Probabilistic Dynamic Belief Logic for Image and Reputation

Isaac Pinyol[a] Jordi Sabater-Mir[a] Pilar Dellunde[a,b]
[a] *IIIA-Artificial Intelligence Research Institute*
CSIC-Spanish Scientific Research Council
Bellaterra, Barcelona, Spain
{ipinyol,jsabater,pilar}@iiia.csic.es

[b] *Univ. Autònoma de Barcelona*
Bellaterra, Barcelona, Spain
pilar.dellunde@uab.cat

Abstract. Since electronic and open environments became a reality, computational trust and reputation models have attracted increasing interest in the field of multi-agent systems (MAS). Some of them are based on cognitive theories of reputation that require cognitive agents to display all their potential. One of them is Repage, a computational system based on a cognitive theory of reputation that focuses on the study of two important concepts: Image and Reputation. The possible integration of these Repage predicates into a cognitive agent architecture, like the well-known *Belief, Desire, Intention* (BDI) approach implies the definition of these two concepts as mental states, as a set of beliefs. In this paper, we specify a belief logic that captures the necessary elements to express Image and Reputation and we study their interplay by analyzing a classical definition of trust on information sources.

Keywords. Image, Reputation, Probabilistic Dynamic Belief Logic, Multi-agent systems

1. Introduction

Computational trust and reputation models have been recognized as one of the key technologies required to design and implement agent systems [12]. These models provide evaluations of the agents' performances towards specific situations (social evaluations) that agents might use to select partners. In recent years, many models have been developed [14,13], but two main approaches currently exist in the literature. On the one hand, centralized approaches consider reputation and trust to be a global and public property of the agent. They are widely used in online web sites such as eBay, Amazon, etc. On the other hand, distributed approaches consider reputation and trust to be a subjective property of each agent. In this case, this system becomes an important part of the agents architecture.

One of these models is Repage [15], a computational system based on a cognitive theory of reputation [3] that describes a model of REPutation and imAGE. Although both are social evaluations, image and reputation are distinct objects. Image is a simple evaluative belief; it tells that a target agent is *good* or *bad* with respect to a norm, a standard or a skill. Reputation is a belief about the existence of a communicated evaluation.

Repage model entails a tight integration with the agent architecture to exploit all its potential. In particular, cognitive agent architectures will allow the agent to reason not only about the final value of trust or reputation but also about all the individual elements that contribute to that value.

This work is a first step in this direction. A tight integration of Repage model with a cognitive agent must start with the representation as mental states of the main Repage predicates: Image and Reputation. To do so in Section 2 we introduce the concept of social evaluation within the context of Repage model. In Section 3 we define the BC logic that we use as a belief logic. In Section 4 we propose a possible description of Repage image and reputation predicates in terms of our BC logic. In Section 5 we study a condition that makes coincide image and reputation mental states of the agents. In Section 6 we state some of the related work regarding trust formalizations, and finally, we conclude in Section 7 by exposing the conclusions and future work.

2. Social Evaluations in Repage: Image and Reputation

A social evaluation is a generic term used to cover the information referring to the evaluation that a social entity might have about the performance, regarding some skill, standard or norm, of another social entity. A social entity is an active member of the society, like a single agent, a group of agents or institutions.

As we mentioned in the introduction, Repage provides social evaluations as image (what agent believes) and reputation (what agents say). Previous works already show the importance of keeping this distinction [3]. Repage builds images from direct experiences and communicated images from other agents, and reputation only from communicated reputation. The influence between them is done, at the moment, at the pragmatic-strategic level of the agent, letting the agent decide which source of information to use.

In Repage, social evaluations are a simplification of the generic view given in [3]. All of them have an owner of the evaluation, a target agent (the agent being evaluated), a role (the context of the evaluation that encapsulates the behavior being evaluated) and the value of the evaluation (how *good* or *bad*).

The role is the object of the evaluation, the context. The evaluation of an agent playing the role of *buyer* can be totally different from playing the role of *car driver*. This concept of role is similar to the one used in electronic institutions [7]. The value of the evaluation is represented with a tuple of five elements, showing a probability distribution over the labels *Very Bad, Bad, Neutral, Good, Very Good* ({VB, B, N, G, VG} from now on). So, the sum of all values is exactly 1.

