on Using Knowledge in its Context, Chambery, France, 1993, pp. 39–49. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.

[5] F. Giunchiglia and C. Ghidini. A Local Models Semantics for Propositional Attitudes. In *Proceedings of the 1st International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97)*, pages 363–372, Rio de Jeneiro, Brazil, 1997. Also IRST-Technical Report 9607-12, IRST, Trento, Italy.

[6] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29–70, 1994. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.

[7] F. Giunchiglia, L. Serafini, E. Giunchiglia, and M. Frixione. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548–554, Chambery, France, 1993. Also IRST-Technical Report 9206-03, IRST, Trento, Italy.

[8] J . Hintikka. *Knowledge and Belief.* Cornell University Press, Ithaca, NY, 1962.

[9] A. P. Martinich. *The philosophy of language.* Oxford University Press, 1985.

[10] D. Prawitz. *Natural Deduction - A proof theoretical study.* Almquist and Wiksell, Stockholm, 1965.

[11] B. Russell. On denoting. *Mind*, XIV, 1905.

[12] L. Serafini and C. Ghidini. Local Models Semantics for Information Integration. Technical Report 9702-04, IRST, Trento, Italy, February 1997. To appear in AAAI FAll 1997 symposium on context in KR and NL. Extended abstract to appear in the proceedings of the 15th IJCAI.

6 Conclusions

The simple case study we have formalized in this paper is clearly an instance of a much more general phenomenon, potentially concerning every belief report where some reference is made to some individual via a definite description. We argued that a proper treatment of this phenomenon does not involve a formalization of the *de re/de dicto* distinction, but rather a formalization of two possible attitudes on the reporter side towards another agent's beliefs. In the opaque reading, the reporter is assumed to just *quote* another agent's beliefs, whereas in the transparent reading the reporter is assumed to *translate* another agent's beliefs in his/her own words. As a consequence, the "same" definite description in the first case is meant to be used in the sense of the agent whose belief is reported, whereas in the second case is meant to be used in the sense of the reporter. This phenomenon is given a general (and quite natural) formalization by exploiting the multi-language features of MC systems.

The MC approach to belief contexts presupposes a deep change in the attitude toward the problem of formalizing beliefs (and reasoning about beliefs). Indeed, it requires to take seriously into account the fact that logical languages can be used to *ascribe beliefs to agents*, and not just to describe their beliefs from the point of view of an external (and – in general – omniscient) observer. The properties of locality and compatibility are motivated by this change of perspective, since they allow us not to hardwire in the logic some assumption which are implausible for an adequate logic of beliefs. For instance, we do not assume that the "same" linguistic token means the same thing for different agents, or that an agent can always understands what another agent meant to communicate. Misunderstandings happen all often in every day conversation, and an adequate formalization must explain how they are possible. We believe that our notion of views as contexts are a contribution toward such an explanation, and that it can throw a new light on some of the most difficult puzzles in formalizing beliefs (*e.g.* omniscience, failure of substitutivity, limitations of reasoning capabilities).

References

[1] A. Bonomi. *Eventi mentali.* Il Saggiatore, 1983.

[2] A. Cimatti and L. Serafini. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In M. Wooldridge and N. R. Jennings, editors, *Intelligent Agents: Proceedings of 1994 Workshop on Agent Theories, Architectures, and Languages*, number 890 in Lecture Notes in Computer Science, pages 71–85. Springer Verlag, 1995. Also IRST-Technical Report 9312-01, IRST, Trento, Italy.

[3] E. Giunchiglia and F. Giunchiglia. Ideal and Real Belief about Belief. In *Practical Reasoning, International Conference on Formal and Applied Practical Reasoning, FAPR'96*, number 1085 in Lecture Notes in Artificial Intelligence, pages 261–275. Springer Verlag, 1996.

[4] F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993. Short version in Proceedings IJCAI'93 Workshop

$$
\dfrac{\dfrac{\dfrac{\dfrac{\boxed{Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”})}_{\epsilon}}{\boxed{Bel_A(\text{“}Corr(\iota P)\text{”})}_{B}}\; R^{B}_{dn}}{\boxed{Corr(\iota P)}_{BA}}\; R^{BA}_{dn}}{\boxed{\iota P = C \supset Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})}}\; Tr^{BA}_{=} \qquad \dfrac{\boxed{Bel_B(\text{“}\iota P = C\text{”})}_{\epsilon}}{\iota P = C}\; R^{B}_{dn}}{\boxed{\dfrac{\iota P = C \supset Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”}) \qquad \iota P = C}{Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})} \supset E}_{B}}
$$

$$
\dfrac{\boxed{\dfrac{\iota P = C \supset Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”}) \qquad \iota P = C}{Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})} \supset E}_{B}}{\boxed{Bel_B(\text{“}Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})\text{”})}_{\epsilon}}\; R^{B}_{up} \tag{11}
$$

The last formula corresponds to a possible transparent reading (performed by ϵ) of the utterance $Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”})$.

