Metareasoning and Social Evaluations in Cognitive Agents

Isaac Pinyol and Jordi Sabater-Mir

IIIA - Artificial Intelligence Research Institute
CSIC - Spanish National Research Council
Bellaterra, Barcelona, Spain
{ipinyol,jsabater}@iiia.csic.es

Abstract. Reputation mechanisms have been recognized one of the key technologies when designing multi-agent systems. They are specially relevant in complex open environments, becoming a non-centralized mechanism to control interactions among agents. Cognitive agents tackling such complex societies must use reputation information not only for selecting partners to interact with, but also in metareasoning processes to change reasoning rules. This is the focus of this paper. We argue about the necessity to allow, as a cognitive systems designers, certain degree of freedom in the reasoning rules of the agents. We also describes cognitive approaches of agency that support this idea. Furthermore, taking as a base the computational reputation model Repage, and its integration in a BDI architecture, we use the previous ideas to specify metarules and processes to modify at run-time the reasoning paths of the agent. In concrete we propose a metarule to update the link between Repage and the belief base, and a metarule and a process to update an axiom incorporated in the belief logic of the agent. Regarding this last issue we also provide empirical results that show the evolution of agents that use it.

Keywords: Reputation, Trust, Cognitive Agents, Metareasoning, BDI agents.

1 Introduction

Reputation mechanisms have been recognized one of the key technologies when designing multi-agent systems (MAS) [1]. In this relatively new paradigm, reputation models have been adapted to confront the increasing complexity that open multi-agent environments bring. Thus, the figure of agents endowed with their own private reputation model takes special relevance as a non-centralized mechanism to control interactions among agents. Following this line, cognitive agents using cognitive reputation models arise as one of the most complete and generic approaches when facing very complex societies. Usually, cognitive agent's architectures, like BDI (*Belief, Desire, Intention*), follow logic-based reasoning mechanisms, providing then a high flexibility and theoretically well-founded reasoning.

Repage [2] is a reputation system based on a cognitive theory of reputation that has been used in logical BDI reasoning processes [3] (BDI+Repage), offering then an integrated reasoning framework. Even when this work faces the field of computational

A.V. Vasilakos et al. (Eds.): AUTONOMICS 2009, LNICST 23, pp. 220–235, 2010.

reputation models, the focus is not on the model itself, but on the integration of the information that it provides with the other elements of the agent. Following this line, a very important aspect of cognitive agents is the capacity to reason about their own reasoning processes. These metareasoning processes act at all levels in the agents' mind. However, we are interested in the aspects related to reputation information.

In this concrete work we justify the use of metarules and metaprocesses in cognitive agents and provide mechanisms to specify them. We apply these ideas to some rules and axioms of the BDI+Repage model. In concrete, we propose a specification that allows the modification at run-time of the rule that relates reputation information with logic belief formulas, and also of an axiom rule integrated in the logical belief based of the agent. Regarding this last point we provide both a metareasoning process to update such axiom, and empirical results of an implementation we develop of the BDI+Repage model placed in a replication of a simple market, populated by buyers, sellers and informant agents. To detail our work, in section 2 we introduce a cognitive theory of agency to justify the use of metarules and metareasoning when modeling cognitive agents. In section 3 we explain the agent model, and in section 4 we provide the tools to specify metareasoning rules and processes. In the same section we apply them to some reasoning rules of the BDI+Repage model. In section 5 we present empirical results to show how the dynamic modification of axioms can produce good results in the agents' level of satisfaction. Finally in section 6 we conclude our analysis and present the future work.

2 Reasoning and Metareasoning: A Cognitive Approach

In this section we briefly get in touch with the cognitive theory that supports the work done in this paper. The theory developed by Castelfranchi and Paglieri [4] is quite generic and focuses on the dynamics of goals processing and its relation with the beliefs of the agent. Although the specific topic of the paper relies on describing which *typology* of beliefs participates in each stage of goal dynamics, we are very interested in the concepts of belief-supporting goals and belief-supporting beliefs that are pointed out by the theory. The authors argue that goals and beliefs have a supported structure of beliefs, i.e., beliefs from which a given goal or belief is activated. Moreover, such structure can be explicitly represented also as beliefs, achieving then a metabelief level.

In our work we are specially interested in the idea of belief-supporting beliefs which has been deeply studied in [5]. When such belief structures are explicitly represented as beliefs, the authors named them *reasons*. Thus, a given belief has a set of *reasons* that continuously supports it. Because of this explicit representation, agents can also reason about them, achieving then a *metareasoning* process that also relies on beliefs.

Reasons are important because allow the agent to *justify* herself *why* the set of beliefs and goals are activated, but also to justify her beliefs to other agents. This last issue has been extensively studied in the field of argumentation [6]. Reasons are also important because they allow the agents to review their own reasoning process, by starting metareasoning processes that *change* the reasoning paths.