Image and reputation predicates are represented as follows:

- $Img_i(t, r, [V_{VB}, V_B, V_N, V_G, V_{VG}])$
- $Rep_i(t, r, [V_{VB}, V_B, V_N, V_G, V_{VG}])$

For instance, the predicate $Img_i(t, seller, [0.5, 0.3, 0.2, 0, 0])$ indicates that agent i evaluates agent t as a seller, and with a probability of 0.5 she acts as *very bad*, with a 0.3, as *bad* and with a 0.2, as *neutral*. Since this social evaluation is an image, this represents what agent i believes. With the same evaluation but as reputation

$Rep_i(t, seller, [0.5, 0.3, 0.2, 0, 0])$ indicates that agent i believes that the evaluation circulates in the society.

In Repage, the role embraces two pieces of information. On one side, it determines the evaluation function, the mapping between the result of a transaction and the sorted set of labels $\{VB, B, N, G, VG\}$, and on the other side, it determines which actions are required to be executed by the agent that intends to evaluate another agent by the specific role. For instance, if the role seller is evaluated with the quality of the attribute obtained after the transaction *buy*, and this quality goes from 0 to 100, we could define VB as qualities between 0 and 20, B as qualities between 20 and 40, and so on.

In the next section we describe a logical framework to express in terms of beliefs these two concepts, allowing then formal reasoning.

3. Defining BC-Logic

3.1. Introduction

The logic we introduce in this section (BC-logic) is a probabilistic dynamic logic with a set of special modal operators B_i, S_{ij} and S_i expressing what is believed by agent i, what has been said from agent i to agent j and what has been said by agent i to however, respectively. The dynamic aspect of this logic is introduced by defining a set Π of actions. Then, for $\alpha \in \Pi$ and $\varphi \in BC$, formulas like $[\alpha]\varphi$ indicate that after the execution of α, the formula φ holds. A very important characteristic of this logic is the inclusion of the special action of communicating a formula. In this case, if φ is a formula and i, j are agents, the expression $[\varphi_{ij}]\phi$ indicates after the communication of φ from agent i to agent j, the formula ϕ holds.

Our language allows explicit reasoning about probability of formulas by means of a new operator Pr, representing the probability of holding a formula. Then, for formulas $\varphi, \phi \in BC$ the expression $Pr\varphi \leq Pr\phi$ indicates that the probability of holding φ is smaller or equal to the probability of holding ϕ. Furthermore, we allow one side of the inequality to be a constant $\overline{r}$ where $r \in [0, 1] \bigcap \mathbb{Q}$. Our language is based on the Logic of Knowledge and Probability introduced by Fagin and Halpern in [8].

BC-logic allows to express formulas like $B_i(0.8 \leq Pr([\alpha]\varphi))$, meaning that agent i believes that the probability of holding φ after the execution of action α is at least 0.8. Thereby, the formula $S_i(0.8 \leq Pr([\alpha]\varphi))$ expresses the same but in terms of what agent i has said. Notice that there is not a necessary implication between both concepts. Like in human societies, people might communicate information that they do not believe.

3.2. The Syntax

To introduce BC-logic we start by defining a countable set P of propositional variables, a finite set Π_0 of atomic programs and a finite set $\mathcal{A}$ of agent identifiers of cardinality n. The set of formulas $Fm(BC)$ and the set Π of programs are defined in Backus-Naur extended form as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid \phi \wedge \varphi \mid [\alpha]\phi \mid \\ B_i\phi \mid S_{ij}\phi \mid Pr\varphi \leq Pr\phi \mid \overline{r} \leq Pr\phi \mid Pr\varphi \leq \overline{r}$$

and

$$\alpha ::= \pi \mid \alpha;\beta \mid \alpha \cup \beta \mid \alpha^*$$

where $p \in P, \pi \in \Pi_0$, ϕ, φ are formulas, α, β are programs, i, j agent's identifiers and $r \in [0,1] \bigcap \mathbb{Q}$.