5.2 Modeling transparency with local models semantics

Definition 5.2 (Transparent model) *An opaque model* $\mathbf{C}$ *is a* transparent model *if and only if for every* $\mathbf{c} \in \mathbf{C}$ *and for every* $i, j \in I$*, with* $i \prec j$*:*

$$
\phi \in V^{\uparrow}(\mathbf{c}_i) \quad implies \quad (\tau = \sigma) \supset Bel^{\mathrm{T}}_{s_k}(\text{“}\phi[\tau/\sigma]\text{”}) \in \Theta(\mathbf{c}_i) \tag{12}
$$

$$
\phi \in V^{\uparrow}(\mathbf{c}_i) \quad implies \quad (\psi \equiv \theta) \supset Bel^{\mathrm{T}}_{s_k}(\text{“}\phi[\psi/\theta]\text{”}) \in \Theta(\mathbf{c}_i) \tag{13}
$$

(12) means that if ϕ is believed by j, *i.e.* the mental image that i has of s_k, then i is able to perform a transparent reading of the formula $Bel_{s_k}(\text{“}\phi\text{”})$ on the basis of the equality $\tau = \sigma$. (13) has an analogous meaning using wffs instead of terms.

We show how (12) allows B to translate A's belief $Corr(\iota P)$ using its own belief $\iota P = C$, instead of using A's belief $\iota P = M$. Since every transparent model is an opaque model, we obtain (7) from (5) as in section 4. Applying constraint (12) between **B** and **BA**, with $\phi = Corr(\iota P)$, we obtain that every $\mathbf{c}_B$ satisfies

$$
\boxed{B \qquad (\iota P = C) \supset Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})} \tag{14}
$$

From the fact that every $\mathbf{c}_B$ satisfies $\iota P = C$ and from soundness of modus ponens in local models semantics [5], we obtain

$$
\boxed{B \qquad Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})} \tag{15}
$$

Finally, applying constraint (3) as from (8) to (10), we obtain

$$
\boxed{\epsilon \qquad Bel_B(\text{“}Bel^{\mathrm{T}}_{A}(\text{“}Corr(C)\text{”})\text{”})} \tag{16}
$$

which corresponds to the transparent reading of the utterance $Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”})$.

Those steps are the model-theoretical is the proof-theoretical counterpart of (11), where applications of constraint (12) correspond to applications of bridge rule $Tr^{is_k}_{=}$. The proof that MV^T systems are sound and complete with respect to the class of trasparent models can be constructed following the methodology showed in [5].

that is **C** satisfies an opaque reading of the utterance $Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”})$.

Those steps are the model-theoretical counterpart of (2), where applications of constraints (4) and (3) correspond to applications of bridge rules $R_{up}^{is_k}$ and $R_{dn}^{is_k}$ respectively. The proof that MV systems are sound and complete with respect to the class of opaque models is a straightforward generalization of the soundness and completeness theorem for the single agent case in [5].

5 Representing transparency

5.1 Formalizing transparency with MultiContext systems

Definition 5.1 (MV^T) *The formal system allowing both opaque and transparent readings is the MC system $MV^T = \langle \{C_i\}_{i \in I}, BR_I \rangle$ such that MV^T is defined as MV and also contains the following bridge rules:*

$$\frac{is_k : \phi}{i : (\tau = \sigma) \supset Bel^{\mathrm{T}}_{s_k}(\text{“}\phi[\tau/\sigma]\text{”})}\ Tr_{=}^{is_k} \qquad \frac{is_k : \phi}{i : (\psi \equiv \theta) \supset Bel^{\mathrm{T}}_{s_k}(\text{“}\phi[\psi/\theta]\text{”})}\ Tr_{\equiv}^{is_k}$$

where $Tr_{=}^{sj}$ and $Tr_{\equiv}^{sj}$ are applicable only if both their premisses and $\tau = \sigma$ [$\psi \equiv \theta$] are closed formulae, and no formula their premisses depend upon has index greater then i.

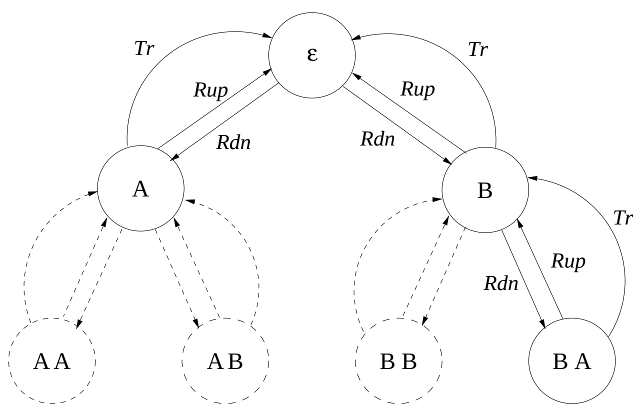

Figure 3: The MV^T system

Figure 3 depicts the MultiContext architecture of MV^T, where arrows corresponding to Tr bridge rules are added. Coming back to our example, by starting from Ω_ϵ, both proof (2) and the following proof of (16) can be performed in MV^T:

Definition 4.2 (Opaque model) *A model* $\mathbf{C}$ *for an MV language is an* opaque model *if and only if for every* $\mathbf{c} \in \mathbf{C}$ *and for every* $i, j \in I$*, with* $i \prec j$*:*

$$Bel_{s_k}(\text{“}V^{\uparrow}(\mathbf{c}_i)\text{”}) \subseteq \Theta(\mathbf{c}_i) \tag{3}$$

$$Bel^{-1}_{s_k}(\text{“}\Theta(\mathbf{c}_i)\text{”}) \subseteq \Theta(\mathbf{c}_j) \tag{4}$$

Condition (3) says that i is a complete observer. Indeed $\mathbf{c}_i$ satisfies all the formulae $Bel_{s_k}(\text{“}\phi\text{”})$ such that ϕ is satisfied by all the models $\mathbf{c}'_j$ related (via compatibility relation) with (a subset of) $\mathbf{c}_i$, *i.e.* such that ϕ is believed by j, the mental image that i has of s_k. Condition (4) says that i is a correct observer because if $Bel_{s_k}(\text{“}\phi\text{”})$ is satisfied by $\mathbf{c}_i$ then ϕ is satisfied by $\mathbf{c}_j$. Notice that (4) is "weaker"[8] than (3). This fact reflects the difference among $R^{is_k}_{dn}$ and $R^{is_k}_{up}$, *i.e.* te fact that $R^{is_k}_{up}$ is a "restricted" bridge rule, and means that every i's belief $Bel_{s_k}(\text{“}\phi\text{”})$ is correct even if it is a supposition of i (and not a theorem).

Coming back to our example, we show how ϵ uses (3) and (4) to perform opaque reading. The initial knowledge of ϵ is

$$\boxed{\begin{array}{ll} \epsilon & Bel_B(\text{“}\iota P = C\text{”}) \\ & Bel_B(\text{“}Bel_A(\text{“}\iota P = M\text{”})\text{”}) \\ & Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”}) \end{array}} \tag{5}$$

that is every $\mathbf{c}_\epsilon$ in $\mathbf{C}$ satisfies (5). From constraint (4) between the pair ϵ, $\mathbf{B}$ we obtain that every $\mathbf{c_B}$ satisfies

$$\boxed{\begin{array}{ll} B & \iota P = C \\ & Bel_A(\text{“}\iota P = M\text{”}) \\ & Bel_A(\text{“}Corr(\iota P)\text{”}) \end{array}} \tag{6}$$

From constraint (4) between the pair $\mathbf{B}$, $\mathbf{BA}$ we obtain that every $\mathbf{c_{BA}}$ satisfies

$$\boxed{\begin{array}{ll} B & \iota P = C \\ & Bel_A(\text{“}\iota P = M\text{”}) \\ & Bel_A(\text{“}Corr(\iota P)\text{”}) \end{array}} \tag{7}$$

From the definition of local semantics as first order semantics, every $\mathbf{c_{BA}}$ satisfies also the following consequence of the equality axioms:

$$\boxed{\begin{array}{ll} BA & Corr(M) \end{array}} \tag{8}$$

Finally, by constraint (3) between the pair $\mathbf{B}$, $\mathbf{BA}$ every $\mathbf{c_B}$ satisfies

$$\boxed{\begin{array}{ll} B & Bel_A(\text{“}Corr(M)\text{”}) \end{array}} \tag{9}$$

and by constraint (3) between the pair ϵ, $\mathbf{B}$, every $\mathbf{c}_\epsilon$ satisfies

$$\boxed{\begin{array}{ll} \epsilon & Bel_B(\text{“}Bel_A(\text{“}Corr(M)\text{”})\text{”}) \end{array}} \tag{10}$$

[8] A condition for correctness, symmetric to (3), is $Bel^{-1}_{s_k}(\text{“}V^{\downarrow}(\mathbf{c}_i)'') \subseteq \Theta(\mathbf{c}_i)$ where $Bel^{-1}_{s_k}(\text{“}\Gamma'')$ and $V^{\downarrow}(\mathbf{c}_i)$ correspond to $Bel_{s_k}(\text{“}\Gamma\text{”})$ and $V^{\uparrow}(\mathbf{c}_i)$. For more details see [5].

