B_{\mathsf{ia}}B_{\mathsf{tv}}B_{\mathsf{ia}},\ p{:}\mathsf{t} \in B_{\mathsf{ia}}B_{\mathsf{tv}}\}$. It is only after line 6 that Tv also becomes aware of this disagreement.

4.2 Resolving Disagreements

Assume there is a disagreement between agent x and y about a formula p with $p{:}\theta_1 \in B_x$ and $p{:}\theta_2 \in B_xB_y$. Proposition $p{:}\xi_1$ resolves the disagreement (viewed from x's perspective), if y would decide to believe $p{:}\xi_1$, and, as a result, y would decide to believe $p{:}\theta_1$ (viewed from x's perspective) due to Reasoning Rule $\mathrm{D2A}_1$. We have:

$$p{:}\theta_1 \in Cn^r(B_xB_y \cup \{p{:}\xi_1\}, \mathrm{D2A}_1)$$

The previous proposition resolves the disagreement from x's perspective if y is to decide to believe the proposition. The following proposition resolves the situation from x's perspective if x herself decides to believe the proposition. Proposition $p:\xi_2$ resolves the disagreement, if x would decide to believe $p:\xi_2$, and, as a result, x would believe $p:\theta_2$ (due to Reasoning Rule $\mathrm{D2A}_1$), thereby resolving the disagreement. Formally, we have:

$$p{:}\theta_2 \in Cn^r(B_xB_yB_x \cup \{p{:}\xi_2\}, \mathrm{D2A}_1)$$

An agent is only interested in the least informative proposition to resolve the situation, that is, the proposition with a truth-value that is the lower bound with respect to $\leq_k$. Remember that in the current dialogue game only additions of information are possible; consequently, resolving a disagreement can only take place by adding sufficient information to one of the two agents' cognitive states, rendering it inconsistent.

4.3 Reasoning Rule to Agree to Disagree

The pre-conditions are given next that state when agents are allowed to decide to believe that they agree to disagree. The reasoning rule to become aware of irresolvable disagreements is denoted $\mathrm{D2A}_3$ which is an instantiation of the generic Reasoning Rule RR. The decision is denoted $d2a_3(x, a2d(x, y, p{:}\theta_1, p{:}\theta_2))$, we abbreviate proposition $a2d(x, y, p\theta_1, p{:}\theta_2)$ with κ for convenience. κ denotes that agent x and y agree to disagree on formula p. In the following paragraphs, we call agent x 'I' and y 'you'. The pre-conditions are the following.

1. I am aware that I have a disagreement with you about formula p.

$$(p{:}\theta_1 \in B_x),\ (p{:}\theta_2 \in B_x B_y),\ (\theta_1 \not\cong \theta_2) \ \in\ pre(d2a_3(x, \kappa))$$

2. I am aware that you are also aware of the disagreement.

$$(p{:}\theta_2 \in B_x B_y),\ (p{:}\theta_3 \in B_x B_y B_x),\ (\theta_3 \not\cong \theta_2) \ \in\ pre(d2a_3(x, \kappa))$$

3. I do *not* believe a set of propositions $\Phi \subseteq B_x$ that I have not offered to you before and that could have resolved the disagreement if you had decided to believe them. Suppose $p{:}\xi_1$ is a proposition that if you had added it to your belief state, then the disagreement would have been resolved. For all sets of beliefs Φ that if you had decided to believe them, then this would have resolved the disagreement, that is, $p{:}\xi_1 \in Cn^r(B_x B_y \cup \Phi, \mathrm{D2A}_1)$. However, I have already offered Φ to you, that is, the post-conditions of an oba apply, $\Phi \subseteq B_x B_y D_x B_y$. In this situation I have no methods left (sets of propositions Φ) to persuade you.

$$\begin{array}{r}\big(\forall \Phi \subseteq B_x\ (p{:}\xi_1 \in Cn^r(B_x B_y \cup \Phi, \mathrm{D2A}_1) \ \Rightarrow \\ \Phi \subseteq B_x B_y D_x B_y)\big)\end{array} \ \in\ pre(d2a_3(x, \kappa))$$

4. I am aware that you do *not* believe a set of propositions $\Psi \subseteq B_x B_y$ that you have not offered to me before that could have resolved the disagreement if I had decided to believe them. Suppose $p:\xi_2$ is the proposition that if I added it to my belief state, then the disagreement had been resolved. For all sets of beliefs Ψ that if they had been accepted by me, then this would have resolved the disagreement, $\psi{:}\xi_2 \in Cn^r(B_x B_y B_x \cup \Psi, \mathrm{D2A}_1)$. However Ψ has already been offered to me, and I seem to have responded negative, that is, the post-conditions of a doba apply, $\Psi \subseteq B_x B_y I_x$. In this situation I think that you have no methods (sets of propositions Ψ) left to resolve the situation.