From now on we assume that a subset $\Pi_1 \subseteq \Pi_0$ of the set of atomic actions are a fixed set of speech acts. Thus, our language contains formulas of the following form $[\varphi_{ij}]\phi$ where φ_{ij} is a BC-formula.

We write $\varphi \vee \phi$ as an abbreviation for $\neg(\neg\varphi \wedge \neg\phi)$, $\varphi \rightarrow \phi$ as an abbreviation for $\varphi \vee \neg\phi$, $\varphi \leftrightarrow \phi$ for $(\varphi \rightarrow \phi) \wedge (\phi \rightarrow \varphi)$, $\bot$ as an abbreviation for $\neg\top$. Finally let $S_i\phi$ be a short cut for $\vee_{j \in \mathcal{A}} S_{ij}\phi$ and $S\varphi$ a short cut for $S_{i_1}\varphi \wedge \cdots \wedge S_{i_n}\varphi$ where $\{i_1, \ldots, i_n\} = \mathcal{A}$. We write $a = b$ as an abbreviation of $a \leq b \wedge b \leq a$.

Also, we would like to point out that this language definition allows an arbitrary number of nested modal operators. When thinking about modeling agent's mind we are very interested in allowing meta-beliefs, so, beliefs on others' beliefs. In this way, agents can model other agents' minds. Following the same idea, why not allowing agents modeling other agents' minds but from the point of view of another third agent? This would require three-level nested beliefs. Even if in real applications more than three levels could seem not useful, from a theoretical point of view we want to keep this possibility open.

3.3. The Semantics

We want to give to the logic a probabilistic interpretation. For this, semantics of BC-logic are given by means of Kripke structures of the following form $M = \langle W, \mathcal{F}, \{R_\alpha : \alpha \in \Pi\}, \{R_{B_i} : i \in \mathcal{A}\}, e, \zeta, \mathcal{C}\rangle$ where:

- W is a non-empty set of possible worlds.
- $\{R_\alpha : \alpha \in \Pi\}$ are the accessibility relations for actions, for each $\alpha \in \Pi$, $R_\alpha \in 2^{(W \times W)}$.
- $\{R_{B_i} : i \in \mathcal{A}\}$ are the accessibility relations for the belief operators, for each agent i, $R_{B_i} \in 2^{(W \times W)}$.
- $e : P \rightarrow 2^W$ assigns to each propositional variable a set of worlds.
- $\mathcal{C} : W \rightarrow 2^{\Pi_1}$ assigns to each world a repository of BC-formulas.
- $\zeta : W \rightarrow PS$ assigns a probability space to each world.

where PS is the class of all probability spaces $(W, \mathcal{G}_w, \mu_w)$ such tha $\mathcal{G}_w$ is the field of all subsets of W and the function $\mu_w : 2^W \rightarrow [0,1]$ is a finitely additive probability measure on 2^W, that is, $\mu_w(\emptyset) = 0$, $\mu_w(W) = 1$ and whenever $X, Y \in 2^W$ and $X \cap Y = \emptyset$ we have $\mu_w(X \cup Y) = \mu_w(X) + \mu_w(Y)$.

Definition A Kripke model $M = \langle W, \mathcal{F}, \{R_\alpha : \alpha \in \Pi\}, \{R_{B_i} : i \in \mathcal{A}\}, e, \zeta, \mathcal{C}\rangle$ is *regular* iff for every $\alpha, \beta \in \Pi$: (1) $R_{\alpha;\beta} = R_\alpha \circ R_\beta$, (2) $R_{\alpha\cup\beta} = R_\alpha \cup R_\beta$ and (3) $R_{\alpha^*} = R^*_\alpha$ (where $*$ is the ancestral of the relation R_α).