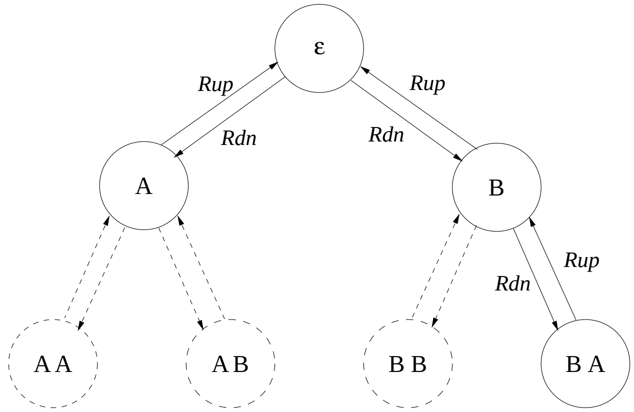

Figure 2: The MV system

$$
\cfrac{
 \cfrac{
  \cfrac{\cfrac{\boxed{Bel_B(``Bel_A(``\iota P = M")")}_{\epsilon}}{\boxed{Bel_A(``\iota P = M")}_{B}}\,R^{B}_{dn}}{\iota P = M}\,R^{BA}_{dn}
  \qquad
  \cfrac{\cfrac{\boxed{Bel_B(``Bel_A(``Corr(\iota P)")")}_{\epsilon}}{\boxed{Bel_A(``Corr(\iota P)")}_{B}}\,R^{B}_{dn}}{Corr(\iota P)}\,R^{BA}_{dn}
 }{Corr(M)}\,Subst \quad {}_{BA}
}{
 \cfrac{\boxed{Bel_A(``Corr(M)")}_{B}}{\boxed{Bel_B(``Bel_A(``Corr(M)")")}_{\epsilon}}\,R^{B}_{up}
}\,R^{BA}_{up} \tag{2}
$$

The proof is reported in a Natural Deduction-like style. Deduction local to a context i is surrounded by a box labeled i. Bridge rules connect local deductions performed in different contexts, their premisses and conclusions being the last formulae of the boxes above and the first formulae in the boxes below them, respectively. The last formula of the deduction corresponds to the opaque reading of the utterance $Bel_B(``Bel_A(``Corr(\iota P)")")$. The rule $Subst \in \Delta_i$ is the standard rule of substitutivity and allows for the MultiContextual version of the RPS principle.

4.2 Modeling opacity with local models semantics

In order to define the model for MV we introduce the following notation for every $i, j \in I$, with $i \prec j$ and $j = is_k$.

- Let Γ be a set of L_j-formulae. $Bel_{s_k}(``\Gamma")$ denotes the set $\{Bel_{s_k}(``\phi") \in L_i \mid \phi \in \Gamma\}$;
- let Γ be a set of L_i-formulae. $Bel^{-1}_{s_k}(``\Gamma")$ denotes the set $\{\phi \in L_j \mid Bel_{s_k}(``\phi") \in \Gamma\}$.

Let $\mathbf{C}$ be a model. $\Theta(\mathbf{c}_i)$ denotes the set $\{\phi \in L_i \mid \phi$ is closed and for every $m \in \mathbf{c}_i$, $m \models_{l_c} \phi\}$ and $V^{\uparrow}(\mathbf{c}_i)$ denotes the set $\{\phi \in L_j \mid \phi$ is closed and for every $\mathbf{c}' \in \mathbf{C}$ if $\mathbf{c}'_i \subseteq \mathbf{c}_i$ then $\phi \in \Theta(\mathbf{c}'_j)\}$.

(i) $L \subseteq L_i$ for each $i \in I$;

(ii) if $\phi \in L_j$ then the term "ϕ" belongs to every L_i with $i \prec j$ (where "ϕ" is the "name" of the formula ϕ);

(iii) if $i \prec j = is_k$ then Bel_{s_k} and $Bel^{\mathrm{T}}_{s_k}$ are unary predicate symbols of L_i (where Bel_{s_k} is the predicate used by i to mean "s_k told me that [...]" and $Bel^{\mathrm{T}}_{s_k}$ is the predicate used by i to mean "I would express the s_k beliefs as [...]")

Notationally, for every $i \in I$, we write $i : \phi$ to mean ϕ and that ϕ is a formula of L_i. We say that $i : \phi$ is a formula and that ϕ is a L_i-formula.