A logical perspective is a good starting point for the explanation of more technical details. Lets consider a propositional logic in which proposition p holds. If formula $p \rightarrow q$ also holds, then for sure q will hold. From the logical view this reasoning has been produced by *modus ponents* by two explicit formulas (p and $p \rightarrow q$). If we are talking about the beliefs of an agent, the belief on q (Bq) is justified by the beliefs Bp and $B(p \rightarrow q)$. Notice that if the agent realizes that p was not the case, then q must be withdrawn. This is the typical belief revision process. However, the agent may also realize that it is not the case that $p \rightarrow q$. Thus, q should be also withdrawn. However, from a cognitive point of view this situation is quite different than the first one, since the formula $p \rightarrow q$ is a *reason*, an explicit belief saying that from p can be deduced q. Therefore, a reasoning concerning the truth of $p \rightarrow q$ could be seen as a metareasoning. It could be argued that the truth of p also affects the reasoning process, but the object of the reasoning is not a representation of an explicit reasoning step, as it is in the other case. See figure Figure 4 for a graphical representation of the structure.

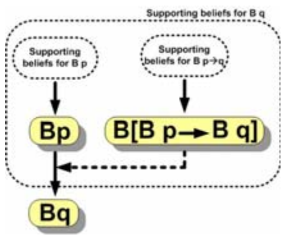

Fig. 1. The generic belief-supporing beliefs structure with explicit representation of a deduction step

The point of the discussion is that logics offers a nice way to construct trees of supporting formulas, through logical reasoning, but in the general case, the links can be modified by the same agent due to other beliefs. As a designers of cognitive agents architectures, we must deal with these concepts to consider real autonomous agents. In this paper we extend a BDI agent architecture that incorporates reputation information [3] by allowing a partial update of the rules that govern the reasoning process, focusing on the belief-supporting beliefs, letting for future work the relationship between desires and intentions. Next section provides the description of the BDI+Repage model.

3 A Multicontext BDI Agent with Repage System

The model of agent we present in this section is a BDI model in which Repage reputation system is also incorporated. To explain it is necessary first to get in touch with the Repage reputation model, and the cognitive theory of reputation that supports it.

3.1 Preliminaries: Social Evaluations, Image and Reputation

Repage [2] is a computational system designed to be part of the agents architecture and based on a cognitive theory of reputation [7]. It provides social evaluations as image and reputation. A social evaluation is a generic term used to encapsulate the information resulting from the evaluation that an agent (evaluator) might make about another agent's (target's) performance regarding some skill, standard or norm (object of the evaluation). The object of the evaluation relies on which property of the target agents is evaluated. The value of the evaluation indicates how *good* or *bad* the performance resulted to be.

A social evaluation in Repage has three elements: a target agent, a role and a probability distribution over a set of labels. The target agent is the agent being evaluated. The role is the object of the evaluation and the probability distribution the value of the evaluation. The evaluator is the agent making the social evaluation.

The role uniquely identifies a kind of transaction and the classification of the possible outcomes. The current implementation of Repage considers five abstract linguistic labels for this classification: *Very bad, Bad, Neutral, Good, Very good*(*VB, B, N, G, VG* from now on), and assigns a probabilistic value to each label, however, we generalize it considering a finite number of labels $w_1, w_2 \ldots$. The *meaning* of each label must be contextualized depending on the role. For instance, we can represent a Repage image predicate as $img_i(j, seller,$[0.4, 0.2, 0.2, 0.1, 0.1]$)$. This indicates that agent i holds an image predicate about agent j in the role of $seller$, and the value of the evaluation is $[0.4,\ 0.2,\ 0.2,\ 0.1,\ 0.1]$. This value reflects a probability distribution over the labels *VB, B, N, G, VG*. Then, it means that agent i believes that in the transaction of buying, when agent j acts as a seller, there is a probability of 0.4 to achieve a VB result (in the context of this transaction, this may mean a very low quality of the product), with a probability of 0.2 a B result, etc. For reputation predicates, it is the same as image, but instead, the agent believes that the evaluation is said by all or most of the agents in the group. We refer to [2] for details on the calculus and the internal architecture.

In the next subsection we detail the BDI model, starting from the basic framework of multicontext systems.

3.2 The Multi-Context BDI Model

Multi-context systems (MCS) provide a framework that allows several distinct theoretical components to be specified together with the mechanisms that link them together [8]. These systems are composed of a set of contexts (or units), and a set of bridge rules. Each context can be seen as a logic and a set of formulas written in that logic. Bridge rules are the mechanisms with which to infer information from one context to another. Each bridge rule has a set of antecedents (preconditions) and a consequent. The consequent is a formula that becomes true in the specific context when each antecedent holds in its respective context.

The specification of our BDI agent as a multi-context system is inspired by the models presented in [9,10]. It is formalized with the tuple $Ag = \langle\{BC,\ DC,\ IC,\ PC,\ CC,\ RC\},\ \triangle_{br}\rangle$. These correspond to Belief, Desire, Intention, Planner, Communication and Repage contexts respectively. The set of bridge rules $\triangle_{br}$ incorporates the rules $1, 2, 3, 4, P, Q$ and B, the bridge rules A_I and A_R shown in Figure 3, and rule B. Figure 2 shows a graphical representation of this multi-context specification. In the next subsections we briefly explain each context and bridge rule.

3.3 Belief Context (BC)

This context contains the beliefs of the agent. For this we use BC-logic [11], a probabilistic dynamic belief logic with a set of special modal operators. We are specially

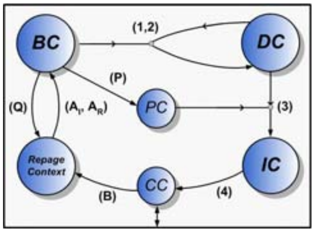

Fig. 2. The Repage context embedded in a multi-context BDI agent. Circles represent context and arrows represent bridge rules.

interested in the operators B_i and S, the first expressing what is believed by agent i, and the latter, what has been said by all the agents in the group 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.