From now on we will only consider regular Kripke models as semantics for our logic. Let M be a model for the logic and $w \in W$ a possible world, given $p \in P$, $\varphi, \phi \in BC$[1], $\alpha \in \Pi$, $i, j \in \mathcal{A}$ and $r \in [0,1] \bigcap \mathbb{Q}$, the truth-value of a BC-formula on model M is defined as follows:

---

[1]We will write BC instead of $Fm(BC)$ to refer to the set of all well-formed formulas

- $M, w \models \top$
- $M, w \models p$ iff $w \in e(p)$
- $M, w \models \varphi \wedge \phi$ iff $M, w \models \varphi$ and $M, w \models \phi$
- $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$
- $M, w \models B_i\varphi$ iff $\forall w_k : (w, w_k) \in \mathcal{R}_{B_i}$ implies $M, w_k \models \varphi$
- $M, w \models S_{ij}\varphi$ iff $\varphi_{ij} \in \mathcal{C}(w)$
- $M, w \models [\alpha]\varphi$ iff $\alpha \in \Pi$ and $\forall w_k : (w, w_k) \in \mathcal{R}_\alpha$ implies $M, w_k \models \varphi$
- $M, w \models Pr\varphi \leq \overline{r}$ iff $\mu_w(\{w_k | M, w_k \models \varphi\}) \leq r$
- $M, w \models \overline{r} \leq Pr\phi$ iff $r \leq \mu_w(\{w_k | M, w_k \models \varphi\})$
- $M, w \models Pr\varphi \leq Pr\phi$ iff $\mu_w(\{w_j | M, w_j \models \varphi\}) \leq \mu_w(\{w_k | M, w_k \models \phi\})$

The semantics of formulas of the form $S_i\varphi$ is introduced by means of the function $\mathcal{C}$, that assigns to every world the set of sentences communicated among the agents up to this moment. Then, $S_i\varphi$ is true if and only if agent i has said φ to somebody. Notice that $[\varphi_{ij}]$ is a modal operator but $S_i\varphi$ is not[2].

Notice that in the semantics we assign a probability space for each world. This is the most general case. In certain conditions and depending on the context we can force a semantic condition stating that all worlds have the same probability space. Nevertheless, in the context of reputation models we are interested in keeping the general case, since some knowledge about probability distributions may depend on the information that the reputation model offers, and this information is totally dynamic and may change.

3.4. Axiomatization

We present now some axioms for BC-logic although our purpose is not to provide a complete axiomatization. Our first axioms and rules are those of classical propositional dynamic logic plus the standard axioms K,D, 4 and 5 of modal logic and necessitation rules for the B_i operators.

- BK: $B_i(\varphi \rightarrow \phi) \rightarrow (B_i\varphi \rightarrow B_i\phi)$
- BD: $B_i\varphi \rightarrow \neg B_i\neg\varphi$
- B4: $B_i\varphi \rightarrow B_iB_i\varphi$
- B5: $\neg B_i\varphi \rightarrow B_i\neg B_i\varphi$

Regarding the operators S_ij, we include the following axioms:

- S1: $[\varphi_{ij}]S_{ij}\varphi$
- S2: $S_{ij}\varphi \rightarrow B_jS_{ij}\varphi$
- S3: $S_{ij} \rightarrow [\alpha]S_{ij}$, for every $\alpha \in \Pi$

For probabilistic formulas, we include the following axioms:

- P1: $Pr\top = 1$
- P2: $Pr\bot = 0$
- P3: $0 \leq Pr\varphi$
- P4: $Pr((\varphi \wedge \phi) = a) \wedge Pr(\varphi \wedge \neg\phi) = b \rightarrow Pr\varphi = a + b$
- P5: $B_i(Pr\varphi = a) \wedge B_i(Pr(\varphi \rightarrow \phi) = b) \rightarrow B_i(max(a + b - 1, 0)) \leq Pr\phi$
- P6: $Pr\varphi = Pr\phi$, if $\varphi \leftrightarrow \phi$ is a theorem.

[2] For the sake of clarity we have written formulas like $S_i\varphi$ and $Pr\varphi \leq \overline{r}$ instead of $S_i\lceil\varphi\rceil$ and $Pr\lceil\varphi\rceil \leq r$. We have not make explicit that formulas in this context are reified, avoiding the use of many-sorted languages

$P1$ and $P2$ state the probabilities of $\top$ and $\bot$. $P3$ claims for the non-negativity of probability formulas. $P4$ is the additivity axiom. $P5$ is the equivalent Lukasiewicz implication for multimodal logic for beliefs and probabilistic formulas. In fact this axiom can be deduced from the previous ones, but we make it explicit because it is very useful for the reasoning process we will achieve in the future work. $P6$ is distributivity. As inference rule we include also: from φ it can be derived $Pr\varphi = 1$.