4.1 Formalizing opacity with MultiContext systems

Definition 4.1 (MV) *The formal system allowing the opaque reading is the MC system* $MV = \langle \{C_i\}_{i \in I}, BR_I \rangle$ *such that*

- I_0 *and* I *are the sets defined above, along with the binary relation* $\prec$*;*
- *for every* $C_i \in \{C_i\}_{i \in I}$*:*
  - L_i *is the language of index* i *in the class of MV languages defined above;*
  - *if* $i \neq \epsilon$*, then* $\Omega_i = \emptyset$*;*
    if $i = \epsilon$*, then* Ω_i *contains the* L_ϵ *formulae* $Bel_B(\text{“}\iota P = C\text{”})$, $Bel_B(\text{“}Bel_A(\text{“}\iota P = M\text{”})\text{”})$, $Bel_B(\text{“}Bel_A(\text{“}Corr(\iota P)\text{”})\text{”})$*;*
  - Δ_i *is the set of inference rules of Natural Deduction [10] plus a sound and complete set of rules for equality;*
- *the set* BR_I*, which presents the compatibility relation of definition 4.2, contains, for any index* $is_k \in I$ *(with* $s_k \in I_0$*), the set* BR_{is_k} *of bridge rules of the form:*

$$\frac{is_k : \phi}{i : Bel_{s_k}(\text{“}\phi\text{”})}\ R_{up}^{is_k} \qquad\qquad \frac{i : Bel_{s_k}(\text{“}\phi\text{”})}{is_k : \phi}\ R_{dn}^{is_k}$$

Restriction: $R_{up}^{is_k}$ *is applicable only if the* L_{is_k}*-formula* ϕ *is a closed formula and no formula the premiss* $is_k : \phi$ *depends upon has index greater*[7] *then* i*.*

Figure 2 depicts the MultiContext architecture of MV: circles correspond to contexts and labeled arrows to bridge rules among contexts. Context ϵ represents the view of the program itself, while each context is_k represents the view of agent s_k from the point of view i. Let us consider our example. The following is a MultiContextual proof of the formula (10) in MV:

[7] An index j is "greater then" index i if and only if i is a subsequence of j.

instance, the bridge rule

$$\frac{C_1 : A_1}{C_2 : A_2} \tag{1}$$

states the derivability of the formula A_2 in context C_2 (written $C_2 : A_2$) from the derivability of A_1 in context C_1 ($C_1 : A_1$). By abuse of notation, in the rest of the section we will often refer to contexts with their indexes (*i.e.* we will refer to context C_i as the context i).

In [7; 3; 2], it is shown how belief contexts can be formalized using MC systems. Our work is built on top of this work. The basic ideas are the following. We imagine that the computer program ϵ is able to handle a set I of views representing the collections of beliefs that it ascribes to Mr. A and Mr. B (the set of views of figure 1). Formally, each view i is thought of as a context C_i, with its own language, set of axioms and local inference rules (locality). Compatibility among different views is modelled by defining a suitable set of bridge rules.

Semantically, locality can be presented by modelling each view i as a set of models of the context C_i. In this way, every context C_i has its own (local) semantics, since the interpretation of the language L_i and of the set of axioms Ω_i is local to the context C_i. Notationally we write $\mathbf{c}_i$ to mean a set of models for C_i. Compatibility can be presented by imposing constraints on the combinations of the sets of models of different context. This is what we call a *compatibility relation.* For instance, the compatibility relation

$$\{\langle \mathbf{c}_1, \mathbf{c}_2 \rangle \mid \text{if } \mathbf{c}_1 \text{ satisfies } A_1, \text{ then } \mathbf{c}_2 \text{ satisfies } A_2\}$$

states that contexts C_1 and C_2 are compatible iff every time A_1 is true in C_1, then A_2 is true in C_2 and corresponds to the bridge rule (1). A model for an MC-system $\langle \{C_i\}_{i \in I}, BR_I \rangle$ is a compatibility relation $\mathbf{C}$ defined over sets of (local) models of $\{C_i\}_{i \in I}$ and modelling the same "links" expressed by BR_I.

In the next two sections we show how to use MC systems to formalize both the opaque and the transparent reading of the belief report "Mr. A believes that the president of the local football team is a corruptor" (the example of section 2). In order to do this, (i) we define the set of contexts (each with its set of local models) representing the views of the example; (ii) we define the bridge rules (compatibility relations) representing the fact that the agent B may performs an opaque or a transparent reading of the belief of A.

4 Representing opacity

In order to present the views system of our case study we first introduce (following [2]) the class of languages for the computer's views. Let L be a first order language containing two constants C and M (Mr. C and Mr. M respectively), a unary predicate $Corr$ (to be a corruptor), and a definite description ιP[6] (the president of the local football team). Let $I_0 = \{A, B\}$ be the set of the agents' names in the scenario, and $I = I_0^*$ the set of finite (possible empty) sequences of elements in I_0. Let $\epsilon \in I$ denote the empty sequence. We may define a partial order relation $\prec$ over I such that for every $i, j \in I$, $i \prec j$ iff j is of the form is_k for some $s_k \in I_0$. An *MV language* (MV stands for Multi View) is defined as a class of first order languages $\{L_i\}_{i \in I}$, where

[6] We use the notation of [11] for definite descriptions.

the sentence "the president of the local football team is a corruptor": for Mr. A, it is true iff Mr. M is a corruptor; for Mr. B, it is true iff Mr. C is. Since the definite description 'the president of the local football team' is such that it can refer only to a single person, the union of the truth conditions apparently result in a contradiction (unless we can prove that Mr. M is the same person as Mr. C). Our approach is based on an idea proposed in [4]: there is not a single language that the program uses to describe its views about other agents' beliefs; instead, *a distinct language is associated with each view, and the interpretation of such a language is local to the view it is associated with.* In this way, we can distinguish the sentence "the president of the local football team is a corruptor" in **B** and the sentence "the president of the local football team is a corruptor" in **BA**. Even though sintactically they "look the same", they do not denote the same proposition, and therefore are satisfied by two different sets of models. This property is called *locality*, and is a key point in the formalization of the intrinsic ambiguity of belief reports.