This logic incorporates specific axioms to reason about the probabilities of formulas by means of the operator Pr and constants $\overline{p}$ such that $p \in [0,1] \bigcap \mathbb{Q}$. It follows that for formulas $\varphi \in BC$, the expression $\overline{p} \leq Pr\varphi$ indicates that the probability of holding φ is higher or equal to p. This logic is based on the Logic of Knowledge and Probability introduced by Fagin and Halpern in [12].

BC-logic allows expressions like $B_i(\overline{p} \leq Pr([\alpha]\varphi))$. This indicates that agent i believes that the probability of holding φ after the execution of action α is at least p. Thereby, the formula $S(\overline{p} \leq Pr([\alpha]\varphi))$ expresses the same but in terms of what all agents have said. To simplify the notation, we will write expressions like $B_i(\overline{p} \leq Pr\varphi)$ as $(B_i\varphi, p)$, and $S(\overline{p} \leq Pr\varphi)$ as $(S\varphi, p)$.

This logic allows us to express image information in terms of beliefs $B_i\varphi$, and reputation information in terms of beliefs about what is said, $B_iS\varphi$ (see section 3.8). By grounding image and reputation into simple elements, we endow the agent with a powerful tool to reason over these concepts.

The complete syntax, semantics and axiomatization of BC-logic can be found at [11]. The belief operator follows the standard K, D, 4 and 5 axioms of modal logic, while operator S has its owns. The most interesting axioms are those that describe the interaction between S and B_i. These are closely related to the concept of *trust* that Demolombe in [13] defined regarding agents as information sources. The relationship of the two operators implies a relation between image and reputation at the belief level [11]. For instance, if for every φ the formula $B_i((S\varphi \rightarrow \varphi), p)$ holds (trust axiom), then agent i believes that what all agents say is really true with a probability p. The trust axiom has big implications in the relation between image and reputation information at the belief level [11].

3.4 Desire Context (DC)

This context deals with the desires of the agent. Like the BDI model described by Rao and Georgeff in [14], they are attitudes that are explicitly represented and that reflect the general objectives of the agent. We consider that desires are graded, and for that, we use the multi-valued logic (DC-logic) based on the Lukasiewicz logic described in [10].

DC-logic includes two fuzzy modal operators[1]: D_i^+ and D_i^-. The intended meaning of $D_i^+\varphi$ is that the formula φ is desired by agent i, and its truth degree, from 0 (minimum) to 1 (maximum), represents the level of satisfaction if φ holds. The intended meaning of $D_i^-\varphi$ is that φ is negatively desired, and the truth degree represents de level of disgust if φ holds. Also, DC-logic includes truth constants $\overline{r}$ where $r \in [0,1] \bigcap \mathbb{Q}$, and the connectives $\&$ and $\Rightarrow$ corresponding to the Lukasiewicz conjunction and implication respectively.

3.5 Intention Context (IC)

This context describes the intentions of the agent. Like in the Rao and Georgeff's BDI model [14], intentions are explicitly represented, but in our case generated from beliefs and desires. Also, we consider that intentions are graded, and for this we use the IC-logic defined in [10].

Similar to DC-logic, IC-logic defines the fuzzy modal operator $I_i\varphi$, indicating that agent i has the intention to achieve φ, and its truth degree (from 0 to 1) represents a measure of the trade-off between the benefit and counter-effects of achieving φ. Furthermore, IC-logic is defined in terms of a Lukasiewicz logic in the same way as DC-logic. Also, formulas like $\overline{r} \Rightarrow I_i\varphi$ will be written as $(I_i\varphi, r)$.

3.6 Planner Context (PC) and Communication Context (CC)

The logic in the Planner context is a first-order logic restricted to Horn clauses. In this first approach, this context only holds the special predicate $action$, which defines a primitive action together with its precondition. We look forward to introducing plans as a set of actions in the future. Communication context is a functional context as well, and its logic is also a first-order logic restricted to Horn clauses with the special predicates $does$ (to perform actions), and rec_{ij} (to notify that agent i has received a communication from agent j).

3.7 Repage Context (RC)

The Repage context contains the Repage model. We can assume that Repage predicates are specified in first-order logic restricted to Horn clauses, where the special predicates Img and Rep are defined. We write them as $img_i\ (j, r, [V_{w_1}, V_{w_2}, \ldots])$ and $rep_i\ (j, r, [V_{w_1}, V_{w_2}, \ldots])$, corresponding to the Image and Reputation of agent j playing the role r, from the point of view of i.

When in Repage the role and its labeled weights are defined, the role uniquely identifies which kind of transaction is part of, and each w_k identifies a predicate. To simplify, we can assume that the transaction identified by a role is summarized in a single action. To state this, we presuppose the definition of a mapping $\mathcal{R}_r$ between each role r and its

[1] The original logic in [10] does not contain the reference to the agent. We include it to remark the desires of agent i.

action. In a similar way, we assume a mapping $\mathcal{T}_{r,w_k}$ between each role r and label w_k to a predicate.