$$\begin{array}{r}\big(\forall \Psi \subseteq B_x B_y\ (p{:}\xi_2 \in Cn^r(B_x B_y B_x \cup \Psi, \mathrm{D2A}_1) \ \Rightarrow \\ \Psi \subseteq B_x B_y I_x)\big)\end{array} \ \in\ pre(d2a_3(x, \kappa))$$

These criteria are the pre-conditions for the reasoning rule to decide to add to a belief state that an agent is stuck in an irresolvable disagreement. If the agent has used the reasoning rule, the post-conditions hold that she actually believes that she agrees to disagree, and that she desires that the agent she disagrees with also believes this proposition.

$$(\kappa{:}\mathsf{t} \in B_x),\ (\kappa{:}\mathsf{t} \in D_x B_y) \ \in\ post(d2a_3(x, \kappa), B_x)$$

The dialogue game to offer information takes care of the communication of κ to y, and possibly reaching an actual agreement on this proposition, making the agreement to disagree common belief.

Example 12 (the Ia and Tv agree to disagree). Abbreviate p for $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{safe})$, ψ for $\mathsf{is\_a}(\mathsf{this\_car}, \mathsf{ferrari}){:}\mathsf{t}$ and κ for $a2d(\mathsf{ia}, \mathsf{tv}, p{:}\mathsf{t}, p{:}\mathsf{f})$. The communicative act

$doba(\mathsf{ia}, \mathsf{tv}, \psi \rightarrowtail p\!:\!\mathsf{f})$ is allowed in CS^8_{ia} (from Example 10) with DOBA rule resulting in CS^9_{tv} from which $oba(\mathsf{ia}, \mathsf{tv}, \kappa)$ is allowed with OBA rule resulting in $CS^1 0$. In the latter state, Tv is allowed to utter $goba(\mathsf{tv}, \mathsf{ia}, \kappa)$.

$$
\begin{aligned}
CS^8_{\mathsf{tv}} &\models \{\, \psi \in D_{\mathsf{tv}} B_{\mathsf{ia}},\ \psi \notin B_{\mathsf{tv}} B_{\mathsf{ia}},\ \varphi \in B_{\mathsf{tv}},\ \varphi \notin B_{\mathsf{tv}} B_{\mathsf{ia}}, \\
&\qquad \psi \in Cn^r(B_{\mathsf{tv}} B_{\mathsf{ia}} \cup \{\varphi\}, \mathrm{D2A}_1) \,\} \\
CS^9_{\mathsf{tv}} &\models \{\, \psi \in B_{\mathsf{tv}} D_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{tv}} I_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{tv}} B_{\mathsf{ia}} \,\} \\
CS^9_{\mathsf{ia}} &\models \{\, \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} D_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} I_{\mathsf{ia}} B_{\mathsf{tv}},\ \psi \in B_{\mathsf{ia}} B_{\mathsf{tv}} B_{\mathsf{ia}} \,\}
\end{aligned}
$$

5 Multi-agent System Architecture

The agent's mental states, reasoning and dialogue games are implemented in SWI-Prolog [35], resulting in flexible multi-agent system architectures. Other reasoning games specifying, for example, when agents are allowed to decide to forget propositions, or dialogue games specifying the semantics of posing and answering questions [14] can be added without the need to change the rules of existing games.

In Section 5.1, the order of execution of the different rules and choices made by agents are described in the agent's deliberation cycle. In Section 5.2, implementations are presented of mental states (MVL theories), reasoning and dialogue rules, which, in Section 5.3, result in a reasoning and dialogue space.

5.1 The Agent's Deliberation Cycle

The agent's deliberation cycle consists of three choices and the execution of five rules, see Figure 3 for a graphical depiction.

1. Check whether cognitive reasoning rules are applicable, select a rule and go to step 2 to execute this rule. If there are no applicable rules, go to step 4 to see whether the agent has received communicative acts.
2. Execute the cognitive reasoning rule that has been found in step 1. The executions of cognitive rules have no observable effects for other agents. Go to step 3 to update the agent's cognitive state accordingly.
3. Execute the appropriate update rule for reasoning rule from step 2. Go to step 1 to check whether more reasoning needs to be done.
4. Check whether communicative acts are received, that is, acts that are directed at the agent. Take the oldest act from the queue of received acts, and go to step 5. If the queue of received acts is empty, go to step 6 to check whether the agent is allowed to utter a communicative act.
5. Execute the appropriate update rule for the received communicative act from step 4. Go to step 1 to check whether reasoning can be done.
6. Check whether dialogue rules are applicable, select a rule and go to step 7 to execute this rule. If there are no applicable rules, go to step 4 to check whether communicative acts have been received.