The contents of different views are obviously related. For instance, it is reasonable to assume that there is a relation between the fact that the program puts into **B** the belief that Mr. A believes that the president of the local football team is a corruptor, and into **BA** the belief that the president of the local football team is a corruptor. Any relation between sets of facts belonging to different views is called a relation of *compatibility*. A very intuitive relation of compatibility between views is the following: a sentence of the form "*Mr. A believes that* ϕ" belongs to **B** iff ϕ belongs to **BA** (where *believes* is a belief operator). In this case, we say that B is a correct and complete observer [3]. But it is very important to realize that this is not the only kind of observers. Indeed, for the transparent reading we need a different kind of observer, defined by the following compatibility relation: if **BA** contains a sentence ϕ and **B** contains an equality (equivalence) of the form $\alpha = \alpha'$ ($\alpha \equiv \alpha'$), then a sentence of the form "*A believes*^{T} *that* $\phi[\alpha/\alpha']$" belongs to **B**, (where *believes*^{T} is a belief operator distinct from *believes*^{O}). In the remain of the paper, we show – both model-theoretically and proof-theoretically – that these two compatibility relations are basically the relations between views that we need in order to model opacity and transparency in belief reports.

3 Belief contexts in MultiContext systems

In order to model the properties of locality and compatibility as discussed in section 2, we formalize belief reports in the framework of *MultiContext systems* (MC systems). In this section we review only those aspects of MC systems that are needed in order ot present our formalization of belief reports. The interested reader can refer to the bibliography for a more complete presentation[5].

Formally, given a set I of indexes, an MC-system is defined as a pair $\langle \{C_i\}_{i \in I}, BR_I \rangle$, where $\{C_i\}_{i \in I}$ is a collection of *contexts* and BR_I is a set of *bridge rules*. Each context is defined as an axiomatic formal system, *i.e.* a triple $\langle L_i, \Omega_i, \Delta_i \rangle$, where L_i is the language of C_i, $\Omega_i \subseteq L_i$ is the set of axioms of C_i and Δ_i is the set of (local) inference rules of C_i. A bridge rule is defined as an inference rule with premisses and conclusion in different contexts. For

[5]MC systems were formally defined in [4] and in [6] (with the name of *MultiLanguage systems*). The semantics of MC systems was presented in [5; 12].

believes that the president of the local football team is Mr. M, the content of the reported belief is that Mr. A believes that Mr. M is a corruptor (opaque reading). If the program ascribes to Mr. B the "translation" attitude, it will reason as follows: since Mr. B believes that the president of the local football team is Mr. C, and since the description 'the president of the local football team' is to be read in Mr. B's sense, the reported belief is that Mr. C is a corruptor (transparent reading).

In order to formalize these intuitions, we introduce the notion of *view*. A view is a representation of a collection of beliefs that a reasoner (in our example, the program) ascribes to an agent (including itself) under a given perspective. Possible perspectives are: the beliefs that the program ascribes to itself (*e.g.* that Mr. B believes that Mr. A believes that the president of the local football team is a corruptor); the beliefs that the program ascribes to Mr. B (*e.g.* that Mr. A believes that the president of the local football team is a corruptor); the beliefs that the program ascribes to Mr. B about Mr. A (*e.g.* that the president of the local football team is a corruptor). As a convention, we use the Greek letter ϵ for the view containing the beliefs that the program ascribe to itself, and bold letters for labelling views. For instance $\epsilon\mathbf{B}$ is the program's view of Mr. B's beliefs, and $\epsilon\mathbf{BA}$ the program's view of the beliefs that Mr. B ascribes to Mr. A[3]. The views that the program can build can be organized in a structure like that presented in figure 1[4]. Each circle represents a view; some circles are dashed because we will not use them in the formalization of the case study.