We illustrate this with an example: In a typical market, the transaction of buying certain product involves two agents, one playing the role of buyer and the other playing the role of seller. From the point of view of the buyer, if she wants to evaluate other agents that play the role of seller, she knows that the associated action is buy. So, $\mathcal{R}_{seller}$ maps to buy. In the same way, the agent must know the meaning of each label w_k of Repage. Then, we can define that $\mathcal{T}_{seller,w_1}$ is $veryBadProduct$, $\mathcal{T}_{seller,w_2}$ is $okProduct$, etc.

In this mapping, the Repage predicate $img_i(j, seller, [0.2, 0.3, \ldots])$ indicates that agent i believes that there is a probability of 0.2 that after executing the action $\mathcal{R}_{seller}$ (buy) with agent j as a seller, she will obtain a $\mathcal{T}_{seller,w_1}$ ($veryBadProduct$); with 0.3 that she will obtain $\mathcal{T}_{seller,w_2}$ ($OKproduct$), etc. With reputation predicates it is similar, but the concept is quite different. In this case it indicates that agent i believes that the corresponding evaluation is said by the agents in the group.

3.8 Bridge Rules

Bridge rules A_I and A_R (see Figure 3) are in charge of generating the corresponding beliefs from images and reputations respectively. Notice that given a Repage social evaluation, these bridge rules generate one belief for each weight w_k. Both bridge rules use the belief operator (B_i) over certain formula, but meanwhile rule A_I states a knowledge that agent i believes as true, A_R states a knowledge that agent i believes to be said. They follow the definition of image and reputation we have given in the Repage context in section 3.7.

The detail of the following rules can be found at [3]. Rules 1,2,3,4 perform the actual BDI reasoning. Bridge rules 1 and 2 transform generic desires to more concrete and realistic desires. To do this, these bridge rules merge generic desires from DC (with absolute values of satisfaction or disgust) with the information contained in BC, which includes the probability to achieve the desire by executing certain action. The result is a desire whose gradation has changed, becoming more realistic. This is calculated by the function g. If we define it as the product of both values, we obtain an expected level of satisfaction/disgust. Notice that we require that the belief information implies the achievement of the desired predicate.

$$\mathbf{A}_I\mathbf{:}\quad \frac{RC : img_i(j, r, [V_{w_1}, V_{w_2}, \ldots])}{\begin{array}{c} BC : (B_i([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1})) \\ BC : (B_i([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2})) \\ \ldots \end{array}}$$

$$\mathbf{A}_R\mathbf{:}\quad \frac{RC : rep_i(j, r, [V_{w_1}, V_{w_2}, \ldots])}{\begin{array}{c} BC : (B_i(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1}))) \\ BC : (B_i(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2}))) \\ \ldots \end{array}}$$

Fig. 3. The bridge rules A_I and A_R (see Figure 2). They translate Image and Reputation predicates respectively into beliefs expressions in BC.

Bridge rule 3 generates intentions. It takes into account both the expected level of satisfaction and the cost of the action. At the same time, executing an action to achieve certain formula can generate undesirable counter-effects. Thus, bridge rule 3 also takes into account the possible negative desires that can be reached by executing this action. In this bridge rule, for each positive realistic desire (D^+), we must include all negative desires (D^-) that can result from the same action. In this way we have the value of the positive desire (δ^+) and the sum of all negative desires (δ^-) that can be achieved by executing the same action. The strength of the intention that is created is defined by a function f. Different f functions would model different behaviors. In our examples we use the following definition: $f(\delta^+, \delta^-) = max(0, \delta^+ - \delta^-)$.

Finally, bridge rule 4 instantiates a unique intention (the one with maximum degree) and generates the corresponding action in the communication context.

4 The Metalevel Specification

In this section we specify a possible metalevel reasoning regarding the trust axiom of the BC-logic and the bridge rules A_I and A_R. For this task we take ideas from the specification of dynamic protocols [15] in the frame of open multiagent organizations. Here, the specification of interaction protocols is described as a set of rules specified at the design time. However, when facing open systems, often environmental or social conditions for instance, may carry the necessity to modify such protocols at run-time. These modifications must be product of a dialog, as a metaprotocol, among the participants. In [15], the author presents an infrastructure to allow agents the modification of a subset of rules. It considers a k-level infrastructure, where at level 0, the main rules of the protocol are specified with certain *degrees of freedom* (Dof). At level 1, a metaprotocol can be specified to allow the discussion about how to change the protocol of level 0. More levels can be specified following the idea that at level i the protocol allows the discussion of the degrees of freedom of level $i - 1$.

4.1 DoF for Reasoning Rules

We apply the same DoF principle to some axioms and bridge rules of our BDI+Repage architecture. Instead of using belief revisions techniques, we encourage the use of DoF to update parts of rules that govern a reasoning process, not only for preserving consistencies, but also for adaptation. Belief revision processes rely on crisp logic and look for the smallest subset of formulas to keep a logical theory consistent when a formula is added in the theory. For our needs this vision is limited because only faces logical theories and because is used to avoid inconsistencies.