4. Grounding Image and Reputation

At this point our interest relies on giving to Repage predicates a description in terms of the BC logic we introduced in Section 3. However, we need to introduce first, the nomenclature we will use in the Repage domain.

4.1. Some Notation

Having the finite set $\mathcal{A}$ of agent identifiers and the finite set $\mathcal{R}$ of role identifiers, the actions are determined by the possible roles and possible agents. As a matter of simplification, we assume that each role r has associated only one generic action, $\Phi(r)$, that at a certain moment of time T may be executed to some agent j, $\Phi(r)[j]_T$. Also, Repage encapsulates the way agents evaluate outcomes from transactions. In this sense, the evaluation of an agent as a seller can be determined, for instance, by the quality of the product obtained and the delivery time. The proposition $\delta_T(r).q$ will refer to the value of the attribute q of the role r obtained in transaction T.

To illustrate this, we could define $\Phi(seller)$ as the action of *buy* and $\Phi(seller)[John]$ as the action of buying to John. We will write it as $Buy(John)$. Then, if the evaluation of the role *seller* is done thorough the attributes quality and delivery time, we will write the proposition $\delta_T(seller).quality$ and $\delta_T(seller).deliveryTime$ to refer to the respective values.

4.2. Image and Reputation Predicates

Let i, j be agent identifiers and r a role Repage image and reputation predicates are represented as $Img_i(j, r, [V_{VB}, V_B, V_N, V_G, V_{VG}])$ and $Rep_i(j, r, [V_{VB}, V_B, V_N, V_G, V_{VG}])$. Agent i is the agent that has generated the predicate, and therefore, that uses Repage. Agent j is the target of the evaluation. Vector $[V_{VB}, V_B, V_N, V_G, V_{VG}]$ represents the five probabilistic values that cover the full space of possible outcomes, which are classified as *Very Bad*, *Bad*, *Neutral*, *Good* and *Very Good*. Following the definition of Image, we have that a Repage image and reputation predicates can be expressed with the following set of beliefs written in BC logic:

Image	Reputation
$B_i(V_{VB} = Pr([\Phi(r)[j]]\Psi_{VB}))$	$B_i(S(V_{VB} = Pr([\Phi(r)[j]]\Psi_{VB})))$
$B_i(V_B = Pr([\Phi(r)[j]]\Psi_B))$	$B_i(S(V_B = Pr([\Phi(r)[j]]\Psi_B)))$
$\ldots$	$\ldots$

The expression Ψ_X is the propositional formula that depends on the specific role r and that evaluates the possible outcome obtained after the execution of action $\Phi(r)[t]$ (the conditions for which an outcome is classified as *Very Bad*, *Bad* etc...). For instance, is Repage has generated the predicate $Img_i(S1, seller, [.5, .2, .1, .1, .1])$, and the role seller is evaluated with the attribute quality (from 0 to 100) obtained after the transaction *buy*, the set of beliefs describing the mental state of the agent regarding seller $S1$ could be:

- $B_i(0.5 = Pr([Buy(S1)]0 \leq \delta(seller).quality < 20))$
- $B_i(0.2 = Pr([Buy(S1)]20 \leq \delta(seller).quality < 40))$
- $\ldots$

If the agent is cognitive, these beliefs participate in the deliberation process. Notice that this mental state does not say anything about which potential seller is *better*. This will be determined by the set of desires. Notice that dealing with more or less condition levels (other that five: VB,...,VG) would not represent a big change.

5. Image, Reputation and Their Interplay

One of the key point of Repage and the cognitive theory of reputation that underlies it [3] is the relationship between image and reputation. The theory states that both are social evaluations but distinct objects. With the representation we give for image and reputation in BC logic this difference depends on the relationship between the belief operator B and the operator S. As a matter of fact, the inclusion of axioms relating both concepts would generate a typology of agents. We discuss some of them in the following subsections.