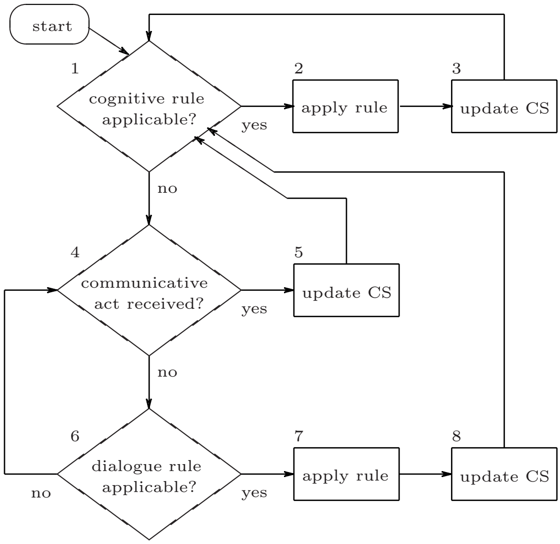

Fig. 3. The agent's deliberation cycle

7. Execute the dialogue rule that has been selected in step 6. The executions have the effect of uttering a communicative act directed at some other agent. Go step 8 to update the agent's cognitive state accordingly.
8. Execute the appropriate update rule for the uttered communicative act from step 7. Go to step 1 to check whether reasoning needs to be done.

Cycle 1-2-3 enforce that an agent's reasoning is done before she engages in conversation; her cognitive state is closed under reasoning before she is to test whether she received communicative acts, that is, $CS_x = Cn^r(CS_x, \mathcal{R})$. Cycle 1-4-5 enforces that all received communicative acts are processed and that the agent's cognitive state is updated accordingly before the agent is to test whether she is can utter communicative acts.

5.2 Implementation in Prolog

Remember that an agent's cognitive state consists of a number of mental states and that these states are theories of MVL. The programming language Prolog, and in particular SWI-Prolog [35], is used to implement the agent's cognitive states, the reasoning and dialogue rules.

A Prolog database is used to store the propositions part of the different theories that compose an agent's cognitive state; the Prolog inference engine is used to enforce that theories are closed under deduction rules as set forward in Definition 2. A Prolog term `prop(T, F, Tv)` states that formula `F` in theory `T` has at least truth-value `Tv`, that is, `F:Tv` $\in$ `T`. The test whether a proposition is part of a theory is implemented as a Prolog call to the database with the corresponding Prolog term for the proposition.

The following Prolog clauses implement complete theories of MVL. Clause leq(k,θ_1,θ_2) implements $\theta_1 \leq_k \theta_2$, and oplus(k, θ_1,θ_2,θ_3) implements $\theta_3 = \theta_1 \oplus_k \theta_2$. The implementation of the bilattice structure is not presented.

```
prop(T, F, u).
prop(T, F, Tv1) :- prop(T, F, Tv2), leq(k, Tv1, Tv2).
prop(T, F, Tv1) :- prop(T, F, Tv2), prop(T, F, Tv3), \+ Tv2 = Tv3,
    oplus(k, Tv1, Tv2, Tv3).
```

Adding a proposition to the agent's cognitive state can be done straightforwardly by asserting the proposition. However, not all propositions need to be asserted, only those that cannot be derived from already asserted propositions.

```
add(T, F, Tv) :- \+ prop(T, F, Tv) -> assert(prop(T, F, Tv)); true.
```

Reasoning and dialogue rules are implemented by taking the conjunction of the pre-conditions of decisions and communicative acts as the body of a clause. The agent's reasoning capabilities are implemented with Prolog's deduction relation ':-'.

Example 13 (The Ia's Initial cognitive state in Prolog). Let ms(b(ia)) denote B_{ia}, and ms(d(ia), b(tv)) denote $D_{\mathsf{ia}}B_{\mathsf{tv}}$. From Example 3 we have:

```
prop(ms(d(ia),b(tv)), is_a(this_car, safe), f).
prop(ms(b(ia)), is_a(this_car, safe), f) :-
    prop(ms(b(ia)), is_a(this_car, ferrari), t).
```