By using degrees of freedom, we bound the space of states by constraining what can be modified and what not. Thus, a main reasoning structure remains constant, but not static. In the case of logic-based BDI agents, this is very clear. For instance, the original model that Rao and Georgeff presented [14] states some basic and untouchable axioms to ensure several properties of the logical reasoning, but also considers other set of axioms that when included in the logic, model totally different behaviors. A clear example is the relation between the three main attitudes: beliefs, desires and intentions.

They define a typology of agents, the main ones being realist, strong realist and weak realist agents.

This is one of the advantages of using BDI models. The flexibility they achieve. By simply adding or erasing some axioms we can model an infinity of agents. However, when facing autonomous agents that must deal with open environments, we need some more flexibility. In Rao and Georgeff's BDI model, could an agent move from a strong realism to weak realism at run-time? From a technical point of view it is just a matter of changing two axioms. From a logical point of view, this process is outside the logic, and must be done at a meareasoning level. The possibility to update or modify some axioms is supported by the cognitive theory presented in the introduction of this paper, in which real autonomous agents should be aware of the way they reason. Due to that, agents can *think* about how they *think* and act in consequence (see section 2 for more details).

Notice that in the model of Rao and Georgeff the switch between strong realism and weak realism implies the totally substitution of a set of axioms for another set. This is the most extreme scenario in which the DoF involves the whole rule, because of the nature of the logic, which is crisp. More complex and expressive logics, like the BC-logic presented in the previous section, can deal with probabilities, which can be incorporated in the axioms to somehow tune their strength in the reasoning process, for instance. In a similar way, brige rules, which in fact are outside the logic, can be also tune by similar elements.

In the following sections we show how a similar formalism used for the DoF of dynamic protocols can be used to specify a metareasoning model for our BDI+Repage model.

4.2 A Metalevel Specification for the Rules A_I and A_R

In this subsection we focus on the relationship between the Repage model and the Belief context. This relation is statically specified by the rules A_I and A_R. As we mentioned, these rules are responsible for translating Image and Reputation predicates into atomic beliefs. Following the specifications of the Repage reputation model and the underlying theory, these rules are a very accurate formalism to generate the belief that in a more atomic way represent the information provided by the reputation model. However such transformation may carry out the logical inconsistency on the belief theory. This is the case when the Trust axiom is present and we have very contradictory information between an image and a reputation predicate of a given agent in a given role.

These inconsistencies always refer to probabilistic issues. To illustrate this, we can assume that the trust predicate is present in the belief context as $B_i(S\varphi \rightarrow \varphi)$. Then, it may occur that rules A_I and A_R have generated the following beliefs:

$$B_i([buy(alice)]VeryBad, 0.9) \quad (1)$$

$$B_i([buy(alice)]VeryGood, 0.1) \quad (2)$$

$$B_iS([buy(alice)]VeryBad, 0.5) \quad (3)$$

$$B_iS([buy(alice)]VeryGood, 0.5) \quad (4)$$

Due to the trust axiom, the formulas 3 and 4 imply the following formulas:

$$B_i([buy(alice)]VeryBad, 0.5) \tag{5}$$
$$B_i([buy(alice)]VeryGood, 0.5) \tag{6}$$

Notice that formula 1 implies formula 5, and formula 6 implies formula 2. The inconsistency relies on that propositions $veryBad$ and $veryGood$ should be mutually disjoint. Then, how is it possible to belief that with a probability higher that 0.9 after execution of the action $buy(alice)$ we will obtain a very bad product, and with a probability higher that 0.5 we will obtain a very good product?

To solve this kind of situations we provide the agent with the capability to modify its bridge rules. To do so, we define one degree of freedom at each one of the rules A_I and A_R. By doing this we are specifying metarules (M(A_I), M(A_R)), which map to a family of different A_I and A_R rules:

$$\textbf{M(A}_I\textbf{):} \quad \frac{RC : img_i(j, r, [V_{w_1}, V_{w_2}, \ldots])}{\begin{array}{c} BC : (B_i(pr([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1})) = \mathcal{X}_{j,r,1}) \\ BC : (B_i(pr([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2}) = \mathcal{X}_{j,r,2})) \\ \ldots \end{array}}$$

$$\textbf{M(A}_R\textbf{):} \quad \frac{RC : rep_i(j, r, [V_{w_1}, V_{w_2}, \ldots])}{\begin{array}{c} BC : (B_i(pr(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1})) = \mathcal{Y}_{j,r,1})) \\ BC : (B_i(pr(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2})) = \mathcal{Y}_{j,r,2})) \\ \ldots \end{array}}$$

Notice that if we set the default value of $\mathcal{X}_{j,r,1}$ and $\mathcal{Y}_{j,r,1}$ to 1, we have exactly the original rules A_I and A_R, since in BC-logic, $\varphi \leftrightarrow pr(\varphi) = 1$. The process by which the agent decides which value to take is a metareasoning process. Different heuristics can be used to perform such task. What it is clear is that such heuristics is a process that depends on a set of beliefs (Trust axiom and the beliefs that refer to the same agent and role must be part of the inputs of such process).