5.1. Honest and Consistent Agents

Let $i \in \mathcal{A}$, we say that the agent i is honest if the formula $S_i\varphi \rightarrow B_i\varphi$ is included in her theory (when she says something, she believes it), and we abbreviate it as $h_i\varphi$. In a similar way, a consistent agent i will hold in her theory the formula $S_i\varphi \rightarrow \neg S_i\neg\varphi$ (when she says something, she never says the opposite) abbreviated as $c_i\varphi$. It is easy to prove that $h_i\varphi \rightarrow c_i\varphi$.

The definition of honesty allows agents to model what other agents think (meta-beliefs) in terms of what they say. In this sense, if the formula $B_i(h_j\varphi)$ holds (agent i believes that agent j is honest), applying definitions and axiom K for B operator we obtain $B_iS_j\varphi \rightarrow B_iB_j\varphi$.

5.2. Trusting Agents

The concept of trust has different connotations and many definitions. From a cognitive point of view, trust can be seen as a mental state [2] of a particular agent that believes that another agent has certain property [6][3]. Considering trust on agents as information sources, Demolombe in [6] introduces six definitions of trust: Sincerity, Credibility, Vigilance, Validity and Completeness. The following definition of trust coincide with the Validity definition, that at the same time, is a conjunction between Sincerity and Credibility:

[3]This vision of trust is somehow a simplification of the general view given in [2] where trust is considered a mental attitude composed of beliefs and goals and strongly related to the concept of delegation.

Let i, j be agents, we define $Trust_{i \to j}\varphi$ as $B_i(S_j\varphi \to \varphi)$. The formula states that agent i trusts agent j when agent i beliefs that whatever j says is true. Applying axiom K, the formula becomes $B_iS_j\varphi \to B_i\varphi$. This means that if agent i believes that j has said something, agent i will believe the same thing. This notion of trust is very strong. In fact, the generalization of this formula to an agent that *trusts* everybody, make collapse the definition of image and reputation we gave for Repage predicates.

Proposition 5.1 *Let i be an agent, if $B_i(S\varphi \to \varphi)$ holds, then the beliefs describing reputation predicates from Repage collapse with the beliefs describing image predicates.*

Proof The proof is quite direct. Applying axiom K to $B_i(S\varphi \to \varphi)$ we obtain $B_iS\varphi \to B_i\varphi$. The antecedent of the implication coincides with the definition of reputation predicates we gave. Since φ is an arbitrary formula, applying modus ponens to each one of the formulas used to describe reputation predicates, for instance, $B_i(S(V_{VB} = Pr([\Phi(r)[j]]\Psi_{VB})))$ we obtain $B_i(V_{VB} = Pr([\Phi(r)[j]]\Psi_{VB}))$. The set of all these new beliefs coincide with the definition of image that we gave.

This result states that we have a condition that makes logically equivalent the mental state of an agent holding an image and a reputation. Notice that the condition $B_i(\varphi \to S\varphi)$ collapses image mental state with reputation mental state.

6. Related Work

Some current state-of-the-art logics inspired us for defining the BC logic. However, none of them seems to be expressive enough for the needs we have described in this paper. The probabilistic and dynamic notions have been mostly treated in epistemic logic ([10], [8]), and in a simpler way in belief logic [1]. Furthermore, some formalizations of trust using belief logic have been done [11], where trust is related to information acquisition in multi-agent systems, but in a crisp way. Similar to this, in [5], modal logic is used to formalize trust in information sources, also with crisp predicates. Here, actions and communicated formulas are also used.

Regarding fuzzy reasoning on trust issues, in [9] it is defined a trust management system in a many-valued logic framework where beliefs are graded. Also, in [4] it is proposed a logic that integrates reasoning about graded trust (on information sources) and belief fusion in multi-agent systems. Our logic does not use graded beliefs. Instead, we use the notion of beliefs on probability sentences, because when deal with image and reputation it seems a more accurate option.