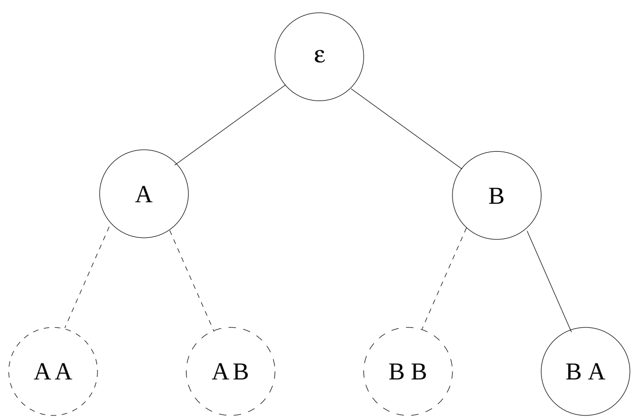

Figure 1: The structure of views

As usual, the propositional content of a belief is represented as a sentence of some language. In most traditional approaches, this language is unique and any belief is represented as a sentence in this language. Of course, a sentence can be believed by an agent and not believed by another, but it denotes one and the same proposition in any possible belief context, namely has the same truth conditions in every belief context. However, different agents may attribute different truth conditions to the "same" sentence. This is the case of

[3]Since in our example no confusion can arise, from now on we will omit the prefix ϵ.

[4]This structure can be easily generalized to n agents. For a more detailed description of the structure, a good reference is [2], where views are used to solve the a well-known puzzle involving reasoning about belief and ignorance, namely the Three-Wise-Men problem.

- Mr. A's belief is referred to Mr. C (since it is B that is speaking), and thus the program concludes that Mr. A believes that C is a corruptor (even though Mr. A does not identify Mr. C as the president of the local football team).

The first is an *opaque reading* of the belief report, and requires an application of RPS: since the two beliefs that the president of the local football team is a corruptor and that the president of the local football team is Mr. M occur in the same belief context (*i.e.* Mr. A's beliefs), the program can apply RPS and conclude that Mr. A believes that Mr. M is a corruptor. The second is a *transparent reading*, and requires a stronger principle than RPS. Indeed, the program must be able to reason about the fact that Mr. B might have used the description 'the president of the local football team' to mean Mr. C (the person Mr. B himself believes to be the president) and not to mean Mr. M (the person Mr. A believes is the president). The question is whether the transparent reading is (intuitively) legitimate. The first reaction is that it is somewhat innatural. Since Mr. A would not agree with the conclusion that Mr. C is a corruptor, the transparent reading is not "faithful" to Mr. A's belief and therefore – one might conclude – incorrect. However, as it is convincingly argued in [1], this conclusion is based on the idea that only Mr. A's beliefs are relevant, and completely disregards the rôle of the reporter. Let us imagine, for example, the following scenario: Mr. B is told by Mr. A: "Mr. C is a corruptor"; since Mr. B believes that Mr. C is the president of the local football team, reporting Mr. A's belief using the description 'the president of the local football team' seems perfectly acceptable (even though Mr. B knows that Mr. A would not agree with this report).

Traditionally, the two readings are explained with the *de re/de dicto* distinction (*e.g.* in [8]): the opaque reading is associated with a *de dicto* belief and the transparent reading with a *de re* belief. However, this does not work in the proposed example. Let us consider another version of the story. Mr. B is told by Mr. A: "The president of the local football team is a corruptor", and Mr. B has good reasons to think that Mr. A is not really referring to Mr. M (perhaps Mr. A is reporting a radio news on the president of the local football team). So, even though Mr. B knows that for Mr. A this entails the false belief that Mr. M is a corruptor, Mr. B can report Mr. A's belief, being sure that the program will understand the description 'the president of the local football team' in the right way (namely, as referring to Mr. C). In this case Mr. A has a *de dicto* belief about the president of the local football team, but Mr. B's report is based on a transparent reading. So the *de re/de dicto* distinction is orthogonal to the opacity/transparency problem.

Our analysis emphasizes the conceptual spaces of all agents involved in the scenario, in particular the conceptual space of the reporter (in the example, Mr. B). Mr. B can assume two attitudes towards Mr. A's reported belief. The first is to *quote* Mr. A's words (something like: He told me that "[...]"). The second is to *translate* the content of Mr. A's belief in his own words (something like: I would express his belief as [...]). If the program is to reason on belief reports, it must be able to ascribe both attitudes to the reporter. In particular, ascribing the first attitude will result in an opaque reading of a belief report and ascribing the second will result in a transparent reading. So let us look again at the example from this perspective. Mr. B tells the program that Mr. A believes that the president of the local football team is a corruptor. If the program ascribes to Mr. B the "quotation" attitude, it will reason as follows: since Mr. B is reporting the exact words of Mr. A, and since Mr. A

Mr. A believes that the president of the local football team is a corruptor
The president of the local football team is Mr. C

Mr. A believes that Mr. C is a corruptor

It is easy to see that the conclusion is not valid. Indeed, even though Mr. C is the president of the local football team, Mr. A could be not aware of this fact. Given the two premisses, we have no means to conclude whether Mr. A believes or not that the president of the local football team is Mr. C. Hence it is possible to think of a model in which the conclusion is false. Only a *restricted principle of substitutivity* (RPS) seems to hold: whenever two terms (two sentences) that occur within the scope of the same belief operator have the same meaning, they can be substituted one for the other in any sentence that occurs in the same context *salva veritate*. More formally, let Φ be a sentence, α a term (sentence), and Bel a belief operator. If $Bel(\alpha = \alpha')$ ($Bel(\alpha \equiv \alpha')$), then $Bel(\Phi) \equiv Bel(\Phi[\alpha/\alpha'])$. In the example above, RPS cannot be applied unless the second premiss is replaced by the fact that Mr. A believes that the president of the local football team is Mr. C. In this case, the conclusion holds.