4.3 A Metalevel Specification for the Trust Axiom

In a similar way, the generic trust axiom that relates what is said with what is believed can be modified, and in fact, it is crucial for the adaptation of the agent. On the one side, if no trust axiom is present in the theory, formulas like $B_i S\varphi$ will never become $B_i\varphi$. On the other side, if the trust axiom is present, $B_i S\varphi$ would imply $B_i\varphi$. However, we also talk about graded trust, $B_i(pr(S\varphi \rightarrow \varphi) = g)$, and the effects on the formulas like $B_i S\varphi$ and $B_i(S(pr(\varphi) \geq p))$[2]. Different values of g model different behaviors of the agent. Thus, we can consider this g as a DoF of the axiom. We can write the meta axiom *M(Trust)* as

$$B_i(pr(S\varphi \rightarrow \varphi) = \mathcal{Z})$$

[2] It can be proved that if the graded trust axiom is present, and the formula $B_i(S(pr(\varphi) \geq p))$ holds, then it can be deduced $B_i(pr(\varphi) \geq p \cdot g)$ [11].

As well, different heuristics as metaprocesses can be considered for the update of the DoF. In section 5 we consider an heuristics for this axiom and show how agents behavior changes during time.

Notice that the inclusion of such axiom with a $\mathcal{Z}$ higher that 0 may cause inconsistencies in the theory, as we explained in the previous subsection. Because of that, this process can start another process to update the values of $\mathcal{X}_{j,r,p}$ and $\mathcal{Y}_{j,r,p}$ agent j, role r and weight p.

4.4 Processes Description

Figure 4 shows a graphical representation of the metarules dependences. Circles $M1$ and $M2$ represent processes:

- **M1:** This process is in charge for deciding the DoF $\mathcal{Z}$, which belongs to $M(Trust)$. In the experimental section we state how this concrete process could be performed. In any case, the graphical representation shows that this process is fed by the actual instantiation of the trust axiom (which is only characterized by the DoF variable $\mathcal{Z}$), and information from the Repage context. The output of the process is a new value for $\mathcal{Z}$ (in the graphic is shown as $\mathcal{Z}'$).
- **M2:** This process only receives the current instantiation of the Trust predicate, characterized by $\mathcal{Z}$. As we argued, this value is the only that can produced inconsistencies in the theory of BC-Context. As output, it provides new values from the DoF $\mathcal{X}_{j,r,p}$ and $\mathcal{Y}_{j,r,p}$, for each different agent j, role r and weight p. Thus, the number of different instantiated rules is computed by $|Ag| \cdot |R| \cdot |W|$. This means that potentially, each agent in each role can have a different inference rule from Repage to BC-Context.

We leave the exploration of the metaprocess $M2$ for future work. Instead, we focus on $M1$, providing a possible mechanism to compute $\mathcal{Z}$.

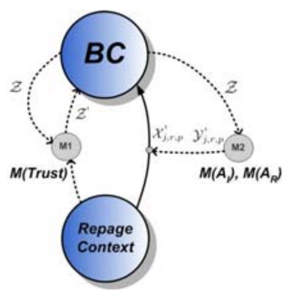

Fig. 4. The metalevel specification. Dot lines specify the metalevel reasoning. Notice that all them come from the belief context.

5 Experimentation

In this section we propose a concrete solution for the $M1$ process, to update the trust axiom represented by the metarule

$$B_i(pr(S\varphi \to \varphi) = \mathcal{Z})$$

As previously shown, this axiom plays a crucial role in the relationship between image and reputation predicates. Different values of the degree of freedom produce a typology of agents. When $\mathcal{Z}$ is 0, the agent only takes into account image information. When $\mathcal{Z}$ is 1, reputation information is as valuable as image information in terms of the impact that the information has in the mind of the agent. Previous work on cognitive theories and simulation of image and reputation dynamics [16,7] reveals that the amount of reputation information that circulates in a society is a lot higher than image-based information, due to the implicit commitment that sending image information carries out.

However, even when reputation information is mostly inaccurate, open societies perform *better* when reputation information is allowed in the system, and also are more robust with respect to certain level of cheating information[3]. This indicates that agents face mostly inaccurate information but that they need to use it to face real uncertain and unpredictable scenarios.

These studies are very helpful when defining a process to decide $\mathcal{Z}$. Our trust axiom is in fact a predicate that indicates how much information that circulates in the society can be considered true. In the way we have defined rules A_I and A_R, settings of $\mathcal{Z}$ tending to 0 could be useful when the number of cheaters is considerably big, meanwhile settings of $\mathcal{Z}$ close to 1 would be helpful in the opposite way.

5.1 Scenario and Simulation Settings

We replicate a simple market where in the society we have a set of buyer, seller and informant agents. In this scenario, all sellers offer the same product, which has a certain quality going from 0 (minimum) to 100 (maximum). Also, a delivery time expressed in weeks is associated with the seller. These agents are completely reactive and sell the products on demand. Buyers are BDI agents following the model described in this paper. Therefore, the main goals of the agents are described in terms of graded desires.

The set of informant agents send out reputation information about the sellers. We control the experiment by setting a percentage of informants that spread *bad* reputation, the number of sellers and the distribution of qualities and delivery times.

The performance of the buyers is evaluated by the level of satisfaction obtained after their decision. As we mentioned, buyer agents state their preferences by a set of graded desires. These desires can be positive or negative, and each one of them has a grade. After an action is performed, the agent receives the fulfillment of the interaction, obtaining the real quality of the product and the delivery time. This information is compared with the objectives of the agent. The *level of satisfaction* of the agent is calculated by summing the grades of positive achieved desires and subtracting the grades of the achieved negative desires.