The dialogue rule to offer information is implemented by taking the pre-conditions of the communicative act as the body of a Prolog clause. The update of the speaker and addressee's cognitive state are implemented as a sequence of actions of asserting propositions. Reasoning rules are implemented analogously.

```
dialogue_rule(oba(X, Y, F:Tv)) :-
    prop(ms(d(X), b(Y)), F, Tv),
 \+ prop(ms(b(X), b(Y)), F, Tv),
    prop(ms(b(X)), F, Tv),
 \+ prop(ms(b(X), b(Y), d(X), b(Y)), F, Tv).

update(oba(X, Y, F:Tv)) :-
    add(ms(b(Y), d(X), b(Y)), F, Tv),
    add(ms(b(Y), i(X), b(Y)), F, Tv),
    add(ms(b(X), b(Y), d(X), b(Y)), F, Tv),
    add(ms(b(X), b(Y), i(X), b(Y)), F, Tv).
```

5.3 Dialogue Space

The implementation of the reasoning and dialogue games provides a computational method to generate the space of all cognitive states reachable from an initial collective cognitive state. Although this space tends to become large for even a small number of agents, graphical depiction may give an intuitive feel whether protocols generate sensible communication. This space can be used by trace checkers to prove formal properties, like, for example, whether dialogues terminate in unique states (confluence property), or whether dialogues terminate at all (normalizing property).

The dialogue from Example 6 is a valid sequence of communicative acts in the dialogue game of offering information and reasoning game of deciding to believe propositions. From collective cognitive state from Example 3, the communicative acts of the dialogue are allowed, this is shown with Example 7, 8, 9, 10 and 12.

The complete space of valid dialogues in a dialogue game can be generated with the aid of software tools from an initial collective cognitive state. From the collective cognitive state in Example 3 the space of valid dialogues has been generated (not presented). The resulting graph has 37 nodes representing collective states and 66 edges representing utterances of communicative acts. This space comprises 177 different dialogues with three different final collective states. One has to remember that agents decide to believe propositions if these are consistent with their *current* belief state. This makes the timing of communicative acts of crucial importance, resulting in the three different endings, that is, the dialogue game does not have the confluent property. Resolving this non-confluence is part of future work.

6 Conclusion

In this paper, a formal semantics for an agent's cognitive state is given that allows agents to believe and desire inconsistent propositions. A reasoning rule is formulated enabling agents to decide to believe propositions if they are aware another agent also believes these propositions. Another reasoning rule is given in which agents decide to believe that they disagree with another agent and that this disagreement is irresolvable from both their perspectives.

A dialogue game is proposed for offering propositions and in particular the proposition that both agents agree to disagree. The semantics of the communicative acts are defined by formulating the rules of usage, being the pre-conditions that need to hold in the speaker's cognitive state, and the post-conditions that need to hold after the communicative act is uttered. With a dialogue game a formal system emerges in which sequences of communicative acts can be checked to be valid dialogues. In addition, dialogues spaces can be generated from dialogue rules, providing the possibility to analyse dialogue games on useful properties. One such property is whether unbalanced desire and belief states are resolved in the terminating cognitive states.

The agent's ability to become aware of irresolvable disagreements with some other agent combined with the ability to communicate this information, enables her to agree with the other that the disagreement is irresolvable. Both agents can then settle the disagreement with an agreement to disagree. Note that the agreement to disagree is based only on the cognitive state of the agents that actually have the disagreement.

Dialogue and reasoning games not only define the semantics of decisions and communicative acts, but also provide rules when to generate decisions and communicative acts. This allows straightforward Prolog implementations with intuitive design. Future research will address agents that strategically select which communicative acts to utter with the goal to arrive at a collective state in which desirable properties hold. Other research will centres around communicative acts for retracting information, that is: an offer to forget.