However, there are cases in which this weaker form of substitutivity is not sufficient in order to model common forms of reasoning about beliefs. A very common example is the way people report other people's beliefs. Following the analysis of [1], we show that many belief reports can be given two readings, called *opaque* and *transparent* respectively, and that the second requires the application of a form of substitutivity stronger than RPS. We argue that both readings are intuitively plausible, and therefore an adequate formalization of beliefs should allow us to model both of them, and not just to eliminate one. To this end, in section 2 we present and discuss in detail a motivating example, which we use also as a case study; then we present a multi-context framework for belief contexts (section 3) and we show that it allows to model both opaque (section 4) and transparent (section 5) readings.

2 Opacity and transparency in belief reports

The example we consider is a slight modification (and translation) of an example from [1]:

> You know that Mr. A believes that the president of the local football team is Mr. M and you know that Mr. B believes that the president is Mr. C. You know also that Mr. B knows that A believes that the president of the local football team is Mr. M. Actually, Mr. B is right, and you know that. Now, B tells you: "Mr. A believes that the president of the local football team is a corruptor". If you want to know whom the sentence is about (besides Mr. A), how will you interpret the sentenceΓ

Suppose that the problem is posed to a computer program. The program is a little puzzled, since the question has two possible answers:

- Mr. A's belief is referred to Mr. M (since Mr. A is the subject of the belief), and thus the program concludes that Mr. A believes that M is a corruptor;

A Multi Context Approach to Belief Report

Massimo Benerecetti Paolo Bouquet Chiara Ghidini
DISA – University of Trento, Via Inama 5, 38100 Trento, Italy.
IRST, Povo, 38100 Trento, Italy
bene@cs.unitn.it bouquet@cs.unitn.it ghidini@cs.unitn.it

Abstract

One of the most interesting puzzles in formalizing belief contexts is the fact that many belief reports can be given both an opaque and a transparent reading. A traditional explanation is that the two readings are related to the failure and success of the *principle of substitutivity* respectively, and this in turn is explained with the *de re/de dicto* distinction. We propose an alternative analysis, based on the idea that another agent's beliefs can just be quoted (preserving opaqueness) or translated into the reporter's language (allowing transparency). We show that MultiContext systems allow the formalization of these two phenomena at the same time, thanks to their multi-language feature.

1 Beliefs and substitutivity

A very important capability is reasoning about our own and other people's beliefs. We have beliefs about our beliefs; we learn from other people's beliefs; we accept (don't accept) other people's beliefs among our beliefs; we share our beliefs with other people; we report other people's beliefs; we infer new information by combining different people's beliefs; we make decisions based on conjectures about other people's beliefs; many communication conventions are based on what we believe other people believe; and so on.

Despite the pervasiveness of the notion of belief and the apparent easiness for humans to deal with it, representing beliefs and formalizing reasoning with and about beliefs has raised very difficult problems. From a logical point of view, a well-known problem is that a very general and intuitive logical principle, the *principle of substitutivity*, fails in its unrestricted version[1]. Intuitively, the *unrestricted principle of substitutivity* (UPS) states that whenever two terms (two sentences) have the same meaning[2], they can be substituted one for the other in any sentence *salva veritate*. More formally, let Φ be a sentence and α, α' two terms (sentences). If $\alpha = \alpha'$ ($\alpha \equiv \alpha'$), then $\Phi \equiv \Phi[\alpha/\alpha']$, where $\Phi[\alpha/\alpha']$ is the result of replacing some occurrences of α with α'. But UPS is not sound (validity preserving) when applied in belief contexts (namely within the scope of an operator of the form *X-believes-that* [...], where [...] stands for any sentence). Consider the following instance of UPS:

[1] This fact was first noticed by Gottlob Frege in [9].
[2] The word 'meaning' is used here informally.

IRST

ISTITUTO PER LA RICERCA
SCIENTIFICA E TECNOLOGICA

38050 Povo (Trento), Italy
Tel.: + 39 461 314575 · Fax: + 39 461 314591
e-mail: prdoc@itc.it · url: http://www.itc.it

A Multi Context Approach to Belief Report

Massimo Benerecetti, Paolo Bouquet, Chiara Ghidini

June 1997

Technical Report # 9706-04