At each turn buyers need to perform an action. In this case, they need to buy a product to some of the available sellers. To simulate the fact that reputation information is more present than image information, at each turn all informants send reputation information

[3] More that 50% of cheaters in a society still produces a benefit in the overall performance when reputation is allowed.

to the buyers. In this experiments we do not consider image communications. Therefore, image information is only calculated through direct experience. In this sense, at each turn one direct experience is contrasted with N reputation communications from the informants (where $N > 1$).

In the specification, we are considering the evolution of a single buyer with 10 sellers and 5 informants. We execute 10 times each experiment and consider the average level of satisfaction for each turn. We state a distribution of qualities and delivery times in such a way that the best qualities and best delivery times are very scarce. If these properties are the norm, the society does not need the exchange of information, since a random choice from the buyer would get already a very good seller[4]

5.2 Static Experiments

It is easy to show the effects of a fix trust axiom in different situations. Figure 5 shows the accumulated average level of satisfaction obtained by a buyer at each turn in an environment where all informants are honest, and when all informants are liars, considering $Z = 0$ and $Z = 1$. Since when $Z = 0$ reputation information is not taken into account, the performance in this case does not depend on the quality of the reputation information.

The graphic shows that when $Z = 1$, in the case of a scenario with honest informants, the level of satisfaction obtained by the agent considerable increases with respect to the case in which $Z = 0$. Assuming normality in the data, from the turn 10, the difference is already statistically significant with a 95% of confidence (p_value≤ 0.05), and from the turn 20 on, the difference becomes significant with a 99% of confidence (p_value$\leq$ 0.01).

Also, when $Z = 1$ and in the scenario all informants spread false reputation, the performance of the buyer decreases considerably with respect to the case in which $Z =$ 0. In fact, from the very first turns, the difference becomes already significant with a confidence of 99%.

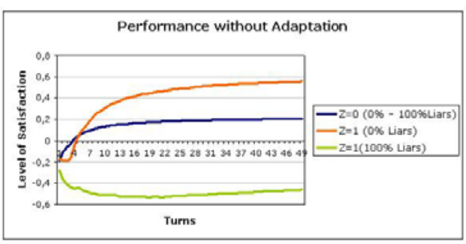

Fig. 5. Level of satisfaction obtained after

[4] We use the JASON platform [17], which offers to logic-based agents (prolog-like) a multiagent communication layer. The source code, together with the exact parameters and the set of desires used to run the experiments can be found at
`http://www.iiia.csic.es/~ipinyol/sourceABSS09.zip`

These results are quite obvious. Since image information is only created from direct experiences (1 at each turn) and reputation information through communicated reputation (5 at each turn) if the communicated information corresponds to the reality and the agent believes what circulates in the society ($Z = 1$) the buyer should discover faster which are the sellers that accomplish her objectives. As well, if reputation information if mostly false, and the agent believes it, for a long time the buyer would not be able to fulfill her objectives.

5.3 Dynamic Adaptation Experiments

The main idea behind the updating of Z is that in scenarios where mostly false reputation information circulates Z should be tend to 0. On the contrary, scenarios where reputation information is mostly accurate, Z should tend to 1. In this very preliminary paper, we study the effects of an adaptation strategy in the same situations tested in the previous extreme experiments.

The strategy is very simple, but effective. As described in the theoretical part of the paper, metarules can be updated from the beliefs that the agent hold. In our case, we theorize that a good metaprocess for updating Z is aggregate the differences for each agent and role of the image and reputation information hold in the Repage system. So, if most of the image information coincide with reputation information (about the same agent/role), the Z value should increase from the current value (in certain proportion). On the contrary, it should decrease. This algorithm contains the parameter *Increment*, which could be also considered as another degree of freedom.for the sake of simplicity we consider it as a constant value.

Figure 6 shows the performance obtained in both scenarios. It can be observed how the final performance tends to the theoretical optimum in each situation. In both scenarios there is no statistical significant difference between the performance and the theoretical optimum, with p_values higher than 0.2 with most of the points of the graph.

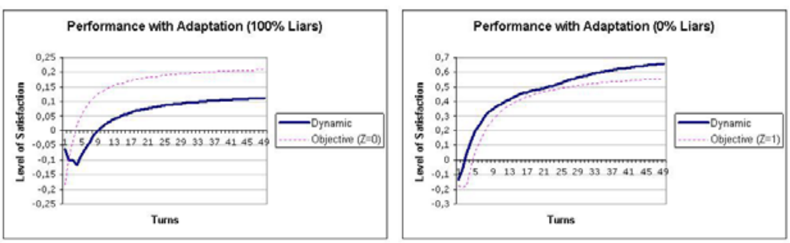

Fig. 6. Level of satisfaction obtained with agents using adaptation in a scenario with 100% of liars (left) and 0% of liars (right). Dot line represents the theoretical best possible performance.

The adaptation process can be clearly observed with the performance of a single execution. Figure 7 shows a typical pattern (usually the period where the level of satisfaction is so low is much shorted. For this reason the final average of 10 executions does not show it) in which after a while, the agent is able increase her level of goal achievement.