References

1. Parsons, S., Wooldridge, M., Amgoud, L.: Properties and complexity of some formal inter-agent dialogues. Journal of Logic and Computation **13** (2003) 347–376
2. Reed, C.A.: In Demazeau, Y., ed.: Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS 98), Paris, IEEE Press (1998) 246–253
3. Walton, D.N., Krabbe, E.C.W.: Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY Press. Albany, NY, USA (1995)
4. McBurney, P., Parsons, S.: Representing Epistemic Uncertainty by means of Dialectical Argumentation. Annals of Mathematics and Artificial Intelligence, Special Issue on Representations of Uncertainty **32** (2001) 125–169
5. Dignum, F., Dunin-Keplicz, B., Verbrugge, R.: Agent theory for team formation by dialogue. In Castelfranchi, C., Lesérance, Y., eds.: Pre-Proceedings of the Seventh International Workshop on Agent Theories, Architectures and Languages (ATAL 2000), Boston, USA (2000) 141–156
6. Prakken, H.: In Ojeda-Aciego, M., de Guzman, M.I.P., Brewka, G., Pereira, L.M., eds.: Proceedings of the 7th European Workshop on Logic for Artificial Intelligence (JELIA'00), Berlin, Germany, Springer Lecture Notes in AI 1919, Springer Verlag (2000) 224–238
7. McBurney, P., van Eijk, R.M., Parsons, S., Amgoud, L.: Journal of Autonomous Agents and Multi-Agent Systems **7** (2003) 232–273
8. Sadri, F., Toni, F., Rorroni, P.: Logic agents, dialogues and negotiation: and abductive approach. In Schroeder, M., Stathis, K., eds.: Proceedings of the Symposium on Information Agents for E-Commerce, Artificial Intelligence and the Simulation of Behaviour Conference (AISB 2001), York, UK, AISB (2001)
9. Hitchcock, D., McBurney, P., Parsons, S.: A framework for deliberation dialogues. In Hansen, H.V., Tindale, C.W., Blair, J.A., Johnson, R.H., eds.: Proceedings of the Fourth Biennial Conference on the Ontario Society of the Study of Argumentation (OSSA 2001), Windsor, Ontario, Canada (2001)
10. McBurney, P., Parsons, S.: Games that agents play: A formal framework for dialogues between autonomous agents. Journal of Logic, Language and Information **11** (2001) 315–334 Special Issue on Logic and Games.
11. Dignum, F., Dunin-Keplicz, B., Verbrugge, R.: Creating collective intentions through dialogue. Logic Journal of the IGPL **9** (2001) 305–319
12. Hulstijn, J.: Dialogue Models for Inquiry and Transaction. PhD thesis, Universiteit Twente, Enschede, The Netherlands (2000)
13. Beun, R.J.: On the Generation of Coherent Dialogue: A Computational Approach. Pragmatics & Cognition **9** (2001) 37–68
14. Lebbink, H.J., Witteman, C., Meyer, J.J.: Dialogue Games for Inconsistent and Biased Information. Electronic Lecture Notes of Theoretical Computer Science **52** (2003)
15. Labrou, Y.: Stand ardizing Agent Communication. LNCS **2086** (2001) 74–98

16. FIPA: Communicative act library specification. standard sc00037j. Technical report, Foundation for Intelligent Physical Agents (2002)
17. Searle, J.R.: Speech Acts: An Essay in the Philosophy of Language. Cambridge University Press, Cambridge, UK (1969)
18. Austin, J.L.: How To Do Things with Words. Harvard University Press, Cambridge Mass. (1962)
19. Audi, R.: Epistemology: a contemporary introduction to the theory of knowledge. Routledge (1998)
20. Latour, B.: Pandora's Hope: Essays on the Reality of Science Studies. Harvard University Press, Cambridge, Mass (1999)
21. Dummett, M.: What is a theory of meaning? (ii). In: The Seas of Language. Clarendon, Oxford, UK (1993) 34–93
22. Wright, C.: Strict finitism. In: Realism, Meaning and Truth. Blackwell, Oxford, UK (1993) 107–175
23. Coffa, J.A.: The Semantic Tradition from Kant to Carnap : To the Vienna Station. Cambridge University Press (1991)
24. Ellenbogen, S.: Wittgenstein's Account of Truth. SUNY series in philosophy. State University of New York Press (2003)
25. Wittgenstein, L.: Philosophical Investigation. Blackwell Publishers, UK (2001) first published in 1953.
26. Ginet, C.: Deciding to believe. In Steup, M., ed.: Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue. Oxford University Press (2001)
27. Rescher, N.: Many-valued Logic. McGraw-Hill (1969)
28. Butler, J.T.: Multiple-Valued Logic in VLSI. IEEE Computer Society Press (1991)
29. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge University Press (1995)
30. Ginsberg, M.L.: Multiv alued Logics: A Uniform Approach to Reasoning in Artificial Intelligence. Computational Intelligence **4** (1988) 265–316
31. Fitting, M.: Bilattices in logic programming. In Epstein, G., ed.: The Twentieth International Symposium on Multiple-Valued Logic, IEEE (1990) 238–246
32. Belnap Jr., N.D.: A useful four-valued logic. In Dunn, J.M., Epstein, G., eds.: Modern Uses of Multiple-Valued Logic, Dordrecht, D. Reidel (1977) 8–37
33. Sowa, J.F.: Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks/Cole Pub Co (2000)
34. Grice, H.P.: Logic and conversation. In Cole, P., Morgan, J.L., eds.: Speech Acts. Volume 11 of Syntax and Semantics. Academic Press, New York, USA (1975) 41–58
35. Wielemaker, J.: Swi-prolog home page. `http://www.swi-prolog.org/` (2004)