7. Conclusion and future work

In this work we have introduced a probabilistic dynamic belief logic to capture the mental states of agents holding image and reputation predicates as defined in Repage model. The logic seems to be expressive enough to describe them and to provide a logical framework in which to define a typology of agents.

In a short period of time, our plans include the study and tentative proof of completeness and soundness of the axiomatization we have given for BC logic. Also, we plan to study more conditions that make image and reputation influence each other at the level of beliefs. In particular, we are very interested in the redefinition of the *trust* predicate by including a grade: $Trust^g_{i \rightarrow j}\varphi$ as $B_i(Pr(S_j\varphi \rightarrow \varphi) = g)$. Also, we plan to study this relationship across the reminding definitions of trust: *Vigilance*, *Cooperativeness*, *Completeness* and the combination of them.

We plan to use this logic as a fundamental part of a BDI agent where desires, intentions and plans are build taking as a base this logic. The importance of the BC logic relies on that from this moment on, we have a logical framework that allows us to express all what we need referring to social evaluations.

8. Acknowledgments

This work was supported by the European Community under the FP6 programme (eRep project CIT5-028575 and OpenKnowledge project FP6-027253), by the projects Autonomic Electronic Institutions (TIN2006-15662-C02-01) and Agreement Technologies (CONSOLIDER CSD2007-0022, INGENIO 2010) and by the Generalitat de Catalunya, under the grant 2005-SGR-00093. Also, we would like to give special thanks to Robert Demolombe for his useful comments.

References

[1] A. Casali, Ll. Godo, and C. Sierra. Graded models for bdi agents. In J. Leite and P. Torroni, editors, *CLIMA V, Lisboa, Portugal*, pages 18—33, 2004. ISBN: 972-9119-37-6.

[2] C. Castelfranchi and R. Falcone. Social trust. In *Proceedings of the First Workshop on Deception, Fraud and Trust in Agent Societies, Minneapolis, USA*, pages 35—49, 1998.

[3] R. Conte and M. Paolucci. *Reputation in artificial societies: Social beliefs for social order*. Kluwer Academic Publishers, 2002.

[4] R. Demolombe and C.J. Liau. A logic of graded trust and belief fusion. In *Proc. of the 4th Workshop on Deception, Fraud and Trust in Agent Societies*, pages 13–25, 2001.

[5] R. Demolombe and E. Lorini. A logical account of trust in information sources. In *Eleventh International Workshop on Trust In Aagent Societies*, 2008.

[6] Robert Demolombe. To trust information sources: a proposal for a modal logical framework. *Trust and deception in virtual societies*, pages 111–124, 2001.

[7] M. Esteva, J. A. Rodriguez-Aguilar, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specifications of electronic institutions. In *Agent Mediated Electronic Commerce, The European AgentLink Perspective.*, pages 126–147. Springer-Verlag, 2001.

[8] R. Fagin and J.Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.

[9] T. Flaminio, G. Michele Pinna, and E. B. P. Tiezzi. A complete fuzzy logical system to deal with trust management systems. *Fuzzy Sets Syst.*, 159(10):1191–1207, 2008.

[10] B. P. Kooi. Probabilistic dynamic epistemic logic. *J. of Logic, Lang. and Inf.*, 12(4):381–408, 2003.

[11] C. J. Liau. Belief, information acquisition, and trust in multi-agent systems: a modal logic formulation. *Artif. Intell.*, 149(1):31–60, 2003.

[12] M. Luck, P. McBurney, O. Shehory, and S. Willmott. *Agent Technology: Computing as Interaction (A Roadmap for Agent Based Computing)*. AgentLink, 2005.

[13] S.D. Ramchurn, D. Hunyh, and N.R. Jennings. Trust in multi-agent systems. *The Knowledge Engineering Review*, 1(19):1—25, 2004.

[14] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artif. Intel. Rev.*, 24(1):33–60, 2005.

[15] J. Sabater-Mir, M. Paolucci, and R. Conte. Repage: Reputation and image among limited autonomous partners. *J. of Artificial Societies and Social Simulation*, 9(2), 2006.