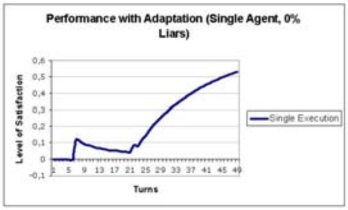

Fig. 7. Performance of a single agent with adaptation in a scenario without liars

6 Conclusions and Future Work

After reading the paper it should be clear the importance of allowing degrees of freedom in the reasoning processes of autonomous agents. As mentioned in the cognitive theory presented by Castelfranchi and Paglieri in [4], cognitive agents are aware of the *reasons* from which certain information if believed, and because of that, they are able to reason about how they reason, and change it if necessary. Thus, we strongly believe that real autonomous agents should be designed taking into account certain degree of granularity. Cognitive designs should be aware that the path that an agent follows to arrive at certain conclusion is as important as the conclusion itself. Therefore, ways to reason about such paths and the capability to modify them should be taken into account, not only for the agent itself, but also for possible explanations to other agents, like in argumentation.

We also encourage the use of logical approaches in the design of cognitive systems. The advantage of such systems is that with a finite set of rules a whole deduction tree can be created, implicitly providing supporting sets. This big advantage has an important counter effect: the static nature of the axiomatization. At a metareasoning level though, similar to belief revision process, certain set of axioms (those which define typology of agents, not that structurally guarantee certain logical properties) can me updated, changing then the whole reasoning tree. We show a possible method to do it in this paper. However, this needs a deeper study in the future.

Getting into the concrete scenario that we faced, it should also be clear that reputation and image information can totally participate in metareasoning processes. We proposed a method in which bridge rules and axioms can be specified as metarules following the idea of degrees of freedom introduced in [15]. We let for future work the proper formalities of the proposed design method. Also, regarding the actual design of BDI+Repage model we plan to provide alternative metaprocesses to update the trust predicate and the relation with the other metarules $M(A_I)$ and $M(A_R)$.

Acknowledgments

This work was supported by the projects AEI (TIN2006-15662-C02-01), AT (CONSOLIDER CSD2007-0022, INGENIO 2010), LiquidPub (STREP FP7-213360), RepBDI (Intramural 200850I136) and by the Generalitat de Catalunya under the grant 2005-SGR-00093.

References

1. Luck, M., McBurney, P., Shehory, O., Willmott, S.: Agent Technology: Computing as Interaction (A Roadmap for Agent Based Computing). AgentLink (2005)
2. Sabater-Mir, J., Paolucci, M., Conte, R.: Repage: Reputation and image among limited autonomous partners. JASSS 9(2) (2006)
3. Pinyol, I., Sabater-Mir, J.: Pragmatic-strategic reputation-based decisions in BDI agents. In: Proc. of the AAMAS 2009, Budapest, Hungary, pp. 1001–1008 (2009)
4. Castelfranchi, C., Paglieri, F.: The role of beliefs in goal dynamics: Prolegomena to a constructive theory of intentions. Synthese 155, 237–263 (2007)
5. Paglieri, F.: Belief dynamics: From formal models to cognitive architectures, and back again. PhD dissertation (2006)
6. Prakken, H.: Formal systems for persuasion dialogue. Knowl. Eng. Rev. 21(2), 163–188 (2006)
7. Conte, R., Paolucci, M.: Reputation in artificial societies: Social beliefs for social order. Kluwer Academic Publishers, Dordrecht (2002)
8. Giunchiglia, F., Serafini, L.: Multilanguage hierarchical logic (or: How we can do without modal logics). Journal of AI 65, 29–70 (1994)
9. Parsons, S., Sierra, C., Jennings, N.: Agents that reason and negotiate by arguing. Journal of Logic and Computation 8(3), 261–292 (1998)
10. Casali, A., Godo, L., Sierra, C.: Graded BDI models for agent architectures. In: Leite, J., Torroni, P. (eds.) CLIMA 2004. LNCS (LNAI), vol. 3487, pp. 126–143. Springer, Heidelberg (2005)
11. Pinyol, I., Sabater-Mir, J., Dellunde, P.: Probabilistic dynamic belief logic for image and reputation. In: Proc. of the CCIA 2008, Empuries, Spain (2008)
12. Fagin, R., Halpern, J.: Reasoning about knowledge and probability. J. ACM 41(2), 340–367 (1994)
13. Demolombe, R.: To trust information sources: a proposal for a modal logical framework. Trust and deception in virtual societies, 111–124 (2001)
14. Rao, A.S., Georgeff, M.P.: Modeling rational agents within a BDI-architecture. In: Allen, J., Fikes, R., Sandewall, E. (eds.) Proc. of KR 1991, pp. 473–484. Morgan Kaufmann Publishers Inc., San Francisco (1991)
15. Artikis, A.: Dynamic protocols for open agent systems. In: Proc. of the AAMAS 2009, Budapest, Hungary, pp. 97–104 (2009)
16. Pinyol, I., Paolucci, M., Sabater-Mir, J., Conte, R.: Beyond accuracy. Reputation for partner selection with lies and retaliation. In: Antunes, L., Paolucci, M., Norling, E. (eds.) MABS 2007. LNCS (LNAI), vol. 5003, pp. 128–140. Springer, Heidelberg (2008)
17. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak Using Jason. John Wiley and Sons, Ltd., Chichester (